



Add Interfaces for LAN Operational Mode, Release 12.2.2/12.2.3

Table of Contents

New and Changed Information	1
Interfaces	2
Adding Interfaces	8
Breakout	10
UnBreakout	10
Editing Interfaces	10
Editing Interfaces Associated with Links	11
Applying Normalize Action on Interfaces	12
Deleting Interfaces	12
Shutting Down and Bringing Up Interfaces	12
Viewing Interface Configuration	13
Rediscovering Interfaces	13
Viewing Interface History	13
Deploying Interface Configurations	14
Creating External Fabric Interfaces	14
Interface Groups	15
Creating an Interface Group	16
Adding and Removing Interfaces from an Interface Group	17
Adding Interfaces to an Interface Group	18
Removing Interfaces from an Interface Group	18
Attaching Networks to an Interface Group	18
Detaching a Network from an Interface Group	19
Deleting an Interface Group	20
Out-of-Band Switch Interface Configurations	21
Guidelines	21
Syncing up Switch Interface Configurations	22
Copyright	25

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
There were no major changes from the previous release.		

Interfaces

The Interfaces option displays all the interfaces that are discovered for the switch, Virtual Port Channels (vPCs), and intended interfaces missing on the device.

Invalid interface error appears on the following scenarios:

- Interface Mode 'routed' is invalid. Allowed mode is trunk & access.
- Access port which is already allocated to other network.
- Interface which is not available in the switch.

You can use the following functions:

- Create, deploy, view, edit and delete a port channel, vPC, Straight-through FEX, Active-Active FEX, loopback, and subinterface.



- The following features are unsupported for the brownfield migration of switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images:

- FEX on switches other than Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with X9500 line cards
- AA-FEX

For information about the platform support for FEX, refer to your platform and NX-OS documentation to check the feature compatibility.

- To edit interfaces associated with fabric links such as intra-fabric links and inter-fabric links, see [Editing Interfaces Associated with Links](#).
- The **flowcontrol** or **priority-flow-control** config is not supported for HIF ports or PO with HIF ports as members.
- When using the REST API for configurations, make sure to set consistent values for the primary fields and NV pairs fields. For example, a REST API post for a single port channel that has different values in certain fields, such as:

- **ifName:** Port-testing123
- **PO_ID:** Port-channel1000

will result in two interfaces being created rather than the intended single interface.

- Create tunnel interfaces for Cisco Cloud Services Router 1000v Series (Cisco CSR 1000v Series).
- Create breakout and unbreakout ports.
- Shut down and bring up interfaces.
- Rediscover ports and view interface configuration history.
- Apply host policies on interfaces and vPCs. For example, int_trunk_host, int_access_host, and so on.
- View interface information such as its admin status, operation status, reason, policy, speed, MTU, mode, VLANs, IP/Prefix, VRF, port channel, and the neighbor of the interface.



The **Neighbor** column provides details of connected switches that are discovered, intent links, and Virtual Machine Manager (VMM) connectivity.

The **Status** column displays the following status of an interface:

- Blue: Pending
- Green: In Sync/Success
- Red: Out-of-Sync/Failed
- Yellow: In Progress
- Grey: Unknown/NA
- If an interface is created out-of-band, you need to perform fabric resync or wait for Config Compliance polling before this interface can be deleted. Otherwise, Config Compliance does not generate the correct diff.

However, you cannot add or edit interfaces for ASR 9000 Series Routers and Arista switches.

You can filter and view information for any of the given fields (such as Device Name). The following table describes the buttons that appear on this page.



- Ensure that appropriate configurations are deployed on the Fabric before deploying from the Interfaces option, including proper vPC pair configurations. If you add or edit an interface before configurations are deployed on the Fabric, the configuration may fail on the device.
- Deploy any underlays including vPC Pairing in the fabric before deploying any configurations from the interface manager.

Field	Description
Create Interface	Allows you to add a logical interface such as a port channel, vPC, Straight-through FEX, Active-Active FEX, and loopback. For more information, see Adding Interfaces .
Create Subinterface	Allows you to add a logical subinterface.
Edit	Allows you to edit and change policies that are associated with an interface.  Access-admin user role cannot edit interfaces associated with link policy such as inter-fabric link or intra-fabric link for easy fabrics. The user role can edit interfaces for LAN classic and IPFM fabrics. If you select the int_monitor interface policy, you are responsible for the interface configuration and management. NDFC does not apply any interface configurations nor perform configuration compliance for interfaces that have the int_monitor policy.
Normalize	Allows you to apply the same interface policies and its parameters on selected interfaces in one shot. For more information, see Applying Normalize Action on Interfaces .

Field	Description
Multi-Attach	Allows you to attach multiple networks when you add an interface from an interface group in the Interfaces window. For more information, see Attaching Networks to an Interface Group .
Multi-Detach	Allows you to detach multiple networks when you remove an interface from an interface group in the Interfaces window. For more information, see Detaching a Network from an Interface Group .
Preview	Allows you to preview the interface configuration.
Deploy	Allows you to deploy or redeploy saved interface configurations.
No Shutdown	Allows you to enable an interface (no shutdown or admin up).
Shutdown	Allows you to shut down the interface.
Add to interface group	Allows you to add an interface to an interface group.
Remove from interface group	Allows you to remove an interface from an interface group.
Breakout	Allows you to <i>breakout</i> an interface.
Un-Breakout	Allows you to unbreakout interfaces that are in <i>breakout</i> state.
Rediscover	Allows you to rediscover or recalculate the compliance status on the selected interfaces.
Show Commands	Allows you to display the interface show commands. A show command requires show templates in the template library.
History	Allows you to display the interface deployment history details.

Field	Description
Import	Allows you to import the edited interfaces. The following are the limitations during importing the interfaces: - You are not allowed to import interfaces with the below policy templates: - All fabric templates with int_fabric or int_ipfm_fabric int_vpc_peer and int_vpc_leaf_tor_assoc int_freeform templates - You must update the mandatory fields fabric name, serial number, interface name, and policy name. - You are not allowed to import the interfaces with the interface name <i>nve</i> and <i>vlan</i> except <i>int_ipfm_vlan</i> policy. You can import the interface with <i>int_ipfm_vlan</i> policy. - The allowed MTU range for integer values is between 576 and 9216. - The allowed MTU string values is either <i>default</i> or <i>jumbo</i>. - The fabric name, serial number, and interface name must be unique. - You can only import a single policy type for an interface per .csv file. Importing a .csv file with multiple policy types is not allowed.  There is a server property to set the maximum number of rows that can be imported. By default, the property is 200 for import.
Export	Allows you to export the selected interfaces with multiple types of policies to a .csv file. While there is technically no limit on the number of interfaces to export, the number of interfaces included in each exported .csv file is limited to the number of rows that are displayed on the page. For example, if you select all of the interfaces in the window and you have 50 as the entry in the Rows per page field at the bottom of the window, then only the 50 interfaces displayed in this page are exported to the .csv file.
Delete	Allows you to delete a logical interface that is created from the Interfaces screen. An interface having a policy that is attached from an overlay and underlay cannot be deleted.

The following table describes the new user role access-admin operations support in the host facing port of **Interfaces** window from Cisco Nexus Dashboard Fabric Controller Release 11.5(1).

Operations	User Roles
	access-admin
Create Interface	Save, Preview, Deploy
Breakout	Blocked
Un-Breakout	Blocked

Edit	Save, Deploy
Delete	Save, Deploy
Shutdown	Save, Deploy
No Shutdown	Save, Deploy
Show commands	Clear Output, Execute
Rediscover	Supported
Deploy	Cancel, Deploy Config
Import	Supported
Export	Supported

You can disable deployments, or freeze, a fabric in Nexus Dashboard Fabric Controller as a network administrator. However, you cannot perform all actions when you freeze the fabric or if the fabric is in monitor mode.

The following table describes the actions you can perform when you freeze a fabric and when you enable the monitor mode for a fabric.

Operations	Nexus Dashboard Fabric Controller Mode	
	Freeze Mode	Monitor Mode
Create	Save, Preview	Blocked
Breakout	Blocked	Blocked
Unbreakout	Blocked	Blocked
Edit	Save, Preview	Blocked
Delete	Save, Preview	Blocked
Shutdown	Save, Preview	Blocked
No Shutdown	Save, Preview	Blocked
Show	Supported	Supported
Rediscover	Supported	Supported
Deploy	Blocked	Blocked
Import	Supported	Supported
Export	Supported	Supported

The buttons for the associated operations are grayed out accordingly.

If you perform admin operations (shutdown/no shutdown) on SVI, which is part of a config profile, successive **Save & Deploy** operations generate **no interface vlan** command.

For SVI with no policy, on performing admin operation, that is, shutdown/no shutdown command pushed from **Interface Manager**, **int_vlan_admin_state** policy is associated with the SVI.

For example, create and deploy the SVI from **switch_freeform**.


```
interface vlan1234
  description test
  no shutdown
  no ip redirects
  no ipv6 redirects
```

If you shutdown the SVI from interface manager, the **int_vlan_admin_state** policy is associated with the SVI.

Pending diff is shown as:

```
interface Vlan1234
  shutdown
  no ip redirects
  no ipv6 redirects
  description test
  no shutdown
```

Remove the **no shutdown** CLI from the free-form config.

If the user has performed admin operation on SVI, device will have interface in running config. Therefore, post network detach **interface vlan** will be still present and interface will be discovered. You need to manually delete the interface from **Interface Manager**.

The following table describes the fields that appear on **Manage > Inventory > Interfaces > Interfaces**.

Field	Description
Fabric Name	Specifies the fabric name.
Device Name	Specifies the device name.
Interface	Specifies the interface name.
Admin Status	Specifies the administrative status of the interface. The status can be either Up or Down.
Oper-Status	Specifies the operational status of the interface. The status can be either Up or Down.
Reason	Specifies the reason.
Policies	Specifies the policy name.
Overlay Network	Specifies the overlay network.
Sync Status	Specifies the sync status. Specifies if the interface status is In-Sync or Out-Of-Sync.
Interface Group	Specifies the interface group to which the interface belongs to.
Port Channel ID	Specified the port channel ID.

Field	Description
vPC ID	Specifies the vPC ID.
Speed	Specifies the interface speed.
MTU	Specifies the MTU size.
Mode	Specifies the interface mode.
VLANs	Specifies the VLANs.
IP/Prefix	Specifies the interface IP/Prefix.
VRF	Specifies virtual routing and forwarding instances (VRFs).
Neighbour	Specifies the interface neighbour.
Description	Specifies the interface description.  There is a known issue where the description entry is truncated to 64 characters. If the interface description is more than 64 characters, run the <code>snmp ifmib ifalias long</code> command on the switch to increase the description entry length to 256 characters.

Adding Interfaces

To add the interfaces from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **LAN > Interfaces > Interfaces**.
2. Click **Actions > Create new interface** to add a logical interface.

The **Create new interface** window appears.

3. From the **Type** drop-down list, choose the type of the interface.

Valid values are Port Channel, virtual Port Channel (vPC), Straight-through (ST) FEX, Active-Active (AA) FEX, Loopback, Subinterface, Tunnel Ethernet, and Switch Virtual Interface (SVI). The respective interface ID field is displayed when you select an interface type.

- o When you create a port channel through Nexus Dashboard Fabric Controller, add interfaces of the same speed. A port channel that is created from interfaces of varying speeds won't come up. For example, a port channel with two *10 Gigabit Ethernet* ports is valid. However, a port channel with a *10-Gigabit Ethernet + 25-Gigabit Ethernet* port combination isn't valid.
- o To add vPC hosts, you must designate vPC switches in the fabric topology and deploy vPC and peer-link configurations using the **Save Deploy** option. After the vPC pair configurations are deployed, it appears in the Select a vPC pair drop-down box.

You can create a vPC using the **int_vpc_trunk_host** policy.

- o When adding a subinterface, you must select a routed interface from the interface table before clicking the Add button.
- o You can preprovision Ethernet interfaces in the Interface window. This preprovisioning feature

is supported in Easy, eBGP, and External fabrics. .

- o After preprovision the Ethernet interface you can preprovision subinterface on a physical interface.

4. In the **Select a device** field, choose a device.

Devices are listed based on the fabric and interface type.. In the case of vPC or Active to Active FEX, select the vPC switch pair.

5. Enter the ID value in the respective interface ID field (**Port Channel ID**,**vPC ID**,**Loopback ID**, **Tunnel ID**, **Interface name**, **VLAN ID**, and **Subinterface ID**) that is displayed, based on the selected interface.

You can override this value. The new value is used only if it's available in the Resource Manager pool. Else, it results in an error.

6. Under the **Policy** field, select a policy to apply on an interface.

The field only lists the Interface Python Policy with tag *interface_edit_policy* and filtered based on the interface type.

You must not create a **_upg** interface policy. For example, you shouldn't create a policy using the **vpc_trunk_host_upg**, **port_channel_aa_fex_upg**, **port_channel_trunk_host_upg**, and **trunk_host_upg** options.



The policies are filtered based on the interface type you choose in the **Type** drop-down list and the device you choose in the **Select a device** drop-down list.

7. Enter values in the required fields under **Policy Options**.

The fields vary according to the interface type you choose.



From Cisco Nexus Dashboard Fabric Controller Release 11.5(1) you can mirror the configurations of Peer-1 on Peer-2 while creating a vPC. When you check the **Enable Config Mirroring** check box, the Peer-2 fields will be grayed out. The configurations that you enter in the Peer-1 fields will be copied to Peer-2 fields.

From Cisco NDFC Release 12.1.2e, you can set Native Vlan for the interface which has **int_trunk_host** or **int_port_channel_trunk_host**, or **int_vpc_trunk_host** policy template.

A trunk port can carry nontagged packets simultaneously with tagged packets. When you assign a default port VLAN ID to the trunk port, all untagged traffic travels on the default port VLAN ID for the trunk port, and all untagged traffic is assumed to belong to this VLAN. This VLAN is referred to as the native VLAN ID for a trunk port. That is, the native VLAN ID is the VLAN that carries untagged traffic on trunk ports.

8. Click **Save** to save the configurations.



To apply QoS policies on the interface, create the interface freeform with references accordingly.

Only saved configurations are pushed to the device. While adding the interface, you can only modify the policy attribute after the first save. If you try to use an ID that is already used, you encounter the *Resource could not be allocated* error.

9. To view configurations for a new interface, double-click on the policy name in the **Policies** tab.
10. (Optional) Click the **Preview** option to preview the configurations to be deployed.
11. Click **Deploy** to deploy the specified logical interface.

The newly added interface appears in the screen.

Breakout

To breakout an interface, from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. On Interface window, click **Actions > Breakout**.

The **Breakout Interfaces** window appears.

2. Choose the required option in the window and click **Breakout**.

The available options are 10g-4x, 25g-4x, 50g-2x, 50g-4x, 100g-2x, 100g-4x, 200g-2x, and Unbreakout.

UnBreakout

You can unbreakout interface that are in breakout state.

On the **Interface** window, click **Actions > UnBreakout**.



The unbreakout option is grayed out for interfaces that are not in breakout state.

Editing Interfaces

To edit the interfaces from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:



The **Edit** interface allows you to change the policy and add or remove an interface from a port channel or vPC.

1. Choose **Manage > Inventory > Interfaces**.

You can break out and unbreak out an interface by using the breakout option in the **Actions** menu.

2. Select the interface check box to edit an interface or vPC.

Select corresponding check boxes for editing multiple interfaces. You cannot edit multiple port channels and vPC. You cannot edit interfaces of different types at the same time.

3. Click **Actions > Edit** to edit an interface.

The variables that are shown in the **Edit interface(s)** window are based on the template and its policy. Select the appropriate policy. Save the policy and deploy the same. This window lists only Interface Python Policy with the tag *interface_edit_policy* and filtered based on the interface type.

In a vPC setup, the two switches are in the order the switch names are displayed in the edit window. For example, if Switch Name is displayed as *LEAF1:LEAF2*, then Leaf1 is peer switch one and Leaf2 is peer switch two.

During overlay network deployment on switches, the network can be associated with trunk interfaces. The trunk interface to network association is reflected in the **Interfaces** tab. You can update such interfaces.

For PVLAN interfaces, you can associate interfaces with the networks only for access and trunk port types.

For interface policies that are not created from the **Manage > Inventory > Interfaces > Interfaces** screen, you can edit some configurations but not change the policy itself. The policy and fields that cannot be edited are grayed out.

The following are some examples of policies that cannot be edited:

- Loopback interface policies - The *int_fabric_loopback* policy is used to create a loopback interface. You can edit the loopback IP address and description but not the *int_fabric_loopback* policy instance.

You cannot edit the loopback IP addresses for loopback interfaces that are created automatically while creating and attaching the VRFs.

- Fabric underlay network interface policies (*int_fabric_num*, for example) and fabric overlay network interface (NVE) policies.
- Policies associated with port channels and member ports of port channels, including the port channels and member ports associated with a vPC.
- SVIs created during network and VRF creation. The associated VLANs appear in the interfaces list.

Editing Interfaces Associated with Links

There are two types of links serving a different purpose within the network:

- Intra-fabric links connect devices within the same Easy Fabric. These links are commonly used for internal connectivity, such as between spine and leaf switches.
- Inter-fabric links, on the other hand, connect an Easy Fabric to external networks or other Easy Fabrics. These links are typically used for WAN connectivity or DCI scenarios.

A policy is associated with each link that effectively states the configuration that is applied to both ends of the link. In other words, the link policy becomes the parent of the individual child interface policies that are associated with the two interfaces that form the link. In this scenario, you must edit the link policy to edit the interface policy fields such as description, IP address, and any per interface freeform config. The following procedure shows how to edit the interfaces associated with links:

1. Choose **Manage > Inventory > Interfaces > Interfaces**.

2. Select a link and click **Actions > More > Rediscover Interface**.

Applying Normalize Action on Interfaces

The Normalize action allows you to apply the same interface policy and its parameters to multiple interfaces in one shot.



- The **Normalize** action has some key differences from the Interface Groups. Normalize action generates an individual policy for each interface, whereas the Interface Group generates a single policy for all the interfaces.
- You can use **Normalize** to apply any available interface policies, whereas Interface Groups can only apply the “share_trunk_host” policy.

To apply **Normalize** action to multiple interfaces from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Manage > Inventory > Interfaces**.
2. Select the check boxes for all interfaces you want to configure in bulk.
3. Click **Actions > Normalize** to update the configuration on all the selected interface.

Deleting Interfaces

To delete the interfaces from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:



- This option allows you to delete only logical ports, port channels, and vPCs. You can delete the interface if it does not have overlay or underlay policy attached.
- When a port channel or vPC is removed, the corresponding member ports get the default policy associated. The default policy can be configured in server.properties file.

1. Choose **Manage > Inventory > Interfaces > Interfaces**.
2. Select the interfaces and choose **Actions > More > Delete Interface**.

You cannot delete logical interfaces created in the fabric underlay.

3. Click **Save**.
4. Click **Deploy** to delete the interface.

Shutting Down and Bringing Up Interfaces

To shut down and bring up the interfaces from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Manage > Inventory > Interfaces > Interfaces**.
2. Select the interfaces that you want to shut down or bring up.

Note: For all interfaces except VLANs, you cannot perform a **Shutdown** or **No Shutdown** on

interfaces that do not have a policy attached.

3. Click **Shutdown** to disable the selected interfaces. For example, you may want to isolate a host from the network or a host that is not active in the network.

A confirmation window appears where you can save, preview, and deploy the changes.

4. Click **Save** to preview or deploy the changes.
5. Click **No Shutdown** to bring up the selected interfaces.

A confirmation window appears where you can save, preview, and deploy the changes.

6. Click **Save** to preview or deploy the changes.

Viewing Interface Configuration

To view the interface configuration commands and execute them from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Manage > Inventory > Interfaces > Interfaces**.
2. Select the interface whose configurations you want to view and choose **Actions > More > Show commands**.
3. In the **Interface show commands** window, select the action from the **Commands** drop-down list and click **Execute**.

The interface configurations are displayed on the right of the screen.

For Show commands, you must have corresponding **show** templates for interface or interface sub types like port channel or vPC, defined in the **Templates**.

Rediscovering Interfaces

To rediscover the interfaces from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Manage > Inventory > Interfaces > Interfaces**.
2. Select the interfaces that you want to rediscover and choose **Actions > More > Rediscover Interface** to rediscover the selected interfaces.

For example, after you edit or enable an interface, you can rediscover the interface.

Viewing Interface History

To view the interface history from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Manage > Inventory > Interfaces > Interfaces**.
2. Select the interface and choose **Actions > More > Deployer History** to view the configuration history on the interface.

3. Click **Status** to view each command that is configured for that configuration instance.

Deploying Interface Configurations

To deploy the interface configuration from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Manage > Inventory > Interfaces > Interfaces**.
2. Select an interface that you want to deploy and choose **Actions > Deploy Interfaces** to deploy or redeploy configurations that are saved for the interface.



You can select multiple interfaces and deploy pending configurations.

After you deploy the interface configuration, the interface status information is updated. However, the overall switch-level state may be in the pending state, which is in blue. The overall switch-level state goes to the pending state whenever there is a change in intent from any module, such as interface, link, policy template update, top-down, or so on. In the pending state, a switch may have pending configurations or switch-level recomputation. The switch-level recomputation occurs when:

- o You deploy for the switch
- o During a deploy
- o During hourly sync

Creating External Fabric Interfaces

You can add and edit port channel, vPC, subinterface, and loopback interfaces for external fabric devices. You cannot add Straight-through FEX and Active-Active FEX functions.

The Breakout port function is only supported for the Cisco Nexus 9000, 3000, and 7000 Series Switches in the external fabric.

When you add an interface to an external fabric device, the Resource Manager is not in sync with the device. So, ensure that the value populated in the ID field (Port-channel ID, vPC ID, Loopback ID, etc) is not previously configured on the switch.

If you want to configure a portchannel in the external fabric, you should add and deploy the **feature_lacp** policy on the switches where the portchannel will be configured.

When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. If you click **Save & Deploy** in the fabric topology screen, it displays an error message. However, the following settings (available when you right-click the switch icon) are allowed:

- vPC pairing – You can designate a vPC switch pair, but it is only for reference.
- View/edit policy – You can add a policy but you cannot deploy it on the switch.
- Manage interfaces – You can only create intent for adding interfaces. If you try to deploy, edit, or delete interfaces, it results in an error message.

Interface Groups

An interface group consists of multiple interfaces with same attributes. You can create an interface group that allows grouping of host-facing interfaces at a fabric level. Specifically, you can create an interface group for physical Ethernet interfaces, Layer 2 port-channels, and vPCs. You can attach or detach multiple overlay networks to the interfaces in an interface group.

Shared Policy

From Cisco NDFC Release 12.1.2e, you can create and add a shared policy for an interface group. This policy ensures update of appropriate configurations for all the interfaces in the Interface group. In the shared policy, all the interfaces will have the same underlay and overlay attributes. When you change the configuration in the shared policy, then that configuration is applied to all the interfaces.

You can see the details of shared policy under the **Policies** column.

Custom policy can also be created by selecting the policy from the template list and **Duplicate Template** to add the additional information. The shared policy must contain the tags **interface_edit_policy**, **interface_edit_shared_policy**, and **int_trunk**.

Guidelines

- Interface groups are only supported for the following fabric template types:
 - **BGP Fabric**
 - **Campus VXLAN EVPN**
 - **Data Center VXLAN EVPN**
 - **Enhanced Classic LAN**
- An interface group is specific to a fabric. For example, consider two fabrics: Fab1 and Fab 2. The interface group IG1 in Fab1 isn't applicable to Fab 2.
- An interface group can only have interfaces of a certain type. For example, you need three separate interface groups if you want to group three types of interfaces such as IG1 for physical Ethernet trunk interfaces, IG2 for Layer 2 trunk port-channels, and IG3 for vPC host trunk ports.
- An interface group can also be created using preprovisioned interfaces.
- Interface groups are supported only to switches with leaf or border roles. For Border Gateway roles, Interface Groups are supported only on vPC BGWs but not on Anycast BGW, BGW Spine, or BGW SuperSpine.
- From Cisco NDFC Release 12.1.2e, you can include Layer 2 ToR interfaces in the interface groups.
- Interfaces added to an interface group with a shared policy replaces the individual policy and get the shared policy.
- You can change description and status of each interface in interface group.
- Interface removed from an interface group with a shared policy will set to a default policy.
- VMs should have the same configuration for all the interfaces under the shared policy.
- Shared policy is supported only for **Ethernet** interface type in interface group.
- Ethernet interface groups now have a common policy.

- **Port-Channel** and **vPC** interface types are not supported for adding shared policy in interface group.
- When the MTU value in the shared policy has to be changed, make sure to update the fabric settings with the same value across all switches of that fabric.
- When the **ptp**, **ttag**, and **ttag-strip** option from the shared policy has to be used, make sure to enable PTP globally in fabric settings.
- When the netflow option has to be used, make sure all the interfaces of interface groups are capable of netflow configuration and that it is enabled globally in fabric settings.
- For Layer 2 port-channels and vPCs that are part of an interface group, they can't be deleted until they are de-associated from the interface group even if there are no networks associated with the interface group. Similarly, a trunk port that has no overlay networks but is part of an IG can't be converted to an access port. In other words, you can't change policies for interfaces that are part of an interface group. However, you can edit certain fields for policies.
- For L4-L7 services configuration on leaf switches, trunk ports that are used for services attachment can't be part of interface groups.
- When you perform a per fabric backup of an easy fabric, if there are interface groups created in that fabric, all the associated interface group state is backed up.
- If an easy fabric contains an interface group, then this fabric can't be imported into the MSO. Similarly, if an easy fabric has been added to the MSO, you can't create interface groups for interfaces that belong to switches in the easy fabric.
- The **Add to Interface Group** and **Remove from Interface Group** button is enabled only for Admin and Stager users. For all other users, this button is disabled.
- The **Interface Group** button is disabled in the following circumstances:
 - Select any other interface apart from vPC, Port-channel, and Ethernet.
 - If the interface has a policy attached from another source, for example:
 - If the interface is member of a port-channel or vPC.
 - If the port-channel is member of vPC.
 - If the interface has a policy from underlay or links.



If you select different types of interfaces, the **Interface Group** button is enabled. However, when you try to create or save different types of interfaces to an interface group, an error is displayed.

Creating an Interface Group

To create an interface group from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Manage > Inventory > Interfaces > Interface Groups**.
2. Choose **Actions > Create new interface group**.
3. From the **Select Fabric** window, select a fabric and click **Select**.
4. In the **Create new interface group** window, provide an interface group name in the **Interface Group Name** field, select an Interface Type, and click **Save**.

An interface group name can have a maximum length of 64 characters.



An interface can be added to single interface group only.

5. Click on the **Policy** field.

The **Select Policy** window appears.

6. Choose **int_shared_trunk_host** policy and then click **Select**.

From Cisco NDFC Release 12.1.2e, a new **Policy** field is introduced in the **Create new interface group** window. You can add a shared policy to the interface group which can be shared by the interfaces existing in that group. Shared Policy is optional, for upgrades, all the existing interface group will not have a policy.



The policy field supports the **Ethernet** interface type only.

7. Enter the mandatory parameters in the text field and click **Save**.
8. In the **Interfaces** tab, select the interfaces to group and choose **Actions > Add to interface Group**.

The **Add new interface Group** window appears.

9. To create a custom interface group, enter an interface group name in the **Select Interface Group** field and click **Create custom**.

If you have already created an interface group, select it from the **Select Interface Group** drop-down list. Also, if an interface is already part of an interface group, you can move it to a different interface group by selecting the new group from the **Select Interface Group** drop-down list.

You can create interface groups from either the **Interfaces Groups** window or the **Interfaces** window under Fabric Overview.

10. Click **Save**.

In the **Interfaces** window, you can see the interface group name under the **Interface Group** column.

11. To edit an interface group, choose **Actions > Edit Interface Group**. You can update the policy options after you assign the shared policy.



You cannot edit or delete the shared policy template.

Adding and Removing Interfaces from an Interface Group

You can add or remove multiple interfaces to Interface Groups from the Cisco Nexus Dashboard Fabric Controller Web UI.

Adding Interfaces to an Interface Group

To add interfaces to an interface group from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Create an interface group.

For information on creating an interface group, see [Creating an Interface Group](#).

2. Choose **Manage > Inventory > Interfaces**.
3. Select the interfaces to add to an interface group and choose **Actions > More > Add to interface Group**.

The Add to Interface Group page appears.

4. Select your Interface Group's name in the **Interface Group** dropdown list.
5. Click **Save** to add the selected interfaces to the Interface Group.

Removing Interfaces from an Interface Group

To remove interfaces from an interface group from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Choose **Manage > Inventory > Interfaces**.
2. Select the interfaces to disassociate from an interface group and choose **Actions > More > Remove from Interface Group**.

A dialog box appears asking whether you want to clear all the associated interfaces.

3. Click **Yes** to proceed.

Note that if there are any networks attached to these interfaces, they are detached as well when you click **Clear**.

Attaching Networks to an Interface Group

To attach networks to an interface group from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. Double click on the fabric to launch **Fabric Overview**.
2. On the **Networks** tab, select the networks that you need to attach to an interface group and click **Interface Group**.



An overlay network can belong to multiple interface groups. You can select only the networks with a VLAN ID. Otherwise, an error message is displayed.]

3. In the **Interface Groups** window, you can perform the following:
 - o Select an existing interface group from the **Select Interface Group** drop-down list and click **Save**.

For example, you select three networks and the interface group **test**, and click the **Save** button, the following operations are performed in the background:

- a. Nexus Dashboard Fabric Controller retrieves interfaces that are part of the interface group **test**.
- b. Nexus Dashboard Fabric Controller determines that three networks are added to the interface group **test**. Therefore, it autoattaches these networks to all the interfaces that are part of the interface group **test**.
- c. For each interface, Nexus Dashboard Fabric Controller pushes the "switchport trunk allowed vlan add xxxx" command three times for each selected network.



Nexus Dashboard Fabric Controller ensures that there's no duplicate configuration intent.

If you click **Clear**, Nexus Dashboard Fabric Controller pushes "switchport trunk allowed vlan remove xxx" config intent.

- o Create a custom interface group by entering an interface group name in the **Select Interface Group** field and click **Create new interface group**. Click **Save**.

If you choose this option, make sure to add interfaces to this Interface Group in the **Interfaces** window. As a result, Nexus Dashboard Fabric Controller performs the following operations:

- Removes all existing overlay networks that don't belong to the interface group from these interfaces.
- Adds new overlay networks to these interfaces that are part of the interface group but not yet attached to these interfaces.

For more information about associating interfaces to interface groups, see [Creating an Interface Group](#).

4. Click **Actions > Recalculate & Deploy** to deploy the selected networks on the switches.

Detaching a Network from an Interface Group

This procedure shows how to detach a network from an interface group in the Networks window. Also, you can detach networks when you remove an interface from an interface group in the **Interfaces** window. For more information, see *Removing Interfaces from an Interface Group*.

1. Double click on the fabric to launch **Fabric Overview**.
2. On the **Networks** tab, select the networks that you need to detach to an interface group and click **Add to Interface Group**.
3. In the **Add to Interface Groups** window, select the interface group from the **Select Interface Group** drop-down list and click **Clear** to detach a network.
4. (Optional) Navigate to **Manage > Inventory > Interfaces**.

Under the **Overlay Network** column, you can see the detached network in the red color for the corresponding interface. Click the network to view the expected config that is struck through.

5. Navigate to the **Networks** tab and choose **Actions > Recalculate & Deploy**.

Deleting an Interface Group

An interface group is automatically deleted when it's not in use. . You can perform an explicit delete by clicking on **Interface Group > Actions > Delete Interface group**. This check is performed whenever you click the **Clear** button in the **Edit Interface Group** window. There may be exception scenarios where you need to clean up the interface groups explicitly.

For example, you create an interface group **storagelG** and add an interface to it. Later, you want to change the interface mapping to another group. Therefore, you select the interface and click **Interface Group** to open the **Edit Interface Group** window. Select the other interface group named **disklG**. Now, the **storagelG** interface group doesn't have any associated member interfaces or networks. In this case, perform the following steps:

1. Select an interface that doesn't belong to an interface group.
2. Click **Interface Group** to open the **Edit Interface Group** window.
3. Select the **StorageIG** interface group from the **Select Interface Group** drop-down list.
4. Click **Clear**.

Out-of-Band Switch Interface Configurations

Any interface level configuration made outside of Nexus Dashboard Fabric Controller (via CLI) can be synced to Nexus Dashboard Fabric Controller and then managed from Nexus Dashboard Fabric Controller. Also, the vPC pair configurations are automatically detected and paired. This applies to the External_Fabric and Classic LAN fabrics only. The vPC pairing is performed with the **vpc_pair** policy.



When Nexus Dashboard Fabric Controller is managing switches, ensure that all configuration changes are initiated from Nexus Dashboard Fabric Controller and avoid making changes directly on the switch.

When the interface config is synced up to the Nexus Dashboard Fabric Controller intent, the switch configs are considered as the reference, that is, at the end of the sync up, the Nexus Dashboard Fabric Controller intent reflects what is present on the switch. If there were any undeployed intent on Nexus Dashboard Fabric Controller for those interfaces before the resync operation, they will be lost.

Guidelines

- Supported in fabrics using the following fabric templates: Data Center VXLAN EVPN, External, and Classic LAN.
- Supported for Cisco Nexus switches only.
- Supported for interfaces that do not have any fabric underlay related policy associated with them prior to the resync. For example, IFC interfaces and intra fabric links are not subjected to resync.
- The time taken by host port resync depends on the number of switches/interfaces to be synchronized.
- Supported for interfaces that do not have any custom policy (policy template that is not shipped with Cisco Nexus Dashboard Fabric Controller) associated with them prior to resync.
- Supported for interfaces where the intent is not exclusively owned by a Cisco Nexus Dashboard Fabric Controller feature and/or application prior to resync.
- Supported on switches that do not have Interface Groups associated with them.
- Interface mode (switchport to routed, trunk to access, and so on) changes are not supported with overlays attached to that interface.

The sync up functionality is supported for the following interface modes and policies:

Interface Mode	Policies
trunk (standalone, po, and vPC PO)	• int_trunk_host • int_port_channel_trunk_host • int_vpc_trunk_host
access (standalone, po, and vPC PO)	• int_access_host • int_port_channel_access_host • int_vpc_access_host

dot1q-tunnel	▪ int_dot1q_tunnel_host ▪ int_port_channel_dot1q_tunnel_host ▪ int_vpc_dot1q_tunnel_host
routed	int_routed_host
loopback	int_freeform
sub-interface	int_subif
FEX (ST, AA)	▪ int_port_channel_fex ▪ int_port_channel_aa_fex
breakout	interface_breakout
nve	int_freeform (only in External_Fabric/Classic LAN)
SVI	int_freeform (only in External_Fabric/Classic LAN)
mgmt0	int_mgmt

In an Easy fabric, the interface resync will automatically update the network overlay attachments based on the access VLAN or allowed VLANs on the interface.

After the resync operation is completed, the switch interface intent can be managed using normal Nexus Dashboard Fabric Controller procedures.

Syncing up Switch Interface Configurations

We recommend that you deploy all switch configurations from NDFC. In some scenarios, it may be necessary to make changes to the switch interface configuration out-of-band. This will cause configuration drift causing switches to be reported Out-of-Sync.

NDFC supports syncing up the out-of-band interface configuration changes back into its intent.

Guidelines and Limitations

The following limitations are applicable after syncing up the switch interface configurations to NDFC:

- This feature is not supported on ToR/Access switches, or on leaf switches with ToR pairing present.
- The port channel membership changes (once the policy exists) are not supported.
- Changing the interface mode (trunk to access and so on) that have overlays attached is not supported.
- Resync for interfaces that belong to **Interface Groups** are not supported.
- The vPC pairing in **External_Fabric** and **Classic LAN** templates must be updated with the **vpc_pair** policy.
- The resync can be performed for a set of switches and repeated as desired.
- The time taken by host port resync depends on the number of switches/interfaces to be synchronized.

- In **Data Center VXLAN EVPN** fabrics, VXLAN overlay interface attachments are performed automatically based on the allowed VLANs.

Before you begin

- We recommend taking a fabric backup before attempting the interface resync.
- In **External Fabric** and **Classic LAN** fabrics, for the vPC pairing to work correctly, both the switches must be in the fabric and must be functional.
- Ensure that the switches are **In-Sync** and switch mode must not be **Migration** or **Maintenance**.
- From the **Actions** drop-list, choose **Discovery > Rediscover** to ensure that NDFC is aware of any new interfaces and other changes.

Procedure

1. Choose **Manage > Fabrics** and double-click on a fabric.

The **Fabric Overview** window appears.

2. Click the **Switches** tab and ensure that switches are present in the fabric and vPC pairings are completed.
3. Click the **Policies** tab and select one or more switches where the interface intent resync is needed.



- If a pair of switches is already paired with either **no_policy** or **vpc_pair**, select only one switch of the pair.
- If a pair of switches is not paired, then select both the switches.

4. From the **Actions** drop-down list, choose **Add Policy**.

The **Create Policy** window appears.

5. On the **Create Policy** window, choose **host_port_resync** from the **Policy** drop-down list.
6. Click **Save**.
7. Check the **Mode** column for the switches to ensure that they report **Migration**. For a vPC pair, both switches are in the **Migration-mode**.
 - After this step, the switches in the **Topology view** are in **Migration-mode**.
 - Both the switches in a vPC pair are in the migration mode even if one of the switches is placed into this mode.
 - If switches are unintentionally put into the resync mode, they can be moved back to the normal mode by identifying the **host_port_resync** policy instance and deleting it from the **Policies** tab.
8. After the configuration changes are ready to sync up to NDFC, navigate to the **Switches** tab and select the required switches.
9. Click **Recalculate & Deploy** to start the resync process.



This process might take some time to complete based on the size of the switch configuration and the number of switches involved.

10. The **Deploy Configuration** window is displayed if no errors are detected during the resync operation. The interface intent is updated in NDFC.



If the External_Fabric or Classic LAN fabric is in **Monitored Mode**, an error message indicating that the fabric is in the read-only mode is displayed. This error message can be ignored and doesn't mean that the resync process has failed.

Close the **Deploy Configuration** window, and you can see that the switches are automatically moved out of the **Migration-mode**. Switches in a vPC pair that were not paired or paired with **no_policy** show up as paired and associated with the **vpc_pair** policy.



The **host_port_resync** policy that was created for the switch is automatically deleted after the resync process is completed successfully.

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)

© 2017–2025 Cisco Systems, Inc. All rights reserved.