

Switches

- Switches, on page 1
- Switch Overview, on page 26

Switches

The following table describes the fields that appear on **Switches** window.

Field	Description
Switch	Specifies name of the switch.
IP Address	Specifies IP address of the switch.
Role	Specifies role assigned on the switch.
Serial Number	Specifies the serial number of the switch.
Fabric Name	Specifies the associated fabric name for the switch.
Config Status	Specifies the configuration status. Status will be either In-Sync or Out-of-sync.
Oper Status	Specifies the configuration status. Status will be either In-Sync or Out-of-sync.
Discovery Status	Specifies the discovery status of the switch.
Model	Specifies the switch model.
vPC Role	Specifies the vPC role of the switch.
vPC Peer	Specifies the vPC peer of the switch.

Adding Switches to a Fabric

UI Path: **LAN** > **Switches** > **Actions** > **Add Switches**

Switches in each fabric are unique, and hence, only one switch can be added to one fabric.

Note Cisco Nexus Dashboard has 2 logical interfaces per node, namely, Management interface (bond1br) and Fabric (also known as data) interface (bond0br). For Cisco Nexus Dashboard Fabric Controller, Nexus Dashboard Management and Fabric interfaces must be in different IP subnets. By default, the route for Nexus Dashboard services is through the fabric interface. An operator must add static routes on Nexus Dashboard Management Network to connect with switches that must be reached over management interface (bond1br). This ensures that a route for the pods uses management interface as the exit interface.

Note Make sure that the switch user role for discovery or add switches or LAN credentials for NDFC must have the network-admin role.

To add switches to the existing fabric, perform below procedures:

1. From Nexus Dashboard Fabric Controller Web UI, choose **LAN** > **Switches**.
2. On Switches tab, Choose **Actions** > **Add Switches**.

   The **Add Switches** window appears.

   Similarly, you can add switches on Topology window. On topology window, choose a fabric, right-click on the fabric and click **Add Switches**.
3. On add switches window, click **Choose Fabric**, click appropriate fabric, and then click **Select**.

   The **Add Switches** window has a default discover tab and other tabs appears based on the fabric selected.

Also, you can pre-provision switches and interfaces. For more information, see pre-provision device and pre-provisioning ethernet interface.

Note NDFC supports switch discovery only for default system-name(serial number).

Note When Nexus Dashboard Fabric Controller discovers a switch with the hostname containing the period character (.), it is treated as a domain-name and truncated. Only the text before the period character (.) is considered as a hostname. For example:

- If hostname is **leaf.it.vxlan.bgp.org1-XYZ**, Nexus Dashboard Fabric Controller shows only **leaf**
- If hostname is **leaf-itvxlan.bgp.org1-XYZ**, Nexus Dashboard Fabric Controller shows only **leafit-vxlan**

Note Ensure that the Switch name or the Host name is unique within the Fabric.

Discovering New Switches

1. When a new Cisco NX-OS device is powered on, typically that device has no startup configuration or any configuration state for that matter. Consequently, it powers on with NX-OS and post initialization, goes into a POAP loop. The device starts sending out DHCP requests on all the interfaces that are up including the mgmt0 interface.
2. As long as there is IP reachability between the device and the Nexus Dashboard Fabric Controller, the DHCP request from the device, will be forwarded to the Nexus Dashboard Fabric Controller. For easy day-0 device bring-up, the bootstrap options should be enabled on the **Fabric Settings** as mentioned earlier.
3. With bootstrap enabled for the fabric, the DHCP request coming from the device will be serviced by the Nexus Dashboard Fabric Controller. The temporary IP address allocated to the device by the Nexus Dashboard Fabric Controller will be employed to learn basic information about the switch including the device model, device NX-OS version, etc.
4. In the Nexus Dashboard Fabric Controller UI, choose **Switch** > **Actions** > **Add Switches**.

   The **Add Switches** window appears with default tabs.
5. Choose **Bootstrap(POAP)** radio button.

   As mentioned earlier, Nexus Dashboard Fabric Controller retrieves the serial number, model number, and version from the device and displays them on the Inventory Management along window. Also, an option to add the IP address, hostname, and password are made available. If the switch information is not retrieved, refresh the window.



Note

- At the top left part of the window, *export* and *import* options are provided to export and import the .csv file that contains the switch information. You can pre-provision devices using the *import* option as well.

From Cisco NDFC Release 12.1.1e, for pre-provisioned and bootstrap switches dummy values can be added for the serial number. After configuring the network successfully, serial number can be changed with the appropriate number of the switch on the Switches tab.

Note: You can change serial number only for Nexus 9000 Series switches.

Select the checkbox next to the switch and enter the switch credentials: IP address and host name.

Add the IPv4 address in the **IP Address** field.

You can provision devices in advance. To pre-provision devices, refer to Pre-provisioning device section.

6. In the **Admin Password** and **Confirm Admin Password** fields, enter and confirm the admin password.

   This admin password is applicable for all the switches displayed in the POAP window.

   You can specify a new user. Choose radio button **Specify a new user** enter **Username**, **Password** and choose **Authentication Protocol** from drop-down list.



Note If you do not want to use admin credentials to discover switches, you can instead use the AAA authentication, that is, RADIUS or TACACS credentials for discovery only.

7. (Optional) Use discovery credentials for discovering switches.
   a. Click the **Add Discovery Credentials** icon to enter the discovery credentials for switches.
   b. In the **Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

      Click **OK** to save the discovery credentials.

      If the discovery credentials are not provided, Nexus Dashboard Fabric Controller uses the admin user and password to discover switches.
8. Click **Bootstrap** at the top right part of the screen.

   Nexus Dashboard Fabric Controller provisions the management IP address and other credentials to the switch. In this simplified POAP process, all ports are opened up.
9. Click **Refresh Topology** to get updated information. The added switch goes through the POAP cycle. Monitor and check the switch for POAP completion.
10. After the added switch completes POAP, the fabric builder topology page is refreshed with the added switch thereby depicting its discovered physical connections. Set the appropriate role for the switch followed by a Deploy Config operation at the fabric level. The Fabric Settings, switch role, the topology etc. are evaluated by the Fabric Builder and the appropriate intended configuration for the switch is generated as part of the Save operation. The pending configuration will provide a list of the configurations that need to be deployed to the new switch in order to bring it IN-SYNC with the intent.

Note

- For any changes on the fabric that results in Out-of-Sync, you must deploy the changes. The process is the same as explained in the *Discovering Existing Switches* section.
- During fabric creation, if you have entered AAA server information (in the **Manageability** tab), you must update the AAA server password on each switch. Else, switch discovery fails.
- When discovering devices using SNMP, if you have configured to use an AAA server for authentication, the command **sync-snmp-password** <*password*> <*username*> is run on the switch through NDFC to generate a cached user. The authentication uses MD5, by default. You must specify the SNMPv3 authentication and privacy protocol attributes in the switch AV-pair as follows: `snmpv3:auth=SHA priv=AES-128`

11. After the pending configurations are deployed, the **Progress** column displays 100% for all switches.
12. Click **Close** to return to the fabric builder topology.
13. Click **Refresh Topology** to view the update. All switches must be in green color indicating that they are functional.
14. The switch and the link are discovered in Nexus Dashboard Fabric Controller. Configurations are built based on various policies (such as fabric, topology, and switch generated policies). The switch image (and other required) configurations are enabled on the switch.
15. In the Nexus Dashboard Fabric Controller GUI, the discovered switches can be seen in the *Standalone* fabric topology. Up to this step, the POAP is completed with basic settings. You must setup interfaces through the **LAN** > **Switches**. Select a switch, a slide-in pane appears, click **Launch** icon. On **Switches Overview** tab, click **Interface** tab for any additional configurations, but not limited to the following:

- vPC pairing.
- Breakout interfaces.
- Port channels, and adding members to ports.

When you enable or disable a vPC pairing/un-pairing or the advertise-pip option, or update Multi-Site configuration, you should use the **Deploy Config** operation. At the end of the operation, an error prompts you to configure the **shutdown** or **no shutdown** command on the nve interface. A sample error screenshot when you enable a vPC setup.

To resolve, go to the **Interfaces** > **Actions** > **Deploy** tab and deploy the Shutdown operation on the nve interface followed by a No Shutdown configuration. This is depicted in the figure below where the up arrow corresponds to a No Shutdown operation while a down arrow corresponds to a Shutdown operation.

You can right-click the switch to view various options:

- **Set Role** - Assign a role to the switch (Spine, Border Gateway, and so on).



Note

- Changing of the switch role is allowed only before executing **Deploy Config**.
- Switch roles can be changed if there are no overlays on the switches, but only as per the list of allowed switch role changes given at switch operations section.

- **Modes** - Maintenance and Active/Operational modes.
- **vPC Pairing** - Select a switch for vPC and then select its peer.

  You can create a virtual link for a vPC pair or change the existing physical link to a virtual link for a vPC pair.
- **Manage Interfaces** - Deploy configurations on the switch interfaces.
- **View/Edit Policies** - See switch policies and edit them as required.
- **History** - View per switch deployment history.
- **History** - View per switch deployment and policy change history.

  The **Policy Change History** tab lists the history of policies along with the users who made the changes like add, update, or delete.

  Under the **Policy Change History** tab, for a policy, click **Detailed History** under the **Generated Config** column to view the generated config before and after.

  The following table provides the summary of generated config before and after for Policy Template Instances (PTIs).

PTI Operations	Generated Config Before	Generated Config After
Add	Empty	Contains the config
Update	Contains config before changes	Contains config after changes

PTI Operations	Generated Config Before	Generated Config After
Mark-Delete	Contains the config to be removed.	Contains the config to be removed with color change.
Delete	Contains the config	Empty

Note When a policy or profile template is applied, an instance is created for each application of the template, which is known as Policy Template Instance or PTI.

- **Preview Config** - View the pending configuration and the side-by-side comparison of the running and expected configuration.
- **Deploy Config** - Deploy per switch configurations.
- **Discovery** - You can use this option to update the credentials of the switch, reload the switch, rediscover the switch, and remove the switch from the fabric.

The new fabric is created, the fabric switches are discovered in Nexus Dashboard Fabric Controller, the underlay configuration provisioned on those switches, and the configurations between Nexus Dashboard Fabric Controller and the switches are synced. The remaining tasks are:

- Provision interface configurations such as vPCs, loopback interface, and subinterface configurations.
- Create networks and deploy them on the switches.

Discovering Existing Switches

To discover existing switches in Cisco Nexus Dashboard Fabric Controller Web UI, perform the following procedure:

Procedure

Step 1 After you click **Add Switches**, click **Discover Switches** to add one or more existing switches into the fabric.

In this case, a switch with known credentials and a pre-provisioned IP address, is added to the fabric.

Step 2 The IP address (Seed IP), username, and password (**Username** and **Password** fields) of the switch are provided as the input by a user. The **Preserve Config** check box is chosen by default. This is the option that a user would select for a brownfield import of a device into the fabric. For a greenfield import where the device configuration will be cleaned up as part of the import process, the user should set the **Preserve Config** check box is not selected.

Note
Easy_Fabric_eBGP does not support brownfield import of a device into the fabric.

Step 3 Click **Discover Switches**.

The **Add Switches** window appears. Since the **Max Hops** field was populated with 2 (by default), the switch with the specified IP address (leaf-91) and switches two hops from that switch, are populated in the **Add Switches** result.

Step 4 If the Cisco Nexus Dashboard Fabric Controller was able to perform a successful shallow discovery to a switch, the status column shows as **Manageable**. Choose the check box next to the appropriate switch(es) and click **Add Switches**.

Though this example describes the discovery of one switch, multiple switches can be discovered at once.

The switch discovery process is initiated. The **Progress** column displays progress for all the selected switches. It displays **done** for each switch on completion.

Note
You must not close the screen (and try to add switches again) until all selected switches are imported or an error message comes up.

If an error message comes up, close the screen. The fabric topology screen comes up. The error messages are displayed at the top right part of the screen. Resolve the errors wherever applicable and initiate the import process again by clicking **Add Switches** in the Actions panel.

Cisco Nexus Dashboard Fabric Controller discovers all the switches, and the Progress column displays **done** for all switches, close the screen. The *Standalone* fabric topology screen comes up again. The switch icons of the added switches are displayed in it.

Note
You will encounter the following errors during switch discovery sometimes.

Step 5 Click **Refresh topology** to view the latest topology view.

When all switches are added and roles assigned to them, the fabric topology contains the switches and connections between them.

Step 6 After discovering the devices, assign an appropriate role to each device. For more information on roles, refer Assign Roles.

If you choose the Hierarchical layout for display (in the Actions panel), the topology automatically gets aligned as per role assignment, with the leaf devices at the bottom, the spine devices connected on top of them, and the border devices at the top.

Assign vPC switch role - To designate a pair of switches as a vPC switch pair, right-click the switch and choose the vPC peer switch from the list of switches.

AAA server password - During fabric creation, if you have entered AAA server information (in the **Manageability** tab), you must update the AAA server password on each switch. Else, switch discovery fails.

When a new vPC pair is created and deployed successfully using Cisco Nexus Dashboard Fabric Controller, one of the peers might be out-of-sync for the **no ip redirects** CLI even if the command exists on the switch. This out-of-sync is due to a delay on the switch to display the CLI in the running configuration, which causes a diff in the configuration compliance. Re-sync the switches in the **Config Deployment** window to resolve the diff.

Step 7 Click **Save**.

The template and interface configurations form the underlay network configuration on the switches. Also, freeform CLIs that were entered as part of fabric settings (leaf and spine switch freeform configurations entered in the Advanced tab) are deployed.

Configuration Compliance: If the provisioned configurations and switch configurations do not match, the **Status** column displays out-of-sync. For example, if you enable a function on the switch manually through a CLI, then it results in a configuration mismatch.

To ensure configurations provisioned from Cisco Nexus Dashboard Fabric Controller to the fabric are accurate or to detect any deviations (such as out-of-band changes), Nexus Dashboard Fabric Controller's Configuration Compliance engine reports and provides necessary remediation configurations.

When you click **Deploy Config**, the **Config Deployment** window appears.

If the status is out-of-sync, it suggests that there is inconsistency between the Nexus Dashboard Fabric Controller and configuration on the device.

The Re-sync button is displayed for each switch in the Re-sync column. Use this option to resynchronize Nexus Dashboard Fabric Controller state when there is a large scale out-of-band change, or if configuration changes do not register in the Nexus Dashboard Fabric Controller properly. The re-sync operation does a full CC run for the switch and recollects "show run" and "show run all" commands from the switch. When you initiate the re-sync process, a progress message is displayed on the screen. During the re-sync, the running configuration is taken from the switch. The Out-of-Sync/In-Sync status for the switch is recalculated based on the intent defined in Nexus Dashboard Fabric Controller.

Click the **Preview Config** column entry (updated with a specific number of lines). The Config Preview screen comes up.

The **PendingConfig** tab displays the pending configurations for successful deployment.

The **Side-by-side Comparison** tab displays the current configurations and expected configurations together.

Multi-line banner motd configuration can be configured in Cisco Nexus Dashboard Fabric Controller with freeform configuration policy, either per switch using **switch_freeform**, or per fabric using leaf/spine freeform configuration. Note that after the multi-line banner motd is configured, deploy the policy by executing the **Deploy Config** option in the (top right part of the) fabric topology screen. Else, the policy may not be deployed properly on the switch. The **banner** policy is only to configure single-line banner configuration. Also, you can only create one banner related freeform configuration/policy. Multiple policies for configuring banner motd are not supported.

Step 8 Close the **screen**.

After successful configuration provisioning (when all switches display a progress of 100%), close the screen.

The fabric topology is displayed. The switch icons turn green to indicate successful configuration.

If a switch icon is in red color, it indicates that the switch and Nexus Dashboard Fabric Controller configurations are not in sync. When deployment is pending on a switch, the switch is displayed in blue color. The pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Deploy Config** to recompute the state of the switch.

Note
If there are any warning or errors in the CLI execution, a notification will appear in the **Fabric builder** window. Warnings or errors that are auto-resolvable have the **Resolve** option.

An example of the **Deploy Config** option usage is for switch-level freeform configurations. Refer for details.

Adding Switches Using Bootstrap Mechanism

When a new Cisco NX-OS device is powered on, typically that device has no startup configuration or any configuration state for that matter. Consequently, it powers on with NX-OS and post initialization, goes into

a POAP loop. The device starts sending out DHCP requests on all the interfaces that are up including the mgmt0 interface.

Starting from Nexus Dashboard Fabric Controller Release 12.0.1a, POAP access user validated key exchange and password-less ssh to limit configuration file access to the specific switch for a finite time. Therefore, you must accept a new key via **Add Switch > Bootstrap** whenever a device attempts POAP.

If there is IP reachability between the device and the Nexus Dashboard Fabric Controller, the DHCP request from the device, will be forwarded to the Nexus Dashboard Fabric Controller. For easy day-0 device bring-up, the bootstrap options should be enabled in the Fabric Settings.

With bootstrap enabled for the fabric, the DHCP request coming from the device will be serviced by the Nexus Dashboard Fabric Controller. The temporary IP address allocated to the device by the Nexus Dashboard Fabric Controller will be employed to learn basic information about the switch including the device model, device NX-OS version, etc.

1. Choose **LAN > Switches > Add Switches**.
2. Choose **Bootstrap(POAP)** radio button.
3. Click **Actions** and add Switches.

   You can add switches one at a time using the **Add** option or add multiple switches at the same time using the **Import** option.

   If you use the **Add** option, ensure you enter all the required details.

   Note: It might take some time for the switches to appear.
4. Choose a required switch.
5. Click **Edit**.

   The **Edit bootstrap switch** dialog appears.
6. Enter the required details.
7. Click **Save**.
8. Choose the switch.
9. Enter the admin password in the **Admin password** field.
10. Click **Import Selected Switches**.

Return Material Authorization (RMA)

This section describes how to replace a physical switch in a Fabric when using Cisco Nexus Dashboard Fabric Controller Easy Fabric mode.

Prerequisites

- Ensure that the fabric is up and running with minimal disruption while replacing the switch.
- To use the POAP RMA flow, configure the fabric for bootstrap (POAP).
- Perform Recalculate config and Deploy more than once, if needed, to copy the FEX configurations for the RMA of switches that have FEX deployed.

Guidelines and Limitations

- To replace the switch, remove the old switch from the fabric and discover the new switch in the fabric. For example, if you want to replace a Cisco Nexus 9300-EX switch with a Cisco Nexus 9300-FX switch, remove the 9300-EX switch from the fabric followed by discovering the 9300-FX switch in the same fabric.
- When GIR is enabled before upgrading Cisco Nexus 7000 Series switches, Nexus Dashboard Fabric Controller pushes the **system mode maintenance** command to the switches when the Nexus Dashboard Fabric Controller RMA procedure is initiated. This command applies the configuration that is present in the default maintenance mode profile to the switches. For more information on performing Graceful Insertion and Removal (GIR) on the Cisco Nexus 7000 Series switches, refer Configuring GIR.

POAP RMA Flow

To provision RMA, follow below procedure:

Procedure

Step 1 Navigate to the Fabric overview.

Step 2 Move the device into maintenance mode. To move a device into maintenance mode, choose the device, click **Actions** > **More** > **Change Mode**. From the Mode drop-down list, select **Maintenance**.

Step 3 Physically replace the device in the network. Physical connections should be made in the same place on the replacement switch as they existed on the original switch.

Step 4 Power on the switch and go through the POAP cycle.

Step 5 Initiate the RMA flow. Choose the device, click **Actions** > **More** > **Provision RMA**.

Step 6 Set the admin password.

(Optional) You can set a AAA user and password for discovery.

Step 7 Select the replacement device.

Step 8 Click **Provision RMA**.

Manual RMA Flow

Use this flow when Bootstrap is not possible (or not desired).

To provision manual RMA, follow below procedure:

Procedure

Step 1 (Optional) Place the device in maintenance mode.

Step 2 Physically replace the device in the network.

Step 3 Log in through Console and set the Management IP and credentials.

Step 4 If you are using AAA, configure AAA commands on the switch.

Ensure you update LAN and discovery credentials in NDFC for the newly configured AAA user, if configured.

Step 5 The Cisco Nexus Dashboard Fabric Controller rediscovers the new device or you can manually choose **Discovery** > **Rediscover**.

Step 6 Deploy the expected configuration using **Actions** > **Deploy**.

Step 7 Depending on the configuration, if breakout ports or FEX ports are in use, you have to deploy again to completely restore the configuration.

Step 8 After a successful deployment, and the device is "In-Sync," you must move the device back to Normal Mode.

RMA for User with Local Authentication

Note This task is only applicable to non-POAP switches.

Use the following steps to perform RMA for a user with local authentication:

Procedure

Step 1 After the new switch comes online, SSH into the switch and reset the local user passwords with the cleartext password using the "username" command. Reset the local user passwords to resync the SNMP password. The password is stored in the configuration file in a nontransferable form.

Step 2 Wait for the RMA to complete.

Step 3 Update Cisco Nexus Dashboard Fabric Controller `switch_snmp_user` policy for the switch with the new SNMP MD5 key from the switch.

Pre-provisioning Support

Cisco NDFC supports provisioning of device configuration in advance. This is specifically applicable for scenarios where devices have been procured, but not yet delivered or received by the Customers. The purchase order typically has information about the device serial number, device model and so on, which in turn can be used to prepare the device configuration in NDFC prior to the device connectivity to the Network. Pre-provisioning is supported for Cisco NX-OS devices in Data Center VXLAN EVPN, External Connectivity Network, and Classic LAN fabrics.

Pre-provisioning a Device

You can provision devices before adding them to fabrics. However, ensure that you enter DHCP details in the Bootstrap tab in the fabric settings.

The pre-provisioned devices support the following configurations in Nexus Dashboard Fabric Controller:

- Base management
- vPC Pairing
- Intra-Fabric links
- Ethernet ports

- Port-channel
- vPC
- ST FEX
- AA FEX
- Loopback
- Overlay network configurations

The pre-provisioned devices do not support the following configurations in Nexus Dashboard Fabric Controller:

- Inter-Fabric links
- Sub-interface
- Interface breakout configuration

When a device is being pre-provisioned has breakout links, you need to specify the corresponding breakout command along with the switch's model and gateway in the **Data** field in the **Add a new device to pre-provisioning** window in order to generate the breakout PTI.



Note The interface breakout CLI in the **Data** key of the pre-provision payload must contain the exact format as is on the 'show running-configuration' output from the switch.

Note the following guidelines:

- Multiple breakout commands can be separated by a semicolon (;).
- The definitions of the fields in the data JSON object are as follows:
  - **modulesModel**: (Mandatory) Specifies the switch module’s model information.
  - **gateway**: (Mandatory) Specifies the default gateway for the management VRF on the switch. This field is required to create the intent to pre-provision devices. You must enter the gateway even if it is in the same subnet as Nexus Dashboard Fabric Controller to create the intent as part of pre-provisioning a device.
  - **breakout**: (Optional) Specifies the breakout command provided in the switch.
  - **portMode**: (Optional) Specifies the port mode of the breakout interface.

The examples of the values in the **Data** field are as follows:

- {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24"}
- {"modulesModel": ["N9K-C93180LC-EX"],"breakout": "interface breakout module 1 port 1 map 10g-4x", "portMode": "hardware profile portmode 4x100G+28x40G", "gateway": "172.22.31.1/24" }
- {"modulesModel": ["N9K-X9736C-EX", "N9K-X9732C-FX", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-SUP-B+", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.22.31.1/24"}

- {"breakout":"interface breakout module 1 port 50 map 10g-4x" , "gateway": "172.16.1.1/24", "modulesModel": ["N9K-C93180YC-EX "]}
- {"modulesModel": ["N9K-X9732C-EX", "N9K-X9732C-EX", "N9K-C9504-FM-E", "N9K-C9504-FM-E", "N9K-SUP-B", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.29.171.1/24", "breakout":"interface breakout module 1 port 1,11,19 map 10g-4x; interface breakout module 1 port 7 map 25g-4x"}
- {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24","breakout": "interface breakout module 1 port 1-4 map 10g-4x","portMode": "hardware profile portmode 48x25G + 2x100G + 4x40G"}

1. Choose **LAN > Switches> Add Switches**.
2. Choose **Pre-provision** radio button.
3. Click **Actions** and add switches.

   You can add switches one at a time using the **Add** option or add multiple switches at the same time using the **Import** option.

   If you use the **Add** option, ensure you enter all the required details.
4. Choose a switch.
5. Enter the admin password in the **Admin password** field.
6. Click **Pre-provision**.

   The pre-provisioned switch is added.

   To bring in the physical device, you can follow the manual RMA or POAP RMA procedure.

   For more information, see Return Material Authorization (RMA).

   If you use the POAP RMA procedure, ignore the error message of failing to put the device into maintenance mode due to no connectivity since it is expected to have no connectivity to a non-existing device.

Pre-provisioning an Ethernet Interface

You can pre-provision Ethernet interfaces in the **LAN Interfaces** window. This pre-provisioning feature is supported in the Easy, External, and eBGP fabrics. You can add Ethernet interfaces to only pre-provisioned devices before they are discovered in NDFC.

Note Before attaching a network/VRF, you must pre-provision the Ethernet interface before adding it to Port-channels, vPCs, ST FEX, AA FEX, loopback, subinterface, tunnel, ethernet, and SVI configurations.

Before you begin

Make sure that you have a pre-provisioned device in your fabric. For information, see Pre-provisioning a Device, on page 11.

Procedure

Step 1 Double-click on the fabric containing the pre-provisioned device from the **LAN Fabrics** window.

The **Fabric Overview** window appears.

Step 2 On the **Interfaces** tab, click **Actions** > **Create Interface**.

The **Create Interface** window appears.

Step 3 Enter all the required details in the **Create Interface** window.

Type: Select **Ethernet** from the drop-down list.

Select a device: Select the pre-provisioned device.

Note
You cannot add an Ethernet interface to an already managed device .

Interface Name: Enter a valid interface name based on the module type. For example, Ethernet1/1, eth1/1, or e1/1. The interface with same name should be available on the device after it is added.

Policy: Select a policy that should be applied on the interface.

For more information, see Adding Interfaces.

Step 4 Click **Save**.

Step 5 Click **Preview** to check the expected configuration that will be deployed to the switch after it is added.

Note
The **Deploy** button is disabled for Ethernet interfaces since the devices are pre-provisioned.

Pre-provisioning a vPC Pair

Before you begin

Ensure that you have enabled **Bootstrap** in the Fabric Settings.

Procedure

Step 1 Import both the devices into the fabric. For more information, refer Pre-provisioning a Device, on page 11.

Two Cisco Nexus 9000 Series devices that are pre-provisioned and added to an existing Fabric. Choose **Add Switches** from the **Actions** drop-down list. On the Inventory Management screen, click **PowerOn Auto Provisioning (POAP)**.

The devices will show up in the fabric as gray/undiscovered devices.

Step 2 Right click and select appropriate roles for these devices similar to other reachable devices.

Step 3 To create vPC pairing between the devices with physical peer-link or MCT, perform the following steps:

a) Provision the physical Ethernet interfaces that form the peer-link.

The vPC peer-link between leaf1-leaf2 comprises of interfaces Ethernet1/44-45 on each device. Choose **LAN > Fabrics > Interfaces** to pre-provision ethernet interfaces. For more information, see

For instructions, see Pre-provisioning an Ethernet Interface, on page 13.

b) Create a pre-provisioned link between these interfaces.

In the **Links** tab, click on **Actions** > **Create**.

Create two links, one for leaf1-Ethernet1/44 to leaf2-Ethernet1/44 and another one for leaf1-Ethernet1/45 to leaf2-Ethernet1/45.

Ensure that you choose **int_pre_provision_intra_fabric_link** as link template. The Source Interface and Destination Interface field names, must match with the Ethernet interfaces pre-provisioned in the previous step.

After the links are created, they are listed in the **Links** tab under **Fabric Overview** window.

c) On the **Topology** window, right click on a switch and choose **vPC Pairing** from the drop-down list.

Select the vPC pair and click vPC pairing for the pre-provisionsed devices.

d) Click **Recalculate & Deploy** to generate the required intended vPC pairing configuration for the pre-provisioned devices.

After completion, the devices will be correctly paired and the vPC pairing intent will be generated for the devices and the policies are generated

Note
Because the devices are not yet operational, Configuration Compliance will not return any IN-SYNC or OUT-OF-SYNC status for these devices.

This is expected as CC requires the running configuration from the devices in order to compare that with the intent and calculate and report the compliance status.

Pre-provisioning a vPC Host Interface

Procedure

Step 1 Create physical ethernet interfaces on the pre-provisioned devices. Add a vPC host interface similar to a regular vPC pair or switches. For more information, see Pre-provisioning an Ethernet Interface, on page 13.

For example, leaf1-leaf2 represents the pre-provisioned vPC device pair, assuming that Ethernet interfaces 1/1 is already pre-provisioned on both devices leaf1 and leaf2.

Step 2 Create a vPC host truck interface .

Preview and **Deploy** actions doesn't yield any result, because both require the device to be present. The vPC host interface is created and displays status as **Not discovered**.

Attaching Overlays to Pre-provisioned Devices

Overlay VRFs and Networks can be attached to pre-provisioned devices similar to any other discovered device.

An overlay network is attached to the pre-provisioned vPC pair of leafs (leaf1-leaf2). It is also attached to the pre-provisioned vPC host interface port-channels created on leaf1-leaf2.

Preview and **Deploy** operations are disabled for the pre-provisioned devices, because the devices are not reachable. After the pre-provisioned device is reachable, all operations are enabled similar to other discovered devices.

On the **Fabric Overview** window, click the **Policies** tab and choose **Actions** > **Edit Policy**. You can view the entire intent generated for the pre-provisioned device, including the overlay network/VRF attachment information.

Previewing Switches

Nexus Dashboard Fabric Controller UI Navigation

- Choose **LAN** > **Switches**.
- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric Summary** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Switches**.

After adding the switches, you can preview the switches with pending configurations, the side-by-side comparison of running configurations, and the expected configurations for the switches. You can select multiple switches and preview them at the same instance. The **Preview** window displays the pending configurations for the successful deployment of a switch.

To preview the switches and resync the ones with pending configurations, perform the following steps:

Procedure

Step 1 In the **Switches** window, use the check boxes next to the switches to select the switches that you want to preview. From the **Actions** drop-down list, choose **Preview**.

The **Preview Config** window appears. This window displays the switch configuration information such as the switch name; its ip address, role, serial number; the fabric status-whether it is in sync, out of sync, or not available; the pending configuration; the status description; and the progress.

Step 2 To only preview the configuration, view the displayed information and click **Close**.

Step 3 To resynchronize the switches with pending configuration, click **Resync**. The progress bar displays the progress of the resynchronization. Click **Close** to close the **Preview Config** window.

Step 4 To view the pending configurations and side-by-side comparison, click the respective link in the **Pending Config** column.

Alternatively, on the **Fabric Overview Actions** drop-down list, select **Recalculate Config**. The **Deploy Configuration** window appears. It displays the configuration status on the switches. You can also view the pending configurations by clicking the respective link in the **Pending Config** column.

The **Pending Config** window appears. The **Pending Config** tab on this window displays the pending configurations on the switch. The **Side-by-Side Comparison** tab displays the running configuration and expected configuration side-by-side.

Close the **Pending Config** window.

Deploy Configuration

This deploy option is a local operation for a switch, that is, the expected configuration or intent for a switch is evaluated against it's current running configuration, and a config compliance check is performed for the switch to get the **In-Sync** or **Out-of-Sync** status. If the switch is out of sync, the user is provided with a preview of all the configurations running in that particular switch that vary from the intent defined by the user for that respective switch.

1. Choose required switch, choose **Actions** > **Deploy** to deploy configuration on a switch.

   The **Deploy Configuration** window appears.

2. Click **Resync** to synchronize configuration.
3. Click **Deploy**.

   The Status column displays FAILED or SUCCESS state. For a FAILED status, investigate the reason for failure to address the issue.

4. Click **Close** to navigate to switch window.

Discovery

This chapter contains below sections:

Update Credentials

Use update discovery credentials for updating discovering switches.

Procedure

Step 1 Choose required switch, choose **Actions** > **Discovery** > **Update Credentials**.

The **Update Discovery Credentials** window appears.

Step 2 In the **Update Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

Step 3 Click **Update** to save the discovery credentials.

If the discovery credentials are not provided, Nexus Dashboard Fabric Controller uses the admin user and password to discover switches.

Rediscover

You can rediscover switch and check the status of it.

To rediscover the switch:

- Choose required switch, choose **Actions** > **Discovery** > **Rediscover** to rediscover switches.

  The **Discovery Status** column shows the status as **Rediscovering** and after discovering it displays the status.

Guidelines and Limitations for Changing Discovery IP Address

From Cisco Nexus Dashboard Fabric Controller Release 12.0.1a, you can change the Discovery IP address of a device that is existing in a fabric.

Guidelines and Limitations

The following are the guidelines and limitations for changing discovery IP address.

- Changing discovery IP address is supported for NX-OS switches and devices that are discovered over their management interface.
- Changing discovery IP address is supported for templates such as:
  - Easy_Fabric
  - Easy_Fabric_eBGP
  - External
  - LAN_Classic
  - LAN_Monitor
- Changing discovery IP address is supported in both managed and monitored modes.
- Only users with the **network-admin** role can change the discovery IP address on Cisco Fabric Controller UI.
- The discovery IP address must not be used on other devices, and it must be reachable when the change is done.
- While changing the discovery IP address for a device in a managed fabric, switches are placed in migration mode.
- When you change the IP address of a switch that is linked to vPC Peer, corresponding changes such as vPC peer, domain configuration will be updated accordingly.
- Fabric configuration restores the original IP address, it reports out of sync post restore and the configuration intent for the device must be updated manually to get the in-sync status.
- Fabric controllers restore that had the original device discovery IP reports the switch as Unreachable post restore. The discovery IP address change procedure must be repeated after the restore.
- Device Alarms associated with the original discovery IP address will be purged after the change of IP address.

Changing Discovery IP Address

Before you begin

You must make the management IP address and route related changes on the device and ensure that the reachability of the device from Nexus Dashboard Fabric Controller.

To change the discovery IP address from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

Procedure

Step 1 Choose **LAN** > **Fabrics**.

Step 2 Click on fabric names to view the required switch.

The **Fabric summary** slide-in pane appears.

Step 3 Click **Launch** icon to view **Fabric Overview** window.

Step 4 On the **Switches** tab, click **Refresh** icon adjacent to the **Action** button on the main window.

Switch with a changed IP address will be in **Unreachable** state in **Discovery Status** column.

Step 5 Click the check box next to the **Switch** column and select the switch.

Note
You can change the IP address for individual switch and not for multiple switches.

Step 6 Choose **Actions** > **Change Discovery IP** on the switches tab area.

The **Change Discovery IP** window appears.

Similarly, you can navigate from **LAN** > **Switches** tab. Choose a required switch, click **Actions** > **Discovery** > **Change Discovery IP**.

Step 7 Enter the appropriate IP address in the **New IP Address** text field and click **OK**.

a) The new IP address must be reachable from Nexus Dashboard Fabric Controller to update successfully.
b) Repeat the above procedures for the devices where the discovery IP address must be changed before proceeding with further steps.
c) If the fabric is in managed mode, the device mode will be updated to migration mode.

Step 8 From the fabric **Actions** drop-down list, click **Recalculate Config** to initiate the process of updating Nexus Dashboard Fabric Controller configuration intent for the devices. Similarly, you can recalculate configuration on topology window. Choose **Topology**, tab right-click on the switch, click **Recalculate Config**.

The Nexus Dashboard Fabric Controller configuration intent for the device management related configuration will be updated and the device mode status for the switch is changed to normal mode. The switch configuration status is displayed as **In-Sync**.

Note
The PM records associated with the old switch IP address will be purged and new record collections take an hour to initiate after the changes.

Update VRF

To update discovery VRF for switches, perform the following steps:

Note If you enable update VRF option, the VRF associated with the interface which has discovery IP address for a switch will be auto discovered in NDFC during importing a switch. You can override VRF settings for required switch with appropriate user role.

Procedure

Step 1 Choose required switch, choose **Actions** > **Discovery** > **Update VRF**.

The **Update Discovery VRF** window appears.

Step 2 In the **Update Discovery VRF** window, choose **New VRF** and **Interface** from drop-down list.

Step 3 Click **OK** to save new VRF details.

Assigning Switch Roles

You can assign roles to switches on Nexus Dashboard Fabric Controller.

1. Choose required switch, choose **Actions** > **Set Role**.
2. The **Select Role** window appears. You can choose appropriate role and click **Select**.

   A confirmation window appears.

Note You must rediscover the switch to view new role assignation in **Role Status** column.

The following roles are supported in Nexus Dashboard Fabric Controller:

- Spine
- Leaf
- Border
- Border Spine
- Border gateway
- Border gateway spine
- Super spine
- Border super spine
- Border gateway super spine
- Access
- Aggregation

- Edge router
- Core router
- TOR

Creating a vPC Setup

You can create a vPC setup for a pair of switches in the external fabric. Ensure that the switches are of the same role and connected to each other.

Procedure

Step 1 Right-click one of the two designated **vPC switches** and choose **vPC Pairing**.

The **Select vPC peer** dialog box comes up. It contains a list of potential peer switches. Ensure that the **Recommended** column for the vPC peer switch is updated as **true**.

Note
Alternatively, you can also navigate to the **Tabular view** from the **Actions** pane. Choose a switch in the **Switches** tab and click **vPC Pairing** to create, edit, or unpair a vPC pair. However, you can use this option only when you choose a Cisco Nexus switch.

Step 2 Click the radio button next to the vPC peer switch and choose **vpc_pair** from the **vPC Pair Template** drop-down list. Only templates with the **VPC_PAIR** template sub type are listed here.

The **vPC Domain** and **vPC Peerlink** tabs appear. You must fill up the fields in the tabs to create the vPC setup. The description for each field is displayed at the extreme right.

vPC Domain tab: Enter the vPC domain details.

vPC+: If the switch is part of a FabricPath vPC + setup, enable this check box and enter the **FabricPath switch ID** field.

Configure VTEPs: Check this check box to enter the source loopback IP addresses for the two vPC peer VTEPs and the loopback interface secondary IP address for NVE configuration.

NVE interface: Enter the NVE interface. vPC pairing will configure only the source loopback interface. Use the freeform interface manager for additional configuration.

NVE loopback configuration: Enter the IP address with the mask. vPC pairing will only configure primary and secondary IP address for loopback interface. Use the freeform interface manager for additional configuration.

vPC Peerlink tab: Enter the vPC peer-link details.

Switch Port Mode: Choose **trunk** or **access** or **fabricpath**.

If you select trunk, then corresponding fields (**Trunk Allowed VLANs** and **Native VLAN**) are enabled. If you select **access**, then the **Access VLAN** field is enabled. If you select **fabricpath**, then the trunk and access port related fields are disabled.

Step 3 Click **Save**.

The **vPC setup** is created.

To update vPC setup details, do the following:

a. Right-click a vPC switch and choose vPC Pairing.

The **vPC peer** dialog box comes up.

b. Update the field(s) as needed.

When you update a field, the **Unpair** icon changes to **Save**.

c. Click **Save** to complete the update.

After creating a vPC pair, you can view vPC details in **vPC Overview** window.

Undeploying a vPC Setup

Procedure

Step 1 Right-click a **vPC** switch and choose **vPC Pairing**.

The vPC peer screen comes up.

Step 2 Click **Unpair** at the bottom right part of the screen.

The vPC pair is deleted and the fabric topology window appears.

Step 3 Click **Deploy Config**.

Step 4 (Optional) Click the value under the **Recalculate Config** column.

View the pending configuration in the **Config Preview** dialog box. The following configuration details are deleted on the switch when you unpair: vPC feature, vPC domain, vPC peerlink, vPC peerlink member ports, loopback secondary IPs, and host vPCs. However, the host vPCs and port channels are not removed. Delete these port channels from the **Interfaces** window if required.

Note
Resync the fabric if it is out of sync.

When you unpair, only PTIs are deleted for following features, but the configuration is not cleared on the switch during **Deploy Config**: NVE configuration, LACP feature, fabricpath feature, nv overlay feature, loopback primary ID. In case of host vPCs, port channels and their member ports are not cleared. You can delete these port channels from the **Interfaces** window if required. You can continue using these features on the switch even after unpairing.

If you are migrating from fabricpath to VXLAN, you need to clear the configuration on the device before deploying the VXLAN configuration.

Performing Actions on Switches

Change Mode

To change mode for the switch, perform the following steps:

1. Choose check box for required switch, choose **Actions** > **More** > **Change Mode**.

   The **Change Mode** window appears.

2. Choose required **Normal** or **Maintenance** from drop-down list.
3. Click **Save and Deploy Now** to change mode or click **Save and Deploy Later** to change mode later.

Provision RMA

To change mode for the switch, perform the following steps:

1. Choose check box for required switch, choose **Actions** > **More** > **Provision RMA**.

   The **Provision RMA** window appears.

2. The Provision RMA UI will show the replacement device 5-10 minutes after it is powered on.

Change Serial Number

Allows you to change the serial number of switches. While pre-provisioning devices, you can provide dummy values for the Serial number of the switch. After you configure the network successfully, you can change the serial number with the appropriate serial number of the switch. Before changing the serial number of switches, on main window, click **Actions > Recalculate and Deploy** to save the latest data on switch.



Note The change of serial number is supported only for Nexus 9000 Series switches. After change of serial number with actual number it is recommended to Re-POAP the device during the power on bootstrap

Copy Run Start

To copy the existing switch configuration to start configuration, perform the following steps:

1. Choose check box for required switch, choose **Actions** > **More** > **Copy Run Start**.

   The **Copy Running Config to Startup Config** screen appears. In the Progress column shows the process in progress and status description shows **Deployment in progress**.

2. A confirmation window appears, click **OK**.

   The status description column displays **Deployment completed** and progress column in green.

3. Click **Close** to close this window.

Reload

To reload required switch, choose **Actions** > **More** > **Reload**.

A confirmation window appears, click **Confirm**.

Restore Switch

You can restore a Cisco Nexus switch in external fabrics and LAN classic fabrics from the Cisco Nexus Dashboard Fabric Controller Web UI. The information you restore at switch-level is extracted from the fabric-level backups. The switch-level restoring doesn't restore fabric-level intents and other configurations applied using the fabric settings. Only switch-level intents are restored. Therefore, after you restore a switch, it might go out-of-sync because the fabric-level intents aren't restored. Perform a fabric-level restore to restore the intents as well. You can restore only one switch at a time. You can't restore a switch if the fabric where it's discovered is part of an MSD fabric.

1. Choose **Actions** > **More** > **Reload**.

   The **Restore Switch** window appears and you are in the **Select a Backup** tab. Refer to Backup Fabric for more information.

2. The **Select a Backup** tab displays the fabric backup details. It includes the following information:

   - Backup Date - Specifies the backup date and time.
   - Backup Version - Specifies the version number of backup.
   - Backup Tag - Specifies the name of backup.
   - NDFC Version - Specifies the NDFC version details.
   - Backup Type - Specifies the type of backup, either manual or automatic.

     You can choose the automatic, manual, or golden backup. These backups are color-coded. Automatic backups are indicated in blue color.

     Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name.

     You can now mark the backups that you don't want to delete even after you reach the archiving limit. These backups are the golden backups. You can't delete golden backups of fabrics. However, Cisco Nexus Dashboard Fabric Controller archives only up to 10 golden backups.

3. Choose radio button for required backup to mark as golden, choose **Actions > Mark as golden**, a confirmation window appears, click **Confirm**.
4. Choose radio button for backup to delete from golden, choose **Actions > Remove as golden**, a confirmation window appears, click **Confirm**.

Note Most of this information is at the fabric level, and may or may not directly impact the proceedings of the switch-level restore.

5. Click **Next** to move to the **Restore Preview** step.
6. You can view information about the switch name, switch serial, IP address, status, restore supported, delta configuration and the VRF details.
7. (Optional) Click **Get Config** to preview device configuration details.

   The **Config Preview** window appears, which has three tabs.

- **Backup Config**: This tab displays the backup configuration for the selected device.
- **Current Config**: This tab displays the current running configuration of the selected device.
- **Side-by-side Comparison**: This tab displays current running configuration on the switch, and the backup configuration, which is the expected configuration.

8. Click **Restore Intent** to proceed to the **Restore Status** step in restoring.

   The restore status and description appears for the switch.

9. Click **Finish** after the restoring process is complete.



Note

- You can't go back to the previous step because the fabric configurations change.
- If the restoring failed, the switch rolls back to the previous configuration.

Show Commands

The following procedure view the commands in Nexus Dashboard Fabric Controller:

1. Choose **Actions** > **More** > **Show commands**.

   The **Switch Show Commands** window appears.

2. Choose required commands from drop-down list and enter required information in text field.
3. Click **Execute** to view the CLI output and to clear the output, click **Clear Output**.

Exec Commands

The commands available in the EXEC mode include the show commands that display the device status and configuration information, the clear commands, and other commands that perform actions that you do not save in the device configuration.

The following procedure shows how to run EXEC commands in Nexus Dashboard Fabric Controller:

1. Choose **Actions** > **More** > **Exec commands**.

   The **Switch Show Commands** window appears.

2. From the **Template** drop-down list, choose **exec_freeform** or **exec_elam_capture**.
3. Enter the commands in the **Freeform CLI** for **exec_freeform** and required IP addresses.
4. Click **Deploy** to run the EXEC commands.
5. In the **CLI Execution Status** window, you can check the status of the deployment. Click **Detailed Status** under the **Command** column to view details.
6. In the **Command Execution Details** window, click the info under the **CLI Response** column to view the output or response.

Delete Switches

You can delete one or more existing switches.

Choose **Actions** > **More** > **Delete switche(s)**. A confirmation window appears, click **Confirm**

Switch Overview

You can perform below operations, from **Actions** icon on Switch Overview window:

- Preview
- Deploy
- Discovery
- Set Roles
- vPC Pairing
- More

Viewing Switch Overview

You can view information about switch along with the switch summary on **Switch Overview** tab. Navigate **LAN** > **Switches**, click on required switch. A slide-in pane appears. Click **Launch** icon to view the **Switch Overview** window.

Field	Description
Switch Info	Specifies the switch information such as switch name, IP address, switch model and other details.
Alarms	Specifies the alarms configured on the selected switch
Performance	Specifies the CPU utilization and memory utilization for the switch.
Interfaces	Specifies the interface details.
Modules/FEX	Specifies the modules and FEX information.
Reports	Specifies the reports.

Hardware

This tab contains below sections:

Modules

To view the inventory information for modules from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

Procedure

Step 1 Choose **LAN > Switch > Switch Overview > Hardware > Modules**.

The **Modules** tab is displayed with a list of all the switches and its details for a selected Scope.

You can view required information in table, enter details in **Filter by Attributes**.

Step 2 You can view the following information.

- **Name** displays the module name.
- **Model** displays the model name.
- **Serial Number** column displays the serial number.
- **Type** column displays the type of the module.
- **Oper. Status** column displays the operation status of the module.
- **Slot** column displays the slot number.
- **HW Revision** column displays the hardware version of the module.
- **Software Revision** column displays the software version of the module.
- **Asset ID** column displays the asset id of the module.

Viewing Bootflash

You can view the following information on Bootflash tab.

- **Primary Bootflash Summary** card displays the total, used and available space.
- **Secondary Bootflash Summary** card displays the total, used and available space.
- **Directory Listing** area displays check box for **Primary Bootflash** and **Secondary Bootflash**.

  This area shows the filename, size, and last modified date for all the files and directories on the switch bootflash. Choose **Actions > Delete** to delete files to increase the available space on the switch.

Links

You can add links between border switches of different fabrics (inter-fabric links) or between switches in the same fabric (intra-fabric links). You can only create an inter-fabric connection (IFC) for a switch that is managed by Nexus Dashboard Fabric Controller.

There are scenarios where you might want to define links between switches before connecting them physically. The links could be inter-fabric or intra-fabric links. Doing so, you can express and represent your intent to add links. The links with intent are displayed in a different color till they are actually converted to functional links. Once you physically connect the links, they are displayed as connected.

Management links might show up in the fabric topology as red colored links. To remove such links, right-click the link and click **Delete Link**.

The Border Spine and Border Gateway Spine roles are added to switch roles for border switches.

You can create links between existing and pre-provisioned devices as well by selecting the pre-provisioned device as the destination device.

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

Field	Description
Fabric Name	Specifies the name of the Fabric.
Name	Specifies the name of the link. The list of previously created links is displayed. The list contains intra-fabric links, which are between switches within a fabric, and inter-fabric links, which are between border switches in this fabric and switches in other fabrics.
Policy	Specifies the link policy.
Info	Provides more information about the link.
Admin State	Displays the administrative state of the link.
Oper State	Displays the operational state of the link.

The following table describes the action items, in the Actions menu drop-down list, that appear on **Fabric Overview > Links > Links**.

Action Item	Description
Create	Allows you to create the following links: • Creating Inter-Fabric Links • Creating Intra-Fabric Links
Edit	Allows you to edit the selected fabric.
Delete	Allows you to delete the selected fabric.
Import	You can import a CSV file containing details of links to add new links to the fabric. The CSV file should have the following details of links: link template, source fabric, destination fabric, source device, destination device, source switch name, destination switch name, source interface, destination interface, and nvPairs. **Note** • You cannot update existing links. • The Import Links icon is disabled for external fabric.

Action Item	Description
Export	Choose the link and select **Export** to export the links in a CSV file. The following details of links are exported: link template, source fabric, destination fabric, source device, destination device, source switch name, destination switch name, source interface, destination interface, and nvPairs. The nvPairs field consists JSON object.

Protocol View

This tab displays the protocols for the links in the selected Fabric.

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

Field	Description
Fabric Name	Specifies the name of the fabric.
Name	Specifies the name of the link.
Is Present	Specifies if the link is present.
Link Type	Specifies the type of link.
Link State	Specifies the state of link.
UpTime	Specifies the time duration from when the link was up.

PTP (Monitoring)

Note From Release 12.1.1e, you can enable PTP feature from Feature Management for IPFM and Classic LAN Fabrics.

UI Navigation

- Choose **LAN** > **Switches**. Click on a switch to open the **Switch** slide-in pane. Click the **Launch** icon. Choose **Switch Overview** > **PTP**.
- Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Switches**. Double-click a switch to open **Switch Overview** > **PTP**.
- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric Summary** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Switches**. Click on a switch to open the **Switch** slide-in pane and then click the **Launch** icon. Alternatively, you can double-click a switch to open **Switch Overview**. Choose **Switch Overview** > **PTP**.

This section explains the preview functionality of the Precision Time Protocol (PTP) monitoring. PTP is a time synchronization protocol for nodes that are distributed across a network. On a local area network, it achieves clock accuracy in the sub-nanosecond range, making it suitable for measurement and control systems.

On the **PTP** tab in the **Switch Overview** window, you can view PTP-related information based on the selected switch. You can click the **Telemetry Switch Sync Status** link to check whether the switches are in sync. The **Sync Status** column displays the status of the switches.

The following tabs are displayed in this window:

- Port Status
- Correction & Mean Path Delay
- Clock Status

Port Status

The **Port Status** table displays the status of the ports in two different views. Click on the arrow at the top-right corner to view/minimize different views.

Topology view displays the network diagram of the selected port. Hover the pointer over the switch icon to view information about role, ports and number of PTP ports. The following information is displayed:

- switch name
- type of port
    - specifies if it's a follower or leader port
    - if it's a follower port, displays the Leader and Grand Leader ports
- interface connecting to the leader
- number of PTP ports

Tabular view displays all the interfaces on the switch, peer link, and admin, operational and port status for the selected switch.

Click on the Filter by attributes field and choose the required attribute and enter a criteria to filter the port status and press **ENTER**.

Correction and Mean Path Delay

The **Correction & Mean Path Delay** tab displays a graph showing the PTP operational statistics: mean path delay, correction, and correction beyond threshold. You can click and drag in the plot area to zoom in and hold the **shift** key to pan. Click the **Reset zoom** button to reset zoom.

By default, the graph is displayed for the threshold value of 500 nanoseconds (ns). You can also display data based on a specific threshold value. In the **Threshold (ns)** field, enter the required value in nanoseconds and click **Apply**. Note that the threshold value is persistent in the Nexus Dashboard Fabric Controller settings, and it is used to generate PTP correction threshold Kafka notifications.

In the **Date** field, you can select the appropriate date to view the data. The PTP data is stored up to the last seven (7) days. The default value for the stored data is 7 days. To change this value, navigate to **Settings > Server Settings > IPFM** and set the updated value for the **IPFM history retention in days** field.

In the **Period** field, you can also select a timeframe over which the data has to be displayed. The values you can choose in the **Period** field are Hour (1 hour), 6 hours, 12 hours, or Day (24 hours).

Note that you can click the legends in the graph to hide or display statistics.

If there are any corrections, you can view them in a tabular format by clicking the **Corrections Beyond Threshold** link.

To perform a refresh, click the **Refresh** icon.

Clock Status and Port Status

The **Clock Status** tab displays information about the Leader Clock and the Grand Leader Clock.

The **Port Status** table displays the status of the ports. Click on the **Filter by attributes** field and choose the required attribute, and enter a criteria to filter the port status and press ENTER.

Policies

Nexus Dashboard Fabric Controller provides the ability to group a set of switches, and allows you to push a set of underlay configurations to the group.

From Cisco NDFC Release 12.1.1e, follow the below navigation path:

- Choose **LAN** > **Fabrics**. Click on a fabric to open the **Fabric** slide-in pane. Click the **Launch** icon. Choose **Fabric Overview** > **Policies**.

  Choose **LAN** > **Fabrics**. Double-click on a fabric to open **Fabric Overview** > **Policies**.

The following table describes the fields that appear on **Fabric Overview > Policies**.

Field	Description
Policy ID	Specifies the policy ID.
Switch	Specifies the switch name.
IP Address	Specifies the IP address of the switch.
Template	Specifies the name of the template.
Description	Specifies the description. **Note** From Cisco NDFC Release 12.1.1e, change of serial number for the switch is allowed, both old and new serial numbers can be viewed in this column.
Entity Name	Specifies the entity name.
Entity Type	Specifies the entity type.
Source	Specifies the source.
Priority	Specifies the priority.
Content Type	Species for the content type.

Field	Description
Fabric Name	Specifies the fabric name.
Serial Number	Specifies the serial number of the switch.
Editable	Specifies a Boolean value to indicate if the policy is editable.
Mark Deleted	Specifies a Boolean value to indicate if the policy is marked to be deleted.

The following table describes the action items, in the **Actions** menu drop-down list, that appear on **LAN > Policies**.

Action Item	Description
Add Policy	
Edit Policy	Choose a policy from the table and choose **Edit Policy** to modify the policy. **Note** • The policies in the italics font cannot be edited. The value under the **Editable** and **Mark Deleted** columns for these policies is false. • A warning appears when you edit a policy whose **Mark Deleted** value is set to *true*. The switch freeform child policies of **Mark Deleted** policies appears in the Policies dialog box. You can edit only **Python** switch_freeform policies. You cannot edit **Template_CLI** switch_freeform_config policies.
Delete Policy	Choose policies from the table and choose **Delete Policy** to delete the policies. **Note** A warning appears when you delete policies whose **Mark Deleted** values are set to *true*. When you delete the policies whose **Mark Deleted** values are set to *true*, these entries are deleted from the NDFC database only but the configs are not deployed to the switch.
Generated Config	Choose policies from the table and choose **Generated Config** to view the delta of configuration changes made by every user.

Action Item	Description
Push Config	Choose policies from the table and choose **Push Config** to push the policy configuration to the device. **Note** • This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric. • A warning appears if you push configuration for a Python policy. • A warning appears when you push configurations for policies whose **Mark Deleted** values are set to *true*.

Adding a Policy

To add a policy, perform the following steps:

Procedure

Step 1 Choose **LAN** > **Fabrics**. Double-click on the required fabric.

The **Fabric Overview** window appears.

Step 2 On the **Policies** tab, choose **Actions > Add Policy**.

The **Create Policy** window appears.

Step 3 Click and choose required switch and click **Select**.

You must deploy the switch in a pending state.

Step 4 Click **Choose Template** and choose appropriate policy template and click **Select**.

Step 5 Specify a priority for the policy.

The applicable values are from 1 to 1000. The default value is 500. The lower number in the **Priority** field means that there is a higher priority for the generated configuration and POAP startup-configuration. For example, features are 50, route-maps are 100, and vpc-domain is 200.

Step 6 Enter the mandatory parameters in the text field and click **Save**.

Event Analytics

Event Analytics includes the following topics:

History

The history tab displays information about the deployment and policy change history. Choose **LAN > Fabrics**. Double-click a fabric name to open the **Fabric Overview** window and then click the **History** tab.

Resources

Cisco Nexus Dashboard Fabric Controller allows you to manage the resources. The following table describes the fields that appear on this page.

Field	Description
Scope Type	Specifies the scope level at which the resources are managed. The scope types can be **Fabric**, **Device**, **Device Interface**, **Device Pair**, and **Link**.
Scope	Specifies the resource usage scope. Valid values are the switch serial numbers or fabric names. Resources with serial numbers are unique and can be used on the serial number of the switch only.
Device Name	Specifies the name of the device.
Device IP	Specifies the IP address of the device.
Allocated Resource	Specifies if the resources are managed with device, device interface, or fabric. Valid values are ID type, subnet, or IP addresses.
Allocated To	Specifies the entity name for which the resource is allocated.
Resource Type	Specifies the resource type. The valid values are **TOP_DOWN_VRF_LAN**, **TOP_DOWN_NETWORK_VLAN**, **LOOPBACK_ID**, **VPC_ID**, and so on.
Is Allocated?	Specifies if the resource is allocated or not. The value is set to **True** if the resource is permanently allocated to the given entity. The value is set to **False** if the resource is reserved for an entity and not permanently allocated.
Allocated On	Specifies the date and time of the resource allocation.
ID	Specifies the ID.

L4-L7 Services Configuration

Cisco Nexus Dashboard Fabric Controller introduces the ability to insert Layer 4-Layer 7 (L4-L7) service devices in a data center fabric, and also enables selectively redirecting traffic to these L4-L7 service devices. You can add a L4-L7 service node, create route peering between the L4-L7 service node and the L4-L7 service leaf switch, and then selectively redirect traffic to these L4-L7 service nodes.