



## Configuring Fabric and Interfaces

---

- [Fabric and Interface Configuration, on page 1](#)
- [Graceful Insertion and Removal \(GIR\) Mode, on page 2](#)
- [Configuring Physical Ports in Leaf Nodes and FEX Devices Using the NX-OS CLI, on page 3](#)
- [Configuring Port Channels in Leaf Nodes and FEX Devices Using the NX-OS CLI, on page 6](#)
- [Configuring Virtual Port Channels in Leaf Nodes and FEX Devices Using the NX-OS CLI, on page 12](#)
- [Configuring FEX Connections Using Profiles with the NX-OS Style CLI, on page 17](#)
- [Reflective Relay \(802.1Qbg\), on page 18](#)
- [Configuring Policy Groups for Interfaces, on page 20](#)
- [Configuring Overrides for Interfaces, on page 23](#)
- [About Forwarding Error Correction, on page 25](#)

### Fabric and Interface Configuration

To form the ACI fabric, Cisco Nexus 9000 Series ACI-mode switches are deployed in a leaf and spine topology managed by the APIC controller. Each leaf node is connected to all spine nodes with no connectivity between the leaf nodes. The interconnecting links between the leaf and spine nodes are called fabric links and the respective ports are called fabric ports. The fabric ports do not require user configuration for normal operation as these are auto discovered and factory default configuration is applied during fabric bring-up. All endpoint devices are connected to the leaf nodes through access ports. The access ports must be configured similar to those in NX-OS switches. Both fabric and access ports are represented as Interfaces as in NX-OS.

The leaf and spine nodes are considered different objects in the ACI model and support different sets of policies. In the CLI, these nodes are represented as leaf and spine respectively while both are commonly referred to as nodes. Leaf and spine node values are unique across all the pods in the fabric. FEX modules, if attached to the leaf nodes, will have fex-id values unique only within each leaf. For example, two leaf nodes can each have a FEX 101 attached.



---

**Note** Configuring FEX connections with FEX IDs 165 to 199 is not supported in the APIC GUI. To use one of these FEX IDs, configure the profile using the NX-OS style CLI. For more information, see *Configuring FEX Connections Using Interface Profiles with the NX-OS Style CLI*.

---

As of Cisco APIC, Release 3.0(1k), connections to FEX modules can be configured as profiles.

### Interface Naming for Leaf and FEX Interfaces

In ACI fabric, most interface configuration is done for physical ports, port-channels, or vPCs (either directly connected to leaf nodes or connected through FEX modules). The general command syntax for each interface type is shown in the following table.

Interface Type	Command Syntax	Examples
Port	<b>interface ethernet</b> <i>slot/port</i>	interface eth 1/1
FEX Port	<b>interface ethernet</b> <i>fex-id/slot/port</i>	interface eth 101/1/1
Port-channel	<b>interface port-channel</b> <i>name</i>	interface port-channel foo
FEX Port-channel	<b>interface port-channel</b> <i>name</i> <b>fex</b> <i>fex-id</i>	interface port-channel foo flex 101
Virtual Port-channel (VPC)	<b>interface vpc</b> <i>name</i>	interface vpc foo
vPC over FEX	<b>interface vpc</b> <i>name</i> <b>fex</b> <i>fex-a</i> <i>fex-b</i>	interface vpc foo flex 101 102

## Graceful Insertion and Removal (GIR) Mode

The Graceful Insertion and Removal (GIR) mode, or maintenance mode, allows you to isolate a switch from the network with minimum service disruption. In the GIR mode you can perform real-time debugging without affecting traffic.

You can use graceful insertion and removal to gracefully remove a switch and isolate it from the network in order to perform debugging operations. The switch is removed from the regular forwarding path with minimal traffic disruption. When you are finished performing the debugging operations, you can use graceful insertion to return the switch to its fully operational (normal) mode. In graceful removal, all external protocols are gracefully brought down except the fabric protocol (IS-IS) and the switch is isolated from the network. During maintenance mode, the maximum metric is advertised in IS-IS within the Cisco Application Centric Infrastructure (Cisco ACI) fabric and therefore the maintenance mode TOR does not attract traffic from the spine switches. In addition, all the front-panel interfaces are shutdown on the switch except the fabric interfaces. In graceful insertion, the switch is automatically decommissioned, rebooted, and recommissioned. When recommissioning is completed, all external protocols are restored and maximum metric in IS-IS is reset after 10 minutes.

The following protocols are supported:

- Border Gateway Protocol (BGP)
- Enhanced Interior Gateway Routing Protocol (EIGRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Open Shortest Path First (OSPF)
- Link Aggregation Control Protocol (LACP)

### Important Notes

- Upgrading or downgrading a switch in maintenance mode is not supported.

- While the switch is in maintenance mode, the Ethernet port module stops propagating the interface related notifications. As a result, if the remote switch is rebooted or the fabric link is flapped during this time, the fabric link will not come up afterward unless the switch is manually rebooted (using the **acidiag touch clean** command), decommissioned, and recommissioned.
- For multi-pod, **IS-IS metric for redistributed routes** should be set to less than 63. To set the **IS-IS metric for redistributed routes**, choose **Fabric > Fabric Policies > Pod Policies > IS-IS Policy**.
- Existing GIR supports all Layer 3 traffic diversion. With LACP, all the Layer 2 traffic is also diverted to the redundant node. Once a node goes into maintenance mode, LACP running on the node immediately informs neighbors that it can no longer be aggregated as part of port-channel. All traffic is then diverted to the vPC peer node.
- For a GIR upgrade, Cisco Application Policy Infrastructure Controller (Cisco APIC)-connected leaf switches must be put into different maintenance groups such that the Cisco APIC-connected leaf switches get upgraded one at a time.

## Removing a Switch to Maintenance Mode Using the CLI

Use this procedure to remove a switch to maintenance mode using the CLI.

### Procedure

	Command or Action	Purpose
Step 1	<code>[no]debug-switch <i>node_id</i> or <i>node_name</i></code>	Removes the switch to maintenance mode.

## Inserting a Switch to Operation Mode Using CLI

Use this procedure to insert a switch to operational mode using the CLI.

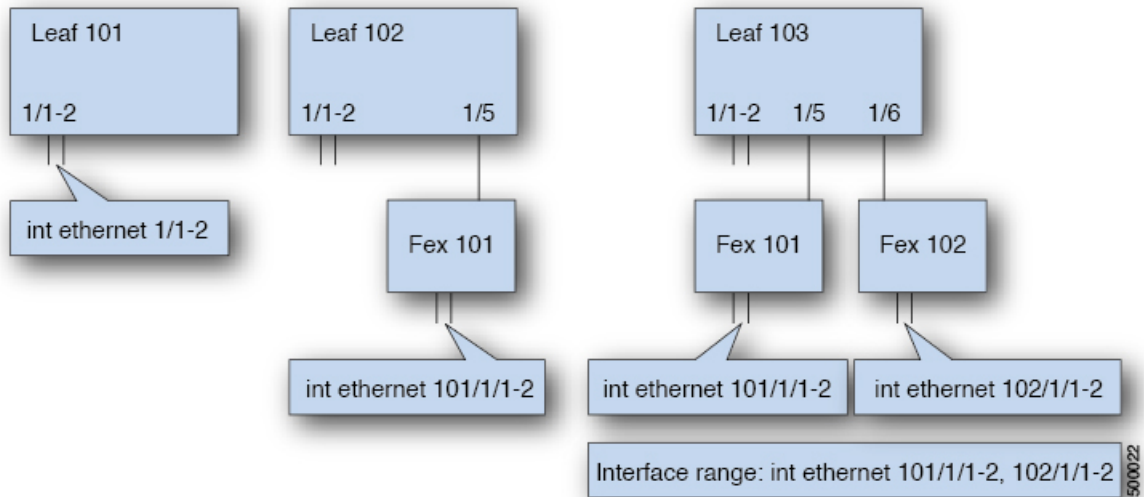
### Procedure

	Command or Action	Purpose
Step 1	<code>[no]no debug-switch <i>node_id</i> or <i>node_name</i></code>	Inserts the switch to operational mode.

## Configuring Physical Ports in Leaf Nodes and FEX Devices Using the NX-OS CLI

The commands in the following examples create many managed objects (MOs) in the ACI policy model that are fully compatible with the REST API/SDK and GUI. However, the CLI user can focus on the intended network configuration instead of ACI model internals.

The following figure shows examples of Ethernet ports directly on leaf nodes or FEX modules attached to leaf nodes and how each is represented in the CLI. For FEX ports, the *fex-id* is included in the naming of the port itself as in **ethernet 101/1/1**. While describing an interface range, the **ethernet** keyword need not be repeated as in NX-OS. Example: **interface ethernet 101/1/1-2, 102/1/1-2**.



- Leaf node ID numbers are global.
- The *fex-id* numbers are local to each leaf.
- Note the space after the keyword **ethernet**.

**Procedure**

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b> <b>Example:</b> apic1# <b>configure</b>	Enters global configuration mode.
<b>Step 2</b>	<b>leaf node-id</b> <b>Example:</b> apic1(config)# <b>leaf 102</b>	Specifies the leaf or leaves to be configured. The <i>node-id</i> can be a single node ID or a range of IDs, in the form <i>node-id1 - node-id2</i> , to which the configuration will be applied.
<b>Step 3</b>	<b>interface type</b> <b>Example:</b> apic1(config-leaf)# <b>interface ethernet 1/2</b>	Specifies the interface that you are configuring. You can specify the interface type and identity. For an Ethernet port, use “ethernet slot / port.”
<b>Step 4</b>	(Optional) <b>fex associate node-id</b> <b>Example:</b> apic1(config-leaf-if)# <b>fex associate 101</b>	If the interface or interfaces to be configured are FEX interfaces, you must use this command to attach the FEX module to a leaf node before configuration.  <b>Note</b> This step is required before creating a port-channel using FEX ports.

	Command or Action	Purpose
<b>Step 5</b>	<b>speed</b> <i>speed</i> <b>Example:</b> apic1(config-leaf-if)# <b>speed 10G</b>	The speed setting is shown as an example. At this point you can configure any of the interface settings shown in the table below.

The following table shows the interface settings that can be configured at this point.

Command	Purpose
[no] shut	Shut down physical interface
[no] speed <i>speedValue</i>	Set the speed for physical interface
[no] link debounce time <i>time</i>	Set link debounce
[no] negotiate auto	Configure negotiate
[no] cdp enable	Disable/enable Cisco Discovery Protocol (CDP)
[no] mcp enable	Disable/enable Mis-cabling Protocol (MCP)
[no] lldp transmit	Set the transmit for physical interface
[no] lldp receive	Set the LLDP receive for physical interface
spanning-tree {bpduguard   bpdufilter} {enable   disable}	Configure spanning tree BPDU
[no] storm-control level <i>percentage</i> [ burst-rate <i>percentage</i> ]	Storm-control configuration (percentage)
[no] storm-control pps <i>packets-per-second</i> burst-rate <i>packets-per-second</i>	Storm-control configuration (packets-per-second)

### Examples

Configure one port in a leaf node. The following example shows how to configure the interface eth1/2 in leaf 101 for the following properties: speed, cdp, and admin state.

```
apic1# configure
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/2
apic1(config-leaf-if)# speed 10G
apic1(config-leaf-if)# cdp enable
apic1(config-leaf-if)# no shut
```

Configure multiple ports in multiple leaf nodes. The following example shows the configuration of speed for interfaces eth1/1-10 for each of the leaf nodes 101-103.

```
apic1(config)# leaf 101-103
apic1(config-leaf)# interface eth 1/1-10
```

```
apicl(config-leaf-if)# speed 10G
```

Attach a FEX to a leaf node. The following example shows how to attach a FEX module to a leaf node. Unlike in NX-OS, the leaf port Eth1/5 is implicitly configured as fabric port and a FEX fabric port-channel is created internally with the FEX uplink port(s). In ACI, the FEX fabric port-channels use default configuration and no user configuration is allowed.




---

**Note** This step is required before creating a port-channel using FEX ports, as described in the next example.

---

```
apicl(config)# leaf 102
apicl(config-leaf)# interface eth 1/5
apicl(config-leaf-if)# fex associate 101
```

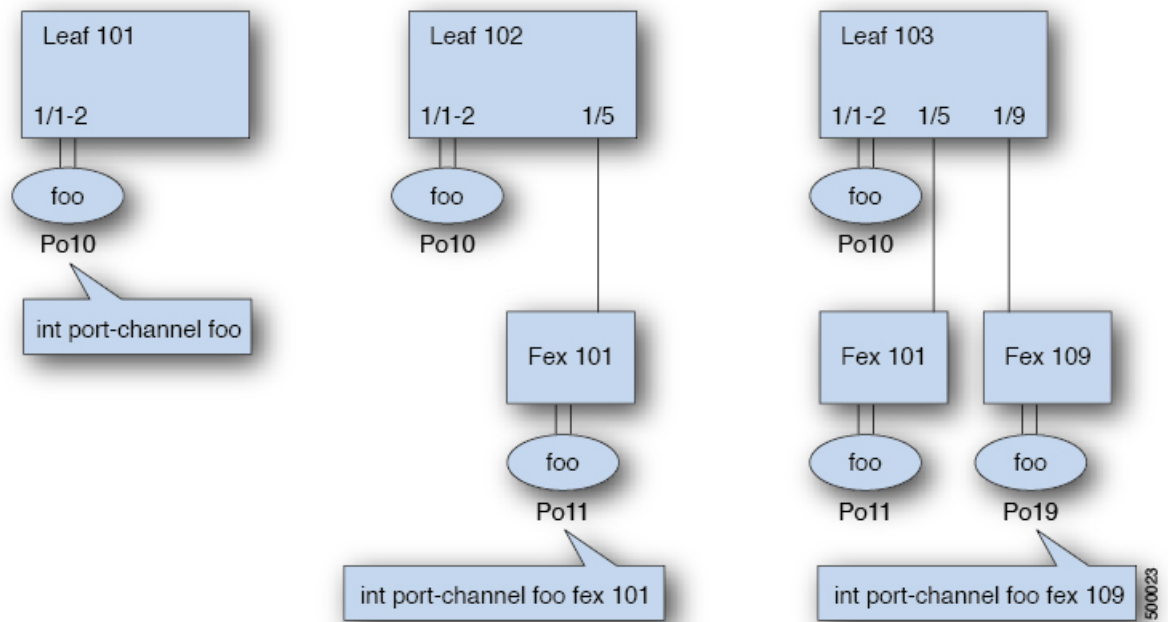
Configure FEX ports attached to leaf nodes. This example shows configuration of speed for interfaces eth1/1-10 in FEX module 101 attached to each of the leaf nodes 102-103. The FEX ID 101 is included in the port identifier. FEX IDs start with 101 and are local to a leaf.

```
apicl(config)# leaf 102-103
apicl(config-leaf)# interface eth 101/1/1-10
apicl(config-leaf-if)# speed 1G
```

## Configuring Port Channels in Leaf Nodes and FEX Devices Using the NX-OS CLI

Port-channels are logical interfaces in NX-OS used to aggregate bandwidth for multiple physical ports and also for providing redundancy in case of link failures. In NX-OS, port-channel interfaces are identified by user-specified numbers in the range 1 to 4096 unique within a node. Port-channel interfaces are either configured explicitly (using the **interface port-channel** command) or created implicitly (using the **channel-group** command). The configuration of the port-channel interface is applied to all the member ports of the port-channel. There are certain compatibility parameters (speed, for example) that cannot be configured on the member ports.

In the ACI model, port-channels are configured as logical entities identified by a name to represent a collection of policies that can be assigned to set of ports in one or more leaf nodes. Such assignment creates one port-channel interface in each of the leaf nodes identified by an auto-generated number in the range 1 to 4096 within the leaf node, which may be same or different among the nodes for the same port-channel name. The membership of these port-channels may be same or different as well. When a port-channel is created on the FEX ports, the same port-channel name can be used to create one port-channel interface in each of the FEX devices attached to the leaf node. Thus, it is possible to create up to N+1 unique port-channel interfaces (identified by the auto-generated port-channel numbers) for each leaf node attached to N FEX modules. This is illustrated with the examples below. Port-channels on the FEX ports are identified by specifying the *fex-id* along with the port-channel name ( **interface port-channel foo fex 101** , for example).



- N+1 instances per leaf of port-channel foo are possible when each leaf is connected to N FEX nodes.
- Leaf ports and FEX ports cannot be part of the same port-channel instance.
- Each FEX node can have only one instance of port-channel foo.

**Procedure**

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b> <b>Example:</b> apicl# <b>configure</b>	Enters global configuration mode.
<b>Step 2</b>	<b>template port-channel channel-name</b> <b>Example:</b> apicl(config)# <b>template port-channel foo</b>	Creates a new port-channel or configures an existing port-channel (global configuration).
<b>Step 3</b>	<b>[no] switchport access vlan vlan-id tenant tenant-name application application-name epg epg-name</b> <b>Example:</b> apicl(config-po-ch-if)# <b>switchport access vlan 4 tenant ExampleCorp application Web epg webEpg</b>	Deploys the EPG with the VLAN on all ports with which the port-channel is associated.

	Command or Action	Purpose
<b>Step 4</b>	<p><b>channel-mode active</b></p> <p><b>Example:</b></p> <pre>apic1(config-po-ch-if) # channel-mode active</pre> <p><b>Note</b> To enable symmetric hashing, enter the <b>lACP symmetric-hash</b> command:</p> <pre>apic1(config-po-ch-if) # lACP symmetric-hash</pre>	<p><b>Note</b> The <b>channel-mode</b> command is equivalent to the mode option in the channel-group command in NX-OS. In ACI, however, this is supported for the port-channel (not on a member port).</p> <p>Symmetric hashing is not supported on the following switches:</p> <ul style="list-style-type: none"> <li>• Cisco Nexus 93128TX</li> <li>• Cisco Nexus 9372PX</li> <li>• Cisco Nexus 9372PX-E</li> <li>• Cisco Nexus 9372TX</li> <li>• Cisco Nexus 9372TX-E</li> <li>• Cisco Nexus 9396PX</li> <li>• Cisco Nexus 9396TX</li> </ul>
<b>Step 5</b>	<p><b>exit</b></p> <p><b>Example:</b></p> <pre>apic1(config-po-ch-if) # exit</pre>	Returns to configure mode.
<b>Step 6</b>	<p><b>leaf node-id</b></p> <p><b>Example:</b></p> <pre>apic1(config) # leaf 101</pre>	Specifies the leaf switches to be configured. The <i>node-id</i> can be a single node ID or a range of IDs, in the form <i>node-id1 - node-id2</i> , to which the configuration will be applied.
<b>Step 7</b>	<p><b>interface type</b></p> <p><b>Example:</b></p> <pre>apic1(config-leaf) # interface ethernet 1/1-2</pre>	Specifies the interface or range of interfaces that you are configuring to the port-channel.
<b>Step 8</b>	<p><b>[no] channel-group channel-name</b></p> <p><b>Example:</b></p> <pre>apic1(config-leaf-if) # channel-group foo</pre>	Assigns the interface or range of interfaces to the port-channel. Use the keyword <b>no</b> to remove the interface from the port-channel. To change the port-channel assignment on an interface, you can enter the <b>channel-group</b> command without first removing the interface from the previous port-channel.
<b>Step 9</b>	<p>(Optional) <b>lACP port-priority priority</b></p> <p><b>Example:</b></p> <pre>apic1(config-leaf-if) # lACP port-priority</pre>	This setting and other per-port LACP properties can be applied to member ports of a port-channel at this point.



	Command or Action	Purpose
	<pre>1000 apicl(config-leaf-if)# lacp rate fast</pre>	<p><b>Note</b> In the ACI model, these commands are allowed only after the ports are member of a port channel. If a port is removed from a port channel, configuration of these per-port properties are removed as well.</p>

The following table shows various commands for global configurations of port channel properties in the ACI model. These commands can also be used for configuring overrides for port channels in a specific leaf in the (config-leaf-if) CLI mode. The configuration made on the port-channel is applied to all member ports.

CLI Syntax	Feature
[no] speed <speedValue>	Set the speed for port-channel
[no] link debounce time <time>	Set Link Debounce for port-channel
[no] negotiate auto	Configure Negotiate for port-channel
[no] cdp enable	Disable/Enable CDP for port-channel
[no] mcp enable	Disable/Enable MCP for port-channel
[no] lldp transmit	Set the transmit for port-channel
[no] lldp receive	Set the lldp receive for port-channel
spanning-tree <bpduguard   bpdufilter> <enable   disable>	Configure spanning tree BPDU
[no] storm-control level <percentage> [ burst-rate <percentage> ]	Storm-control configuration (percentage)
[no] storm-control pps <packet-per-second> burst-rate <packets-per-second>	Storm-control configuration (packets-per-second)
[no] channel-mode { active   passive   on   mac-pinning }	LACP mode for the link in port-channel 1
[no] lacp min-links <value>	Set minimum number of links
[no] lacp max-links <value>	Set maximum number of links
[no] lacp fast-select-hot-standby	LACP fast select for hot standby ports
[no] lacp graceful-convergence	LACP graceful convergence
[no] lacp load-defer	LACP load defer member ports
[no] lacp suspend-individual	LACP individual Port suspension
[no] lacp port-priority	LACP port priority

CLI Syntax	Feature
[no] lacp rate	LACP rate

### Examples

Configure a port channel (global configuration). A logical entity foo is created that represents a collection of policies with two configurations: speed and channel mode. More properties can be configured as required.



**Note** The channel mode command is equivalent to the mode option in the channel group command in NX-OS. In ACI, however, this is supported for the port-channel (not on member port).

```
apic1(config)# template port-channel foo
apic1(config-po-ch-if)# switchport access vlan 4 tenant ExampleCorp application Web epg webEpg
apic1(config-po-ch-if)# speed 10G
apic1(config-po-ch-if)# channel-mode active
```

Configure ports to a port-channel in a FEX. In this example, port channel foo is assigned to ports Ethernet 1/1-2 in FEX 101 attached to leaf node 102 to create an instance of port channel foo. The leaf node will auto-generate a number, say 1002 to identify the port channel in the switch. This port channel number would be unique to the leaf node 102 regardless of how many instance of port channel foo are created.



**Note** The configuration to attach the FEX module to the leaf node must be done before creating port channels using FEX ports.

```
apic1(config)# leaf 102
apic1(config-leaf)# interface ethernet 101/1/1-2
apic1(config-leaf-if)# channel-group foo
```

In Leaf 102, this port channel interface can be referred to as interface port-channel foo FEX 101.

```
apic1(config)# leaf 102
apic1(config-leaf)# interface port-channel foo fex 101
apic1(config-leaf)# shut
```

Configure ports to a port channel in multiple leaf nodes. In this example, port channel foo is assigned to ports Ethernet 1/1-2 in each of the leaf nodes 101-103. The leaf nodes will auto generate a number unique in each node (which may be same or different among nodes) to represent the port-channel interfaces.

```
apic1(config)# leaf 101-103
apic1(config-leaf)# interface ethernet 1/1-2
apic1(config-leaf-if)# channel-group foo
```

Add members to port channels. This example would add two members eth1/3-4 to the port-channel in each leaf node, so that port-channel foo in each node would have members eth 1/1-4.

```
apic1(config)# leaf 101-103
apic1(config-leaf)# interface ethernet 1/3-4
apic1(config-leaf-if)# channel-group foo
```

Remove members from port channels. This example would remove two members eth1/2, eth1/4 from the port channel foo in each leaf node, so that port channel foo in each node would have members eth 1/1, eth1/3.

```
apic1(config)# leaf 101-103
apic1(config-leaf)# interface eth 1/2,1/4
apic1(config-leaf-if)# no channel-group foo
```

Configure port-channel with different members in multiple leaf nodes. This example shows how to use the same port-channel foo policies to create a port-channel interface in multiple leaf nodes with different member ports in each leaf. The port-channel numbers in the leaf nodes may be same or different for the same port-channel foo. In the CLI, however, the configuration will be referred as interface port-channel foo. If the port-channel is configured for the FEX ports, it would be referred to as interface port-channel foo fex <fex-id>.

```
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/1-2
apic1(config-leaf-if)# channel-group foo
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
apic1(config)# leaf 102
apic1(config-leaf)# interface ethernet 1/3-4
apic1(config-leaf-if)# channel-group foo
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
apic1(config)# leaf 103
apic1(config-leaf)# interface ethernet 1/5-8
apic1(config-leaf-if)# channel-group foo
apic1(config-leaf-if)# exit
apic1(config-leaf)# interface ethernet 101/1/1-2
apic1(config-leaf-if)# channel-group foo
```

Configure per port properties for LACP. This example shows how to configure member ports of a port-channel for per-port properties for LACP.



**Note** In ACI model, these commands are allowed only after the ports are member of a port channel. If a port is removed from a port channel, configuration of these per-port properties would be removed as well.

```
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/1-2
apic1(config-leaf-if)# channel-group foo
apic1(config-leaf-if)# lacp port-priority 1000
apic1(config-leaf-if)# lacp rate fast
```

Configure admin state for port channels. In this example, a port-channel foo is configured in each of the leaf nodes 101-103 using the channel-group command. The admin state of port-channel(s) can be configured in each leaf using the port-channel interface. In ACI model, the admin state of the port-channel cannot be configured in the global scope.

```
// create port-channel foo in each leaf
apicl(config)# leaf 101-103
apicl(config-leaf)# interface ethernet 1/3-4
apicl(config-leaf-if)# channel-group foo

// configure admin state in specific leaf
apicl(config)# leaf 101
apicl(config-leaf)# interface port-channel foo
apicl(config-leaf-if)# shut
```

Override config is very helpful to assign specific vlan-domain, for example, to the port-channel interfaces in each leaf while sharing other properties.

```
// configure a port channel global config
apicl(config)# interface port-channel foo
apicl(config-if)# speed 1G
apicl(config-if)# channel-mode active

// create port-channel foo in each leaf
apicl(config)# leaf 101-103
apicl(config-leaf)# interface ethernet 1/1-2
apicl(config-leaf-if)# channel-group foo

// override port-channel foo in leaf 102
apicl(config)# leaf 102
apicl(config-leaf)# interface port-channel foo
apicl(config-leaf-if)# speed 10G
apicl(config-leaf-if)# channel-mode on
apicl(config-leaf-if)# vlan-domain dom-foo
```

This example shows how to change port channel assignment for ports using the channel-group command. There is no need to remove port channel membership before assigning to other port channel.

```
apicl(config)# leaf 101-103
apicl(config-leaf)# interface ethernet 1/3-4
apicl(config-leaf-if)# channel-group foo
apicl(config-leaf-if)# channel-group bar
```

## Configuring Virtual Port Channels in Leaf Nodes and FEX Devices Using the NX-OS CLI

A virtual port channel (vPC) is an enhancement to port-channels that allows connection of a host or switch to two upstream leaf nodes to improve bandwidth utilization and availability. In NX-OS, vPC configuration is done in each of the two upstream switches and configuration is synchronized using peer link between the switches.



**Note** When creating a vPC domain between two leaf switches, both switches must be in the same switch generation, one of the following:

- Generation 1 - Cisco Nexus N9K switches without “EX” on the end of the switch name; for example, N9K-9312TX
- Generation 2 – Cisco Nexus N9K switches with “EX” on the end of the switch model name; for example, N9K-93108TC-EX

Switches such as these two are not compatible vPC peers. Instead, use switches of the same generation.

The ACI model does not require a peer link and vPC configuration can be done globally for both the upstream leaf nodes. A global configuration mode called **vpc context** is introduced in ACI and vPC interfaces are represented using a type **interface vpc** that allows global configuration applicable to both leaf nodes.

Two different topologies are supported for vPC in the ACI model: vPC using leaf ports and vPC over FEX ports. It is possible to create many vPC interfaces between a pair of leaf nodes and similarly, many vPC interfaces can be created between a pair of FEX modules attached to the leaf node pairs in a straight-through topology.

vPC considerations include:

- The vPC name used is unique between leaf node pairs. For example, only one vPC 'corp' can be created per leaf pair (with or without FEX).
- Leaf ports and FEX ports cannot be part of the same vPC.
- Each FEX module can be part of only one instance of vPC corp.
- vPC context allows configuration
- The vPC context mode allows configuration of all vPCs for a given leaf pair. For vPC over FEX, the *fex-id* pairs must be specified either for the vPC context or along with the vPC interface, as shown in the following two alternative examples.

```
(config)# vpc context leaf 101 102
(config-vpc)# interface vpc Reg fex 101 101
```

or

```
(config)# vpc context leaf 101 102 fex 101 101
(config-vpc)# interface vpc Reg
```

In the ACI model, vPC configuration is done in the following steps (as shown in the examples below).



**Note** A VLAN domain is required with a VLAN range. It must be associated with the port-channel template.

1. VLAN domain configuration (global config) with VLAN range
2. vPC domain configuration (global config)

3. Port-channel template configuration (global config)
4. Associate the port-channel template with the VLAN domain
5. Port-channel configuration for vPC (global config)
6. Configure ports to vPC in leaf nodes
7. Configure L2, L3 for vPC in the vpc context

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b> <b>Example:</b> <code>apicl# configure</code>	Enters global configuration mode.
<b>Step 2</b>	<b>vlan-domain</b> <i>name</i> [dynamic] [ <b>type</b> <i>domain-type</i> ] <b>Example:</b> <code>apicl(config)# vlan-domain dom1 dynamic</code>	Configures a VLAN domain for the virtual port-channel (here with a port-channel template).
<b>Step 3</b>	<b>vlan range</b> <b>Example:</b> <code>apicl(config-vlan)# vlan 1000-1999</code> <code>apicl(config-vlan)# exit</code>	Configures a VLAN range for the VLAN domain and exits the configuration mode. The range can be a single VLAN or a range of VLANs.
<b>Step 4</b>	<b>vpc domain explicit</b> <i>domain-id leaf node-id1 node-id2</i> <b>Example:</b> <code>apicl(config)# vpc domain explicit 1 leaf 101 102</code>	<p>Configures a vPC domain between a pair of leaf nodes. You can specify the vPC domain ID in the explicit mode along with the leaf node pairs.</p> <p>Alternative commands to configure a vPC domain are as follows:</p> <ul style="list-style-type: none"> <li>• <b>vpc domain [consecutive   reciprocal]</b>  The consecutive and reciprocal options allow auto configuration of a vPC domain across all leaf nodes in the ACI fabric.</li> <li>• <b>vpc domain consecutive</b> <i>domain-start leaf start-node end-node</i>  This command configures a vPC domain consecutively for a selected set of leaf node pairs.</li> </ul>
<b>Step 5</b>	<b>peer-dead-interval</b> <i>interval</i> <b>Example:</b>	Configures the time delay the Leaf switch waits to restore the vPC before receiving a response from the peer. If it does not receive

	Command or Action	Purpose
	<code>apic1(config-vpc)# peer-dead-interval 10</code>	a response from the peer within this time, the Leaf switch considers the peer dead and brings up the vPC with the role as a master. If it does receive a response from the peer it restores the vPC at that point. The range is from 5 seconds to 600 seconds. The default is 200 seconds.
<b>Step 6</b>	<b>exit</b> <b>Example:</b> <code>apic1(config-vpc)# exit</code>	Returns to global configuration mode.
<b>Step 7</b>	<b>template port-channel</b> <i>channel-name</i> <b>Example:</b> <code>apic1(config)# template port-channel corp</code>	Creates a new port-channel or configures an existing port-channel (global configuration).  All vPCs are configured as port-channels in each leaf pair. The same port-channel name must be used in a leaf pair for the same vPC. This port-channel can be used to create a vPC among one or more pairs of leaf nodes. Each leaf node will have only one instance of this vPC.
<b>Step 8</b>	<b>vlan-domain member</b> <i>vlan-domain-name</i> <b>Example:</b> <code>vlan-domain member dom1</code>	Associates the port channel template with the previously configured VLAN domain.
<b>Step 9</b>	<b>switchport access vlan</b> <i>vlan-id</i> <b>tenant</b> <i>tenant-name</i> <b>application</b> <i>application-name</i> <b>epg</b> <i>epg-name</i> <b>Example:</b> <code>apic1(config-po-ch-if)# switchport access vlan 4 tenant ExampleCorp application Web epg webEpg</code>	Deploys the EPG with the VLAN on all ports with which the port-channel is associated.
<b>Step 10</b>	<b>channel-mode active</b> <b>Example:</b> <code>apic1(config-po-ch-if)# channel-mode active</code>	<b>Note</b> A port-channel must be in active channel-mode for a vPC.
<b>Step 11</b>	<b>exit</b> <b>Example:</b> <code>apic1(config-po-ch-if)# exit</code>	Returns to configure mode.
<b>Step 12</b>	<b>leaf</b> <i>node-id1</i> <i>node-id2</i> <b>Example:</b> <code>apic1(config)# leaf 101-102</code>	Specifies the pair of leaf switches to be configured.

	Command or Action	Purpose
<b>Step 13</b>	<b>interface</b> <i>type leaf/interface-range</i> <b>Example:</b> <pre>apicl(config-leaf)# interface ethernet 1/3-4</pre>	Specifies the interface or range of interfaces that you are configuring to the port-channel.
<b>Step 14</b>	<b>[no] channel-group</b> <i>channel-name vpc</i> <b>Example:</b> <pre>apicl(config-leaf-if)# channel-group corp vpc</pre>	Assigns the interface or range of interfaces to the port-channel. Use the keyword <b>no</b> to remove the interface from the port-channel. To change the port-channel assignment on an interface, you can enter the <b>channel-group</b> command without first removing the interface from the previous port-channel.  <b>Note</b> The <b>vpc</b> keyword in this command makes the port-channel a vPC. If the vPC does not already exist, a vPC ID is automatically generated and is applied to all member leaf nodes.
<b>Step 15</b>	<b>exit</b> <b>Example:</b> <pre>apicl(config-leaf-if)# exit</pre>	
<b>Step 16</b>	<b>exit</b> <b>Example:</b> <pre>apicl(config-leaf)# exit</pre>	
<b>Step 17</b>	<b>vpc context leaf</b> <i>node-id1 node-id2</i> <b>Example:</b> <pre>apicl(config)# vpc context leaf 101 102</pre>	The vPC context mode allows configuration of vPC to be applied to both leaf node pairs.
<b>Step 18</b>	<b>interface vpc</b> <i>channel-name</i> <b>Example:</b> <pre>apicl(config-vpc)# interface vpc blue fex 102 102</pre>	
<b>Step 19</b>	(Optional) <b>[no] shutdown</b> <b>Example:</b> <pre>apicl(config-vpc-if)# no shut</pre>	Administrative state configuration in the vPC context allows changing the admin state of a vPC with one command for both leaf nodes.

### Example

This example shows how to configure a basic vPC.



```

apic1# configure
apic1(config)# vlan-domain dom1 dynamic

apic1(config-vlan)# vlan 1000-1999
apic1(config-vlan)# exit
apic1(config)# vpc domain explicit 1 leaf 101 102
apic1(config-vpc)# peer-dead-interval 10

apic1(config-vpc)# exit
apic1(config)# template port-channel corp
apic1(config-po-ch-if)# vlan-domain member dom1

apic1(config-po-ch-if)# channel-mode active

apic1(config-po-ch-if)# exit
apic1(config)# leaf 101-102
apic1(config-leaf)# interface ethernet 1/3-4
apic1(config-leaf-if)# channel-group corp vpc
apic1(config-leaf-if)# exit
apic1(config)# vpc context leaf 101 102

```

This example shows how to configure vPCs with FEX ports.

```

apic1(config-leaf)# interface ethernet 101/1/1-2
apic1(config-leaf-if)# channel-group Reg vpc
apic1(config)# vpc context leaf 101 102
apic1(config-vpc)# interface vpc corp
apic1(config-vpc-if)# exit
apic1(config-vpc)# interface vpc red fex 101 101
apic1(config-vpc-if)# switchport
apic1(config-vpc-if)# exit
apic1(config-vpc)# interface vpc blue fex 102 102
apic1(config-vpc-if)# shut

```

## Configuring FEX Connections Using Profiles with the NX-OS Style CLI

Use this procedure to configure FEX connections to leaf nodes using the NX-OS style CLI.



**Note** Configuring FEX connections with FEX IDs 165 to 199 is not supported in the APIC GUI. To use one of these FEX IDs, configure the profile using the following commands.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	configure Example:	Enters global configuration mode.

	Command or Action	Purpose
	<code>apic1# configure</code>	
<b>Step 2</b>	<b>leaf-interface-profile</b> <i>name</i> <b>Example:</b> <code>apic1(config)# leaf-interface-profile fexIntProf1</code>	Specifies the leaf interface profile to be configured.
<b>Step 3</b>	<b>leaf-interface-group</b> <i>name</i> <b>Example:</b> <code>apic1(config-leaf-if-profile)# leaf-interface-group leafIntGrp1</code>	Specifies the interface group to be configured.
<b>Step 4</b>	<b>fex associate</b> <i>fex-id</i> [ <b>template</b> <i>template-type</i> <i>fex-template-name</i> ] <b>Example:</b> <code>apic1(config-leaf-if-group)# fex associate 101</code>	Attaches a FEX module to a leaf node. Use the optional template keyword to specify a template to be used. If it does not exist, the system creates a template with the name and type you specified.

### Example

This merged example configures a leaf interface profile for FEX connections with ID 101.

```
apic1# configure
apic1(config)# leaf-interface-profile fexIntProf1
apic1(config-leaf-if-profile)# leaf-interface-group leafIntGrp1
apic1(config-leaf-if-group)# fex associate 101
```

## Reflective Relay (802.1Qbg)

Reflective relay is a switching option beginning with Cisco APIC Release 2.3(1). Reflective relay—the tagless approach of IEEE standard 802.1Qbg—forwards all traffic to an external switch, which then applies policy and sends the traffic back to the destination or target VM on the server as needed. There is no local switching. For broadcast or multicast traffic, reflective relay provides packet replication to each VM locally on the server.

One benefit of reflective relay is that it leverages the external switch for switching features and management capabilities, freeing server resources to support the VMs. Reflective relay also allows policies that you configure on the Cisco APIC to apply to traffic between the VMs on the same server.

In the Cisco ACI, you can enable reflective relay, which allows traffic to turn back out of the same port it came in on. You can enable reflective relay on individual ports, port channels, or virtual port channels as a Layer 2 interface policy using the APIC GUI, NX-OS CLI, or REST API. It is disabled by default.

The term *Virtual Ethernet Port Aggregator* (VEPA) is also used to describe 802.1Qbg functionality.

### Reflective Relay Support

Reflective relay supports the following:

- IEEE standard 802.1Qbg tagless approach, known as reflective relay.

Cisco APIC Release 2.3(1) release does not support the IEEE standard 802.1Qbg S-tagged approach with multichannel technology.

- Physical domains.

Virtual domains are not supported.

- Physical ports, port channels (PCs), and virtual port channels (vPCs).

Cisco Fabric Extender (FEX) and blade servers are not supported. If reflective relay is enabled on an unsupported interface, a fault is raised, and the last valid configuration is retained. Disabling reflective relay on the port clears the fault.

- Cisco Nexus 9000 series switches with *EX* or *FX* at the end of their model name.

## Enabling Reflective Relay Using the NX-OS CLI

Reflective relay is disabled by default; however, you can enable it on a port, port channel, or virtual port channel as a Layer 2 interface policy on the switch. In the NX-OS CLI, you can use a template to enable reflective relay on multiple ports or you can enable it on individual ports.

### Before you begin

This procedure assumes that you have set up the Cisco Application Centric Infrastructure (ACI) fabric and installed the physical switches.

### Procedure

---

Enable reflective relay on one or multiple ports:

#### Example:

This example enables reflective relay on a single port:

```
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/2
apic1(config-leaf-if)# switchport vepa enabled
apic1(config-leaf-if)# exit
apic1(config-leaf)# exit
```

#### Example:

This example enables reflective relay on multiple ports using a template:

```
apic1(config)# template policy-group grp1
apic1(config-pol-grp-if)# switchport vepa enabled
apic1(config-pol-grp-if)# exit
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/2-4
apic1(config-leaf-if)# policy-group grp1
```

#### Example:

This example enables reflective relay on a port channel:

```
apic1(config)# leaf 101
apic1(config-leaf)# interface port-channel po2
apic1(config-leaf-if)# switchport vepa enabled
apic1(config-leaf-if)# exit
```

```
apicl(config-leaf)# exit
apicl(config)#
```

**Example:**

This example enables reflective relay on multiple port channels:

```
apicl(config)# template port-channel po1
apicl(config-if)# switchport vepa enabled
apicl(config-if)# exit
apicl(config)# leaf 101
apicl(config-leaf)# interface ethernet 1/3-4
apicl(config-leaf-if)# channel-group po1
apicl(config-leaf-if)# exit
apicl(config-leaf)# exit
```

**Example:**

This example enables reflective relay on a virtual port channel:

```
apicl(config)# vpc domain explicit 1 leaf 101 102
apicl(config-vpc)# exit
apicl(config)# template port-channel po4
apicl(config-if)# exit
apicl(config)# leaf 101-102
apicl(config-leaf)# interface eth 1/11-12
apicl(config-leaf-if)# channel-group po4 vpc
apicl(config-leaf-if)# exit
apicl(config-leaf)# exit
apicl(config)# vpc context leaf 101 102
apicl(config-vpc)# interface vpc po4
apicl(config-vpc-if)# switchport vepa enabled
```

## Configuring Policy Groups for Interfaces

In data center networks, typically configuration of many interfaces is the same across multiple nodes. This can be achieved in the ACI Policy Model by creating policy-groups to be shared by groups of interfaces across multiple leaf nodes. The policy-group is identified by a name similar to the port-channel; however, in case of port-channel the policies shared with the group of ports create one logical interface in each leaf while in case of a policy-group, each of the ports sharing the policies are individual physical interfaces. The policy-group concept is very similar to a port-profile in NX-OS.

**Procedure**

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b> <b>Example:</b> apicl# <b>configure</b>	Enters global configuration mode.
<b>Step 2</b>	<b>template policy-group</b> <i>policy-group-name</i> <b>Example:</b> apicl(config)# <b>template policy-group pgl</b>	Creates a new policy group or edits an existing policy group.

	Command or Action	Purpose
<b>Step 3</b>	<p>[no] <b>switchport access vlan</b> <i>vlan-id</i> <b>tenant</b> <i>tenant-name</i> <b>application</b> <i>application-name</i> <b>epg</b> <i>epg-name</i></p> <p><b>Example:</b></p> <pre>apicl(config-pol-grp-if) # <b>switchport</b> <b>access vlan 4 tenant ExampleCorp</b> <b>application Web epg webEpg</b></pre>	
<b>Step 4</b>	<p>(Apply configuration commands)</p> <p><b>Example:</b></p> <pre>apicl(config-pol-grp-if) # <b>speed 10G</b> apicl(config-pol-grp-if) # <b>cdp enable</b></pre>	The table at the end of these steps shows various commands for configurations of policy-group for interfaces.
<b>Step 5</b>	<p><b>exit</b></p> <p><b>Example:</b></p> <pre>apicl(config-pol-grp-if) # <b>exit</b></pre>	Returns to configure mode.
<b>Step 6</b>	<p><b>leaf</b> <i>node-id</i></p> <p><b>Example:</b></p> <pre>apicl(config) # <b>leaf 101-103</b></pre>	Specifies the leaf or leaves to be configured. The <i>node-id</i> can be a single node ID or a range of IDs, in the form <i>node-id1 – node-id2</i> , to which the configuration will be applied.
<b>Step 7</b>	<p><b>interface</b> <i>type</i></p> <p><b>Example:</b></p> <pre>apicl(config-leaf) # <b>interface ethernet</b> <b>1/1-24</b></pre>	Specifies the interface or range of interfaces to which you will apply the policy group.
<b>Step 8</b>	<p>[no] <b>policy-group</b> <i>policy-group-name</i> [<b>force</b>]</p> <p><b>Example:</b></p> <pre>apicl(config-leaf-if) # <b>policy-group pgl</b></pre>	<p>Applies the policy-group to the interface or range of interfaces. Use the keyword <b>no</b> to remove the policy-group from the interface. Use the keyword <b>force</b> to delete any override configurations on the interfaces.</p> <p>If the specified policy-group was not configured prior to this command, this command would not implicitly create the policy-group. However, the policy-group would take effect on the interface after the policy-group has been configured in the global scope.</p> <p>To change the policy-group assignment on an interface, you can enter the <b>policy-group</b> command without first removing the previous policy-group from the interface.</p>

	Command or Action	Purpose
		<b>Note</b> If you apply a policy-group to an interface and then assign the interface to a port-channel, the interface will lose the policy-group configuration and the policies in the port-channel will be applied.

The following table shows various commands for configurations of policy-group for interfaces.

CLI Syntax	Feature
[no] speed <speedValue>	Set the speed for Physical Interface
[no] link debounce time <time>	Set link debounce for Physical Interface
[no] negotiate auto	Configure Negotiate for Physical Interface
[no] cdp enable	Disable/Enable CDP for Physical Interface
[no] mcp enable	Disable/Enable MCP for Physical Interface
[no] lldp transmit	Set the LLDP transmit for Physical Interface
[no] lldp receive	Set the LLDP receive for Physical Interface
spanning-tree <bpduguard   bpdufilter> <enable   disable>	Configure spanning tree BPDU
[no] storm-control level <percentage> [ burst-rate <percentage> ]	Storm-control configuration (percentage)
[no] storm-control pps <packet-per-second> burst-rate <packets-per-second>	Storm-control configuration (packets-per-second)

### Example

This example shows how to configure a policy-group and apply it to a range of ports in each of the leaf nodes 101-103. Each of the ports sharing the policy-group in each leaf will have the same configuration as defined in the policy-group pg1.

```
apic1# configure
apic1(config)# template policy-group pg1
apic1(config-pol-grp-if)# switchport access vlan 4 tenant ExampleCorp application Web epg
webEpg
apic1(config-pol-grp-if)# speed 10G
apic1(config-pol-grp-if)# cdp enable
apic1(config-pol-grp-if)# exit
apic1(config)# leaf 101-103
apic1(config-leaf)# interface ethernet 1/1-24
apic1(config-leaf-if)# policy-group pg1
```

# Configuring Overrides for Interfaces

When policy-groups are used with large number of interfaces, it may be useful to have the option to configure a set of ports for specific properties that will override the configuration in the assigned policy-group. Override configuration is allowed only if the port is assigned to a policy-group. Override configuration is not allowed for member ports of a port-channel. When a port is added to a port-channel, the override configuration is automatically removed. However, during policy-group assignment to a port that has overrides configured, the override configuration is not removed automatically and the user can decide to remove the override configuration with the force option, if required, in the policy-group command.

When a policy-group assignment is removed from a port, the override config, if exists, does not change. Similarly, the override config does not change if the port is assigned to a different policy-group (without the force option). The override config takes effect once configured and it is not removed even if the user assigns default values to all the properties in the override. To remove the override config, the user can reapply the policy-group assignment with force option. The force option, however, is not displayed in the show running-config as it is used to just remove the override config in the ACI model.

In the ACI model, overrides can be configured for a policy which may contain one or more properties. If a policy has more than one property, it is not possible to override only one property within a policy. In the CLI framework, when the user intends to override a property for which the corresponding policy has more than one property, all other properties in the policy except the override property would be implicitly copied to the override configuration to avoid ambiguity. Such implicit copy of configuration would be reflected in the output of show running-config regardless of the value (including default values). Also, the copy is done only once during the configuration of the override policy and any subsequent change to the policy-group for any of the properties in that policy would have no effect on the port(s) on which the override is configured.

If the policy-group assigned to a port is not configured when the override is created, the implicit copy of properties noted above is not possible; instead, default values are assigned to properties in the override config for which the corresponding policy has more than one property. These properties shall not change for the override config when the policy-group is configured afterwards. It is recommended that user create overrides after configuring the policy-group itself or the user may need to configure the overrides in addition to the config in policy-group to get desired configuration if the config for properties in override are set to default implicitly before the configuration of the policy-group with non-default values for those properties.

## Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> apicl# <b>configure</b>	Enters global configuration mode.
<b>Step 2</b>	<b>leaf node-id</b>  <b>Example:</b> apicl(config)# <b>leaf 102</b>	Specifies the leaf or leaves to be configured. The <i>node-id</i> can be a single node ID or a range of IDs, in the form <i>node-id1 – node-id2</i> , to which the configuration will be applied.
<b>Step 3</b>	<b>interface type</b>  <b>Example:</b> apicl(config-leaf)# <b>interface ethernet 1/2</b>	Specifies the interface or range of interfaces with an override configuration.

	Command or Action	Purpose
<b>Step 4</b>	<b>policy-group</b> <i>policy-group-name</i> <b>force</b> <b>Example:</b> apic1(config-leaf-if)# <b>policy-group pgl force</b>	Forces the policy-group to the interface or range of interfaces, deleting any override configurations on the interfaces.

### Examples

This example shows how to apply a policy-group and then override the speed configuration for port eth1/1 in leaf node 101. In the ACI model, speed is part of a policy that also contains properties autoneg and link debounce time. As a result, those properties are copied from the speed policy-group when the override of pgl is configured.

```

apic1# configure
apic1(config)# interface policy-group pgl
apic1(config-pol-grp-if)# speed 10G
apic1(config-pol-grp-if)# cdp enable
apic1(config-pol-grp-if)# exit
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/1-2
apic1(config-leaf-if)# policy-group pgl
apic1(config-pol-grp-if)# exit
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/1
apic1(config-leaf-if)# speed 1G
apic1(config-leaf-if)# show running-config

leaf 101
  interface ethernet 1/1
    policy-group pgl
    speed 1G
    autoneg on
    link debounce time 100

  interface ethernet 1/2
    policy-group pgl

```

This example shows how to remove the override configuration from port eth1/1 in leaf node 101.

```

apic1# configure
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/1
apic1(config-leaf-if)# policy-group pgl force
apic1(config-leaf-if)# show running-config

leaf 101
  interface ethernet 1/1
    policy-group pgl

```



## About Forwarding Error Correction

Forwarding Error Correction (FEC) is a method of obtaining error control in data transmission over an unreliable or noisy channel in which the source (transmitter) encodes the data in a redundant way using Error Correcting Code, and the destination (receiver) recognizes it and corrects the errors without requiring a retransmission. The available options are as follows:

- CL74-FC-FEC—Supports 25 Gbps speed.
- CL91-RS-FEC—Supports 25 and 100 Gbps speeds.
- Disable-FEC—Disables FEC.
- Inherit—The switch uses FEC based on the port transceiver type. All copper (CR4) transceivers have FC-FEC enabled on 25G. All interfaces with 100G transceivers have RS-FEC enabled.

The default is "Inherit".



**Note** FEC is only configurable on the front port and not on fabric ports.

## Configuring FEC Using NX-OS Style CLI

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	Enter the configure mode. <b>Example:</b> apicl# <b>configure</b>	Enters the configuration mode.
<b>Step 2</b>	Enter the switch mode. <b>Example:</b> apicl(config)# <b>leaf 104</b>	Enters the switch mode.
<b>Step 3</b>	Specify the interface and port. <b>Example:</b> apicl(config-leaf)# <b>int eth 1/4</b>	Specifies the interface and port.
<b>Step 4</b>	Configure FEC. <b>Example:</b> apicl(config-leaf-if)# <b>forward-error-correction cl91-rs-fec</b>	Configures RS-FEC. <b>Note</b> The default forward-error-correction value is inherit.
<b>Step 5</b>	Exit the interface mode. <b>Example:</b> apicl(config-leaf-if)# <b>exit</b>	Exits the interface mode.

