Working with Configuration Policies for Your Nexus Dashboard LAN or IPFM Fabrics, Release 4.2.1

Table of Contents

New and changed information

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

Release Version	Feature	Description
Nexus Dashboard 4.2.1	Support for advanced DLB configuration	Beginning with Nexus Dashboard 4.2.1, you can apply Dynamic Load Balancing (DLB) configuration at fabric level using the **Apply Fabric Level Setting** option. Nexus Dashboard now supports the **Dynamic_Load_Balancing_S1** policy templates for Silicon One switches, in addition to the **Dynamic_Load_Balancing_CS** policy template for the CloudScale platform. For more information, see Add a Dynamic Load Balancing (DLB) policy template.

Navigate to the Configuration policies page

Follow these steps to navigate to the **Configuration policies** page to view or edit information on configuration policies.

1. Click **Manage > Fabrics**.
2. Click the appropriate fabric on the **Fabrics** page.
3. Click the **Configuration policies** tab.

Policies

Nexus Dashboard manages configuration of devices using Policies. Nexus Dashboard policies are way of grouping all the required CLIs and variables to achieve certain configuration on the devices. These policies can be defined either using CLI commands or Python scripts. Nexus Dashboard generates the configuration for a device based on the policies attached to the device.

Nexus Dashboard provides the ability to create policy groups which can be applied to multiple switches. Policy groups let you create policies that define specific switch parameters that are common to switches and apply them to multiple switches in a fabric.

Dynamic load balancing policy template

When you create a policy, you can choose the Dynamic_Load_Balancing policy template, which enables dynamic load balancing for Layer 3 equal cost multi-path (ECMP) routing group members. Dynamic load balancing adjusts the traffic allocations according to congestion levels of the outgoing links. It measures the congestion across the available paths and places the flows on the least congested paths, which results in an optimal or near-optimal placement of the data.

Supported switches for dynamic load balancing

You can configure dynamic load balancing for these switches:

- N9K-C9348GC-FX3
- N9K-C9348GC-FX3PH
- N9K-C93108TC-FX3
- N9K-C93108TC-FX3P
- N9K-C93180YC-FX3
- N9K-C9316D-GX
- N9K-C93600CD-GX
- N9K-C9364C-GX
- N9K-C9332D-GX2B
- N9K-C9348D-GX2A
- N9K-C9364D-GX2A
- N9K-C9332D-H2R
- N9K-C93400LD-H1

Guidelines and limitations for dynamic load balancing

These limitations apply for dynamic load balancing.

- You can have only one instance of the Dynamic_Load_Balancing policy template. You must modify an existing template instance to make any changes.
- The template does not inherit any fabric-level settings. Dynamic load balancing configuration deployment is identical regardless of the fabric type.

- If you configured dynamic load balancing out-of-band, then the corresponding configuration appears in freeform.
- When you add a new DLB policy or edit an existing DLB policy to change the DLB interface, you must copy the running configuration to the startup configuration on the switch and reload the switch. For more information on copying the running configuration to the startup configuration, see the section " Copy run start" in Perform actions on switches.

Access the Policies page

Follow these steps to access the **Policies** page.

1. Navigate to the Configuration policies page.
2. Click the **Policies** tab.

   The following table describes the fields that appear on the **Policies** page.

Field	Description
Template	Specifies the name of the policy template.
Description	Specifies the description, if available.  Because a change of the serial number for the switch is allowed, you can see both old and new serial numbers in this column.
Content type	Specifies for the template content type. The supported content types are **TEMPLATE_CLI**, **PYTHON**, and **PYTHON_CLI**.
Switch	Specifies the name of the switch the policy has been applied to. If you are configuring a policy group, this field provides a link that specifies the number of switches that are linked to the policy. Click the link to open the policy group details dialog box with details such as the number of switches that are linked to the policy, the IP address, the fabric name, serial number, and mark deleted state.
Entity name	Specifies the switch or the interface name to which the policy has been applied to.
Entity type	Specifies if the entity is a switch or an interface.
Source	Specifies the source.

Priority	Specifies the priority. During an **Edit membership** operation to remove one or more switches from an existing policy group which uses **Template_CLI** content type, the **Priority** column displays the value **Mixed** indicating that the policy group has mixed priorities and mark deleted states. Whereas, when you edit switch_freeform policies of **Content Type PYTHON** (where multiple CLI policy templates are combined with a common source), after an edit operation, Nexus Dashboard removes the occurrence of the switch from the source policy and displays the source and the child policies as different entries. The **Mark Deleted** value for these switches in a child policy indicates the value **true** and the **Priority** indicates a negative value. For a policy group, click on the link to view the group policy details of all the associated switches.
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 column displays **true** indicating that the policy is marked for deletion. All the configurations for a policy with the **Mark Deleted** value **true** will be negated. The **Generated Config** for the policy displays the configuration to be removed from the switch. For a policy group, click on the link to view the group policy details of all the associated switches.
Policy ID	Specifies the policy ID. The policy ID for a policy group begins with the term **POLICY-GROUP**. While searching for a policy group, you can filter the policy ID using this term.
IP address	Specifies the IP address of the switch. If you are configuring a policy group, this field provides a link that specifies the IP addresses for the number of switches that are linked to the policy. Click on the link to view the group policy details of all the associated switches.
Serial number	Specifies the serial number of the switch. If you are configuring a policy group, this field provides a link that specifies the serial numbers for the switches that are linked to the policy. Click on the link to view the group policy details of all the associated switches.
Created on	Specifies the date the policy was created.
Modified on	Specifies the date the policy was modified.

This table describes the action items, in the **Actions** drop-down list, that appear on the **Policies** tab.

Action Item	Description

Add policy	Allows you to create the following types of policies: • Regular policies. To add a regular policy, see Add a policy. • Policy Group. To add a policy group, see Create a policy group.
Edit policy	To modify the policy, choose a policy from the table and choose **Edit policy**.  The policies in italics cannot be edited. The value under the **Editable** and **Mark Deleted** columns for these policies will indicate **false**. You cannot perform **Edit policy** for policies 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.
Edit membership	Lets you edit the membership configuration for a policy group. You can add or remove switches from a policy group using this option. If you remove switches from a policy group, the **Switch** column in the policy details table still displays the original number of switches in the policy. However the **Mark Deleted** column in the details view dialog box displays **true**. You cannot immediately edit a policy after an edit membership operation. The system displays an error message indicating to deploy the pending membership configuration changes to the switch before proceeding with any other configuration changes. If you choose to not deploy the membership configuration changes and would like to edit the policy, ensure you perform **Preview** and proceed with the edit policy operation. Preview operation removes any pending configuration changes from the system. You cannot perform **Edit membership** for policies whose **Mark Deleted** value is set to **true**.

Delete policy	To delete policies, choose the policies from the table and choose **Delete policy**. The following are the points to consider while deleting group policies: • For **TEMPLATE_CLI** policies, removing a policy group removes all the child policies from the switch. • For **Python** policies which has a source and multiple child policies, removing a policy group removes the source policy template instance (PTI) from the switch and displays only the child policies. The system shows the **Generated Config** as negative for both the child policies. You cannot delete the child polices without deploying the configuration. The child policies are deleted automatically after deploying all the pending configuration. A warning appears when you delete policies whose **Mark Deleted** values are set to **true**. Deleting a **TEMPLATE_CLI** policy removes the policy directly from the switch and sets the **Mark Deleted** value to **true**. When you delete policies whose **Mark Deleted** values are set to **true**, these entries are only removed from the Nexus Dashboard database; the configs are not deployed to the switch. These policies do not have any intent and hence you need not deploy the config to the switch.
Generated Config	To view the delta of configuration changes made by every user, choose policies from the table and choose **Generated Config**.
Push Config	To apply the policy configuration to the device, choose policies from the table and choose **Push Config**. This option is grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric. A warning appears if you apply the configuration for a Python policy. You cannot perform a **Push Config** for policies whose **Mark Deleted** value is set to **true**.

Add a policy

Follow these steps to add a policy.

1. Navigate to the Configuration policies page.
2. Click the **Policies** tab.
3. On the **Policies** page, choose **Actions > Add policy**.

   The **Create Policy** page appears.
4. Choose the required switches and click **Next**.

   You must deploy the switch in pending state.

5. Click **No Policy Selected**, choose the appropriate policy template, and click **Select**.

   You can enable or disable PTP high-correction notification when the system encounters a high-correction event. Whenever the correction value exceeds the configured value then that correction is called a high-correction. By default, a high-correction notification is disabled. Enable it manually to generate the notification. Perform the following steps to enable the high-correction notification:

   a. Put a check in the **Enable PTP Telemetry** check box to enable telemetry for PTP.
   b. Put a check in the **Is Large-Scale Fabric?** check box to generate the high-correction notification.

      If there are more than 35 devices in a fabric, PTP events will be used if the switch version is 9.3(5) or higher, or else PTP correction data will be pushed periodically.

   c. Enter the wait time between two successive notifications in the **PTP High-Correction Interval** field.

      The duration value is in seconds.

   d. Set the correction range threshold value (ns) in the **PTP Correction Range** field.

      The default is 100000 (100us).

6. Enter the priority value for the policy in the **Priority** field.

   The applicable values are from 1 to 1000. The default value is 500. A lower number in the **Priority** field indicates 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.

7. If you chose the **ipv4_prefix_list** or **ipv6_prefix_list** policy template, perform these steps to include the prefix-list entries.
   a. Enter the required name in the **Prefix List Name** field.
   b. On the **Prefix-list Entries** card, click **ActionsAdd**.

      The **Add Item** page appears.

   c. Configure the mandatory fields on the **Add Item** dialog box and click **Save**.
   d. Repeat this step to add the required number of prefix-list entries.

      

      The value in the **Sequence Number** must be higher than the previous prefix-list entry. If not, an error message is displayed.

   e. Select the appropriate prefix-list entry and click **Actions > Insert Above** to insert a new prefix-list entry.

      

      The value in the **Sequence Number** must be lower than the below prefix-list entry. If not, an error message is displayed.

8. If you chose the **Dynamic_Load_Balancing** policy template, perform these steps:
   a. For **DLB Interfaces**, specify the interfaces to use for dynamic load balancing.

b. For **DLB MAC Address**, specify the MAC address that is shared by the dynamic load balancing interfaces.

c. If you want to use per-packet load balancing, put a check in the **Per Packet Load Balancing** check box.

d. If you want to use static pinning, in the **Static Pinning** area, choose **Actions > Add**, fill out the port fields, and click **Save**. Repeat this step for each source port and destination port that you want to use for static pinning. You cannot use static pinning if you enabled per-packet load balancing.

e. For **Flowlet Aging**, specify the aging period in microseconds.

f. For the various **DRE Threshold Level** fields, enter the threshold for each level. The thresholds must total 100.

9. For all other policy templates, fill out the fields as necessary.
10. Click **Save**.

Add a Dynamic Load Balancing (DLB) policy template

With this release, you can apply DLB configuration at the fabric level using the **Apply Fabric Level Setting** option. Nexus Dashboard now supports the **Dynamic_Load_Balancing_S1** policy templates for Silicon One switches, including the N93C64E-SG2-Q and N9364E-SG2-O models, in addition to the **Dynamic_Load_Balancing_CS** policy template for the CloudScale platform.

The **Dynamic_Load_Balancing_mode** configuration for **Dynamic_Load_Balancing_S1** template configuration supports three policy-driven modes. The **Dynamic_Load_Balancing_CS** template configuration now supports two additional modes, **policy driven per packet** and **policy driven flowlet**.

For more information on Dynamic Load Balancing configuration details on the Silicon One switches and the CloudScale switches, see the Dynamic Load Balancing on Silicon One switches and Dynamic Load Balancing on CloudScale switches documents.

Follow these steps to add a dynamic load balancing policy template.

1. Navigate to the Configuration policies page.
2. Click the **Policies** tab.
3. On the **Policies** page, choose **Actions > Add policy**.

   The **Create Policy** page appears.

4. Choose the required switches and click **Next**.

   You must deploy the switch in pending state.

5. Enter the priority value for the policy in the **Priority** field.

   The applicable values are from 1 to 2000. The default value is 500. A lower number in the **Priority** field indicates 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.

6. (Optional) Enter a description for your template in the **Description** field.

7. Click **No Policy Selected >** in the **Select Template** field to choose the appropriate policy template, and click **Select**.
   a. Choose one of these **Dynamic_Load_Balancing** policy templates.
      - **Dynamic_Load_Balancing_CS**—applies ISL configuration for for CloudScale.
      - **Dynamic_Load_Balancing_S1**—applies the ISL configuration for G200.
   b. Specify the interfaces to use for dynamic load balancing in the **DLB Interfaces** field.

      

      This DLB Interfaces field is disabled if you choose to the apply the DLB policy template configuration at fabric-level using the option.

   c. For **DLB MAC Address**, specify the MAC address that is shared by the dynamic load balancing interfaces.
   d. Choose a mode from the **Dynamic_Load_Balancing_mode** drop-down list.
      - **flowlet** —allows load balancing to occur at flowlet level based on port load. This is the default mode when DLB is enabled.
      - **per packet** —allows load balancing decision to occur at a per-packet level instead of the flowlet level.
      - **policy driven flowlet**—uses network policies to determine when and how flowlet-based load balancing is applied to traffic. When the mode is not matching the policy, applied interfaces go to err-disabled state.
      - **policy driven per packet** —forwards packets based on policies that allow load balancing decisions to be made on a per-packet basis. When the mode is not matching the policy, applied interfaces go to err-disabled state.
      - **policy driven mixed mode** —allows to use both per-flow and per-packet load balancing methods, based on defined policies.

        

        The **policy driven mixed mode** option is only supported on the **Dynamic_Load_Balancing_S1** policy templates.

   e. Choose an option for **DLB Mixed Mode Default**. Default load balancing mode for policy driven mixed mode DLB. The DLB Mixed Mode Default field is inactive for other modes, ecmp, flowlet, and per packet are the options for this field.
   f. Specify the aging period in the **Flowlet aging time (in microseconds)** field. The default value is 256.
   g. If you want to use per-packet load balancing, put a check in the **Per Packet Load Balancing** check box.
   h. If you want to use static pinning, in the **Static Pinning** area, choose **Actions > Add**, fill out the port fields, and click **Save**. Repeat this step for each source port and destination port that you want to use for static pinning. You cannot use static pinning if you enabled per-packet load balancing.
   i. For the various **DRE Threshold Level** fields, enter the threshold for each level. The thresholds must total 100.
   j. **Enable adaptive routing**--allows the network to automatically adjust the forwarding path of

packets based on real-time network conditions, such as link congestion or failures.

The following additional configuration fields are not applicable to the **Dynamic_Load_Balancing_CS** policy template.

k. **Decay factor**--decay factor determines how much weight is given to recent traffic measurements compared to historical data when calculating load. A higher decay factor means recent measurements have a stronger influence, allowing the system to respond more quickly to changes in network load.

l. **Sampling interval (in nanoseconds)**--sampling interval specifies how often the system collects traffic statistics to assess link utilization and make load balancing decisions.

m. **Load awareness**-- when load awareness is enabled the system makes routing decisions based on current network load, distributing traffic more evenly and avoiding congestion. .(Optional) Toggle the **Group** field to enable the policy for the policy group.

8. Check the **Apply Fabric Level Setting** check-box to uniformly apply switch configuration on all the switches in the fabric.
9. Click **Save**.

What's next: You can enable DLB configuration for AI fabrics, see the "AI settings" section in Editing AI Data Center VXLAN Fabric Settings for more information. Similarly, you can enable DLB configuration for VXLAN fabrics as described in **Advanced** settings section in Fabric Management (for fabrics with iBGP overlay routing protocol) navigate to Advance settings .

Create a policy group

Policy groups provides a method of configuring and managing switches collectively. This feature enables you to create group polices for switches that share common configurations. You can create a policy group and add multiple switches to the policy at the time of creating the policy group or later. Similarly, policy groups also let you edit or delete policies for multiple switches simultaneously.

Follow these steps to create a policy.

1. Navigate to the Configuration policies page.
2. Click the **Policies** tab.
3. On the **Policies** page, choose **Actions > Add policy**.

   The **Create Policy** page appears.

4. To create a policy group, select the required switches to which you need to apply the policy and click **Next**.

   Ensure you select switches that are part of the same fabric.

5. Enter the priority value for the policy in the **Priority** field.

   The applicable values are from 1 to 1000. The default value is 500. A lower number in the **Priority** field indicates that there is a higher priority for the generated configuration and POAP startup-configuration. For example, the priority for vPC related policies are as follows: base_feature_vpc is 100, vpc-domain_mgmt is 150, for policies for interfaces on vPC (int_vpc_peer_link_po) is

202.

6. Use the toggle switch to enable or disable the **Group** option, as required.

   If you have selected multiple switches, the **Group** toggle switch is enabled by default. If you select one switch initially and choose to add additional switches later, you can select the **Group** toggle switch to create a policy group and add additional switches later.

   Not all templates provide support for creating policy groups. If you choose a template that does not support a policy group, Nexus Dashboard generates an error message. Ensure you uncheck the **Group** toggle switch and create regular policies for templates that do not support a policy group.

7. Click **Choose Template** and choose the appropriate policy template and click **Select**.

   The available policy templates are **TEMPLATE_CLI**, **PYTHON**, and **PYTHON_CLI**.

   Note that nested Python policies are not supported. Additionally, when configuring policy groups, make sure that you do not add policies that can be applied only on a single switch. Choose policies that you can apply on multiple switches.

8. Depending on the policy template that you have chosen, enter all the necessary field values to create a policy and click **Save**.

   The new policy group appears on the **Fabric Overview > Policies** page.

9. To deploy the configuration to the switches, choose the new policy that you have created and choose **Actions > Push Config**.

   The **Generated Config** page displays with the pending configuration changes.

10. Click **Push Config** to push the pending configuration to the device.
11. Alternatively, to deploy the configuration, navigate to the **Inventory > Switches** page and choose **Actions > Recalculate and deploy**.

    Note that the **Push Config** option does not go through configuration compliance checks. Use the **Push Config** option only when you want to deploy commands that are ignored during configuration compliance checks.

Advertise a PIP on a vPC

Follow these steps to enable the advertise PIP feature on a vPC.

1. Choose the required LAN fabric and navigate to **Edit Fabric Settings > Fabric Management > vPC** and check the **vPC advertise-pip** check box to enable advertising the primary IP address (PIP) feature on all vPCs in a fabric.
2. Choose the **vpc_advertise_pip_jython** policy to enable the advertise PIP feature on specific vPCs in a fabric.

Guidelines and limitations for advertising a PIP on a vPC

- If you do not globally enable **vPC advertise-pip** or a vPC peer is not using fabric peering, only then can you create the **vpc_advertise_pip_jython policy** on specific peers.
- You can apply the policy **vpc_advertise_pip_jython** only when switches are part of vPC pairing.
- Ensure that you configure the vpc advertise-pip command during a maintenance period, as it involves a BGP next-hop rewrite. Enabling this feature with EVPN type 5 uses the switch primary IP address as the next-hop while EVPN type 2 continues to use a secondary IP address.
- Disabling **vPC advertise-pip** for a fabric doesn't affect this policy.
- Unpairing of switches deletes this policy.
- You can manually delete this policy from the peer switch where it was created.

Follow these steps to advertise a PIP on a vPC.

1. Navigate to the Configuration policies page.
2. Click the **Policies** tab.
3. On the **Policies** page, choose **Actions > Add policy** and then choose a switch with a vPC.
4. Click **Actions > Add** and choose the switch from the **Switch List** drop-down list.
5. Choose the **vpc_advertise_pip_jython** policy template and enter the mandatory parameters.

You can add this policy on one vPC peer, and the policy creates the respective commands for vPC advertisement on both peers.

6. Click **Save** and then deploy this policy.

Navigate to the Inventory page

Follow these steps to navigate to the **Inventory** page to view or edit device information. You can navigate to the **Inventory** page using either of these methods.

To view inventory information at the Nexus Dashboard level, click **Manage > Inventory**.

Follow these steps to view inventory information at an individual fabric level.

1. Click **Manage > Fabrics**.
2. Click the appropriate fabric on the **Fabrics** page.
3. Click the **Inventory** tab.

Custom maintenance mode profile policy

Nexus Dashboard configures only a fixed set of BGP and OSPF isolate CLIs in the maintenance mode profile when you place a switch in maintenance mode. You can create a custom_maintenance_mode_profile policy with customized configurations for maintenance mode and normal mode profiles, deploy the policy to the switch, and then move the switch to maintenance mode.

Create and deploy a custom maintenance mode profile policy

Follow these steps to create and deploy a custom maintenance mode profile policy from Nexus Dashboard.

1. Navigate to the Inventory page.

   If you navigate to the **Inventory** page from the fabric, click **Inventory > Switches**.

   The **Inventory** page shows information on already-configured switches.

2. Click the appropriate switch.
3. Under the **Configuration policies** tab, choose **Policies**.
4. From the **Actions** drop-down list, choose **Add policy** to add a new policy.
5. In the **Create Policy** page, click **Select Template**.
6. Choose **custom_maintenance_mode_profile** from the **Select Policy Template** list and click **Select**.
7. Fill in the **Maintenance mode profile contents** with the desired configuration CLIs.

   Example:

   ```
   configure maintenance profile maintenance-mode
   ip pim isolate
   ```

   Fill in the **Normal mode profile contents** with the desired configuration CLIs.

   Example:

   ```
   configure maintenance profile normal-mode
   no ip pim isolate
   configure terminal
   ```

8. Click **Save**.
9. From the **Switch Overview** page, click **Actions > Preview**.
10. Click on **Pending Config** lines to view the **Pending Config** and **Side-by-Side Comparison**.
11. Click **Close**.
12. From the **Switch Overview** page, click **Actions > Deploy** and then click **Deploy All** to deploy the new policy configuration on the switch.

    Click **Close** after the deployment is complete.

13. Choose the policy and navigate to **Actions > More > Change Mode**.
14. In the **Mode** drop-down list, choose **Maintenance**.
15. Click **Save and Deploy Now** to move the switch to maintenance mode.

When you apply the default maintenance profile to a device functioning as an anycast border gateway (BGW) within a VXLAN fabric, it can result in dropping network traffic specifically for multi-fabric BUM traffic. This issue affects a subset of VNIs, particularly those for which the given BGW is designated as the forwarder. To address this issue, the maintenance mode profile must always use include-local for BGP isolation. In such cases, custom_maintenance_mode_profile policy must be created and deployed following the above steps, and then the necessary configuration CLIs should be modified.

The following is a sample custom_maintenance_mode_profile policy content for an anycast BGW in a VXLAN fabric with an OSPF underlay and multicast replication mode.

Normal mode profile contents example:

```
configure maintenance profile normal-mode
  router ospf UNDERLAY
    no isolate
  router bgp 65001
    no isolate include-local
  no ip pim isolate
```

Maintenance mode profile contents example:

```
configure maintenance profile maintenance-mode
  ip pim isolate
  router bgp 65001
    isolate include-local
  router ospf UNDERLAY
    isolate
```

Delete a custom maintenance mode profile policy

The switch has to be moved to active, operational, or normal mode before deleting the custom maintenance mode profile policy.

Follow these steps to delete a custom maintenance mode profile policy from the **Switch Overview** page.

1. Choose the desired switch to navigate to the **Switch Overview** page.
2. From the **Switch Overview** page, choose **Actions > More > Change Mode**.
3. In the **Mode** drop-down list, choose **Normal**.
4. Click **Save and Deploy Now** to move the switch to normal mode.
5. After the switch has been moved to normal mode, choose the **custom_maintenance_mode_profile** policy that has to be deleted.
6. Choose **Actions > Edit policy**.
7. Choose **Actions > Delete policy** and click **Confirm** to mark the policy for deletion.

The **Mark Deleted** column shows **true** indicating that the policy is marked for deletion.

8. Again, choose **Actions > Delete policy** and click **Confirm** to delete the policy.
9. From the **Switch Overview** page, choose **Actions > Deploy**.
10. Click **Deploy All** to delete the policy configuration on the switch.
11. Click **Close** after the deployment is complete.

Allocations

The **Allocations** page allows you to manage your resources.

This table describes the fields that appear on the **Allocations** 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.
VRF name	Specifies the VRF name associated with the resource allocation.
ID	Specifies the ID.

Release a resource

Follow these steps to release a resource from Nexus Dashboard.

1. Choose **Manage > Fabrics**.
2. Click the LAN or IPFM fabric where you want to gather information on resource allocations.

   The **Fabric Overview** page displays.
3. Click the **Configuration policies** tab.
4. Click the **Resources** tab.
5. Choose a resource that you want to delete.

   

   You can delete multiple resources at the same time by choosing multiple resources.
6. Click **Actions > Release resource(s)** to release the resource.

A confirmation dialog box displays.

7. Click **Confirm** to release the resource.

Copyright

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

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