

Working with Configuration Policies for Your Nexus Dashboard LAN or IPFM Fabrics, Release 4.3.1

Table of Contents

New and changed information . . . 1
Navigate to the Configuration policies page . . . 2
Policies . . . 3
- Simplified policy template selection and combined configuration view . . . 3
- Dynamic load balancing policy template . . . 3
  - Supported switches for dynamic load balancing . . . 3
  - Guidelines and limitations for dynamic load balancing . . . 4
- Access the Policies page . . . 4
- Add a policy . . . 9
- Add a Dynamic Load Balancing (DLB) policy template . . . 11
- Create a policy group . . . 13
- Advertise a PIP on a vPC . . . 15
- Guidelines and limitations for advertising a PIP on a vPC . . . 15
- Navigate to the Inventory page . . . 16
- Custom maintenance mode profile policy . . . 16
  - Create and deploy a custom maintenance mode profile policy . . . 16
  - Delete a custom maintenance mode profile policy . . . 18

Resources . . . 19
- Release a resource . . . 19

Copyright . . . 21

New and changed information

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

Release Version	Feature	Description
Nexus Dashboard 4.3.1	Policy simplification	Beginning in Nexus Dashboard 4.3.1, policy template selection now supports category-based filtering, and the generated configuration view consolidates related parent and child policy configuration into a single show run-like display. Generated configuration for multiple non-group policies is shown by policy ID, and policy groups use Group instances to identify switch-level policy instances. For more information, see Simplified policy template selection and combined configuration view.

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 device configuration using policies. Nexus Dashboard policies group the required CLIs and variables to achieve specific configuration on devices. These policies can be defined either by 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.

Simplified policy template selection and combined configuration view

The policy template selection experience now includes category-based filtering. In the **Select Policy Template** panel, policy templates are organized by category so that you can narrow the template list to the type of configuration you want to create. You can select one or more categories to show templates from any selected category.

Use the category control with the **Filter** field to locate a policy template by category, name, or description. The generated configuration display is also updated.

Previously, generated configuration could show separate, repeated blocks for related child policies. For example, if several child policies were associated with the same interface, the interface configuration could appear multiple times.

The updated display consolidates related parent and child policy configuration into a single show run-like view. When generated configuration includes multiple non-group policies, Nexus Dashboard shows the generated configuration by policy ID so that you can identify which policy contributed each configuration block. For policy groups, Nexus Dashboard shows group-level generated configuration and uses **Group instances** to identify the switch-level policy instances associated with the group.

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. For a policy group, this field shows the number of switches associated with the policy group. Open the policy group entry to view switch-level details in the policy configuration summary and **Group instances** area. The switch count is not a separate link to policy group details.
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 policy priority. For policy groups with mixed switch-instance states, the **Priority** column can display **Mixed**. Open the policy group entry and review **Group instances** for per-switch priority and **Mark Deleted** state. For switch_freeform policies with the **PYTHON** content type, Nexus Dashboard can display the source and child policies as separate entries after an edit operation. In the child policy entry, **Mark Deleted** can be **true** and **Priority** can be a negative value.
Editable	Specifies a Boolean value to indicate if the policy is editable.
Mark deleted	Specifies whether the policy is marked for deleted. The column displays **true** indicating that the policy is marked for deletion. Configurations for a policy with the **Mark Deleted** value **true** are negated, and the **Generated Config** for the policy displays the configuration to be removed from the switch. For policy groups, open the policy group entry and review **Group instances** for the switch-level **Mark Deleted** state. The **Mark Deleted** value in the **Policies** table is not a separate link to policy group details.
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. Policy IDs also appear in generated configuration views so you can identify the policy or policy group that contributed a configuration block.
IP address	Specifies the IP address of the switch. For policy groups, open the policy group entry and review **Group instances** for the IP addresses of associated switch instances. IP address values in the **Policies** table are not separate links to policy group details.

Serial number	Specifies the serial number of the switch. For policy groups, open the policy group entry and review **Group instances** for the serial numbers of associated switch instances. Serial number values in the **Policies** table are not separate links to policy group details.
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 membership for a policy group. You can add or remove switches from a policy group using this option. When you remove a switch from a policy group, Nexus Dashboard keeps the removed switch visible as a group instance until the pending membership change is deployed. The removed instance is marked for deletion, and the generated configuration shows the configuration that will be removed from the switch. When membership changes leave switch instances in different states, the **Priority** column can display **Mixed**. Open the policy group entry and review **Group instances** to see each switch instance, including whether **Mark Deleted** is **true** and the priority value for that instance. After an **Edit membership** operation, deploy the pending membership configuration changes before making additional policy configuration changes. 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. Deleting a user-related policy (such as **switch_user_strong_encrypted** or **secure_switch_snmp_user_sha**) does not remove the corresponding user account from the switch either on a **Recalculate and Deploy** or **Config push** of the policy marked for deletion. This behavior applies to all policies associated with user creation.
Generated Config	To view generated configuration for one or more policies, choose the policies from the table and choose **Generated Config**. The **Generated Config** page displays a show run-like view that consolidates related parent and child policy configuration and reduces repeated configurations. When multiple non-group policies are selected, Nexus Dashboard shows generated configuration as separate sections by policy ID. For policy groups, the page shows group-level generated configuration, and the **Group instances** area identifies the switch-level policy instances associated with the group.

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. 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.
6. Enter a description for the policy in the **Description** field.
7. In the **Policy template** field, click **No Policy Selected**.

   The **Select Policy Template** panel appears.

   a. Use the category control to narrow the list of policy templates to one or more configuration areas, such as **BGP**, **Manageability**, **QoS**, **Dynamic load balancing**, **AI**, **VRF lite**, **Routing policies**, or **Others**. If you select more than one category, the panel shows templates from any selected category.

   b. (Optional) In the **Filter** field, enter all or part of a policy template name, category, or description.

   c. Select the policy template that you want to use.

   d. Click **Select**.

      The selected policy template appears in the **Policy template** field. Depending on the selected policy template, additional fields and tabs appear on the **Create Policy** page.

      The following steps describe examples of fields that appear for selected policy templates. If your selected template does not show these fields, fill in the fields that apply to that template

and click **Save**.

If the selected policy template includes PTP high-correction notification fields, you can enable or disable notification when the system encounters a high-correction event. A correction is considered high when the correction value exceeds the configured value. By default, high-correction notification is disabled.

Perform the following steps to enable the high-correction notification:

e. Put a check in the **Enable PTP Telemetry** check box to enable telemetry for PTP.

f. 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.

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

The duration value is in seconds.

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

The default is 100000 (100us).

8. 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 **Actions > Add**.

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.

9. If you chose the **Dynamic_Load_Balancing** policy template, perform these steps for the template-specific DLB fields. For the full DLB procedure, see Add a Dynamic Load Balancing (DLB) policy template.

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.

10. For all other policy templates, fill out the fields that apply to the selected template, and then click **Save**.

    After you save the policy, the policy appears on the **Policies** page. To review and deploy pending configuration, select the policy and choose **Actions > Push Config**.

    The **Generated Config** page displays the pending configuration in a consolidated, show run-like view. When you review generated configuration for multiple non-group policies, Nexus Dashboard shows a separate generated-configuration section for each policy ID. For policy groups, the generated configuration is associated with the policy group. See Create a policy group.

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

      

      The **DLB Interfaces** field is disabled when you apply the DLB policy template configuration at the fabric level by using **Apply Fabric Level Setting**.

   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. For policy driven mixed mode, choose the default load-balancing mode in the **DLB Mixed Mode Default** field. This field is inactive for other modes. The available options are **ecmp**, **flowlet**, and **per packet**.
   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.

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

After you save the DLB policy, select the policy on the **Policies** page and choose **Actions > Push Config** to review and deploy the pending DLB configuration.

The **Generated Config** page uses the same consolidated generated-configuration view and policy ID grouping behavior as other policies.

What's next: To enable DLB configuration for AI fabrics, see the "AI settings" section in Editing AI Data Center VXLAN Fabric Settings. To enable DLB configuration for VXLAN fabrics with iBGP overlay routing protocol, see the **Advanced** settings section in Fabric Management.

Create a policy group

Policy groups provide a method for configuring and managing switches collectively.

Use a policy group when the same policy template and shared field values must be applied to multiple switches in the same fabric. You can add switches when you create the policy group or later by editing membership.

On the **Policies** page, a policy group appears as a single policy group entry. Open the policy group entry to view the policy configuration summary and the **Group instances** area for switch-level policy instances. For policy groups, use the policy group entry to open details. Switch count, priority, IP address, serial number, and **Mark Deleted** values are not separate links to policy group details.

Follow these steps to create a policy group.

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. Select the switches to which you want to apply the policy group, and click **Next**.

   The switches must be part of the same fabric.

5. Enter the priority value for the policy group 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 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 the policy group in the **Description** field.
7. Ensure that the **Group** toggle is enabled.

   If you selected multiple switches, the **Group** toggle is enabled by default. If you selected one switch initially and add more switches later, enable the **Group** toggle before saving to create a policy group.

   Not all policy templates support policy groups. If the selected policy template does not support policy groups, Nexus Dashboard displays an error. Disable the **Group** toggle and create regular policies for templates that do not support policy groups.

8. In the **Policy template** field, click **No Policy Selected**.

   The **Select Policy Template** panel appears.

   Use the category control to narrow the list of policy templates to one or more configuration areas, such as **BGP**, **Manageability**, **QoS**, **Dynamic load balancing**, **AI**, **VRF lite**, **Routing policies**, or **Others**. If you select more than one category, the panel shows templates from any selected category.

   (Optional) In the **Filter** field, enter all or part of a policy template name, category, or description.

   Select a policy template that supports policy groups.

   Click **Select**.

   The selected policy template appears in the **Policy template** field. Policy group templates can use **TEMPLATE_CLI**, **PYTHON**, or **PYTHON_CLI** content types. Nested Python policies are not supported.

   Choose templates that can be applied to multiple switches; do not use templates that apply only to a single switch.

9. Depending on the selected policy template, enter the required field values and click **Save**.

   The new policy group appears on the **Policies** page as one policy group entry. The policy ID begins with **POLICY-GROUP**.

   Use the policy group entry to review the policy configuration summary and **Group instances** for switch-level details such as switch name, IP address, serial number, priority, **Mark Deleted** state,

and policy ID.

10. To deploy the configuration to the switches, choose the policy group that you created and choose **Actions > Push Config**.

    The **Generated Config** page displays the pending configuration changes. For policy groups, the page displays consolidated, show run-like generated configuration for the group, and **Group instances** identifies the switch-level policy instances associated with the group.

11. Click **Push Config** to push the pending configuration to the switches.

12. 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 **Policy template** field, click **No Policy Selected**.
6. Choose **custom_maintenance_mode_profile** from the **Select Policy Template** panel 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.

Resources

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

This table describes the fields that appear on the **Resources** 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

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

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

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

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

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

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

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

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

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

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

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