

Breakout Ports

This chapter contains the following sections:

- Configuration of Breakout Ports, on page 1
- Guidelines and Limitations for Dynamic Breakout Ports on Down Links, on page 2
- Guidelines and Limitations for Auto Breakout Ports on Fabric Links, on page 6
- Configuring Breakout Ports With Profiles and Selectors Using the GUI, on page 8
- Configuring Breakout Ports With Profiles and Selectors Using the GUI, on page 11
- Configuring Breakout Ports With the Interface Configuration Using the GUI, on page 14
- Configuring Dynamic Breakout Ports Using the NX-OS Style CLI, on page 14

Configuration of Breakout Ports

Breakout cables are suitable for very short links and offer a cost effective way to connect within racks and across adjacent racks. Breakout enables a 40 Gigabit (Gb) port to be split into four independent and logical 10Gb ports, a 100Gb port to be split into four independent and logical 25Gb ports, or a 400Gb port to be split into four independent and logical 100Gb ports.

You configure breakout on the down links (also known as the access-facing ports or downlink ports) and fabric links of the switches. Fabric links form the connections between the leaf switches and spine switches, or between the tier 1 leaf switches and tier 2 leaf switches for a multi-tier topology.

You can configure breakout ports in the following ways:

- You can use port profiles and selectors. With this method, you configure a breakout leaf port with an leaf interface profile, associate the profile with a switch, and configure the sub-ports.
- Beginning with the Cisco Application Policy Infrastructure Controller (APIC) 6.0(1) release, you can use the **Fabric** > **Access Policies** > **Interface Configuration** workflow.
- You can use the **Fabric** > **Inventory** > ***pod*** > ***leaf_name*** workflow. Beginning with the Cisco APIC 6.0(1) release, the inventory view configuration also uses the interface configuration.

Guidelines and Limitations for Dynamic Breakout Ports on Down Links

The 40Gb to 10Gb dynamic breakout feature is supported on the access-facing ports of the following switches:

- N9K-C93180LC-EX
- N9K-C93180YC-FX
- N9K-C9336C-FX2
- N9K-C93360YC-FX2
- N9K-C93216TC-FX2
- N9K-C93108TC-FX3P
- N9K-C93180YC-FX3
- N9K-C93600CD-GX
- N9K-C9364C-GX
- N9K-C9408 (beginning in the 6.0(2) release)
- N9K-C9348D-GX2A (beginning in the 6.0(3) release)
- N9K-C9364D-GX2A (beginning in the 6.0(3) release)
- N9K-C9332D-GX2B (beginning in the 6.0(3) release)

The 100Gb to 25Gb breakout feature is supported on the access ports of the following switches:

- N9K-C93180LC-EX
- N9K-C9336C-FX2
- N9K-C93180YC-FX
- N9K-C93360YC-FX2
- N9K-C93216TC-FX2
- N9K-C93108TC-FX3P
- N9K-C93180YC-FX3
- N9K-C93600CD-GX
- N9K-C9364C-GX
- N9K-C9408 (beginning in the 6.0(2) release)
- N9K-C9348D-GX2A (beginning in the 6.0(3) release)
- N9K-C9364D-GX2A (beginning in the 6.0(3) release)
- N9K-C9332D-GX2B (beginning in the 6.0(3) release)

The 400Gb to 100Gb breakout feature is supported on the access ports of the following switches:

- N9K-C9348D-GX2A
- N9K-C9364D-GX2A
- N9K-C9332D-GX2B
- N9K-C93600CD-GX
- N9K-C9316D-GX
- N9K-C9408 (beginning in the 6.0(2) release)
- Beginning with the 6.0(2) release, the QDD-400G-SR4.2-BD optic is supported on 400Gb ports. The peer node with the 100Gb speed must use QSFP-100G-SR1.2 optics.
- Beginning with the 6.0(3) release, 400Gb-to-100Gb breakout using a QDD-4ZQ100 cable of any length is supported only with the following platforms:
  - N9K-C9332D-GX2B, N9K-C9348D-GX2A, N9K-C9364D-GX2A, N9K-C9316D-GX, and N9K-C93600CD-GX leaf and spine switches
  - N9K-X9716D-GX line card
- The N9K-C9400-SW-GX2A switch supports 400Gb breakout with the N9K-X9400-8D line card expansion module and the peer 100G-PCC is supported with the N9K-X9400-8D and N9K-X9400-16W line card expansion modules.
- QDD-400G-DR4-S, QDD-4X100G-FR-S, QDD-4X100G-LR-S optics are supported on 400Gb ports. The peer node with the 100Gb speed can use the following optics:
  - QSFP-100G-DR-S
  - QSFP-100G-FR-S
  - QSFP-100G-LR-S

Before you configure breakout ports, connect a 40Gb port to four 10Gb ports, a 100Gb port to four 25Gb ports, or a 400Gb port to four 100Gb ports with one of the following cables:

- Cisco QSFP-4SFP10G

  Beginning in the 6.0(3) release, GX2 switches support this direct-attach cable.
- Cisco QSFP-4SFP25G

  Beginning in the 6.0(3) release, GX2 switches support this direct-attach cable.
- Cisco QSFP-4X10G-AOC

  Beginning in the 6.0(3) release, GX2 switches support this direct-attach cable.
- Cisco QDD-4ZQ100-CU (1M, 2M, 2.5M, and 3M) (beginning in the 6.0(3) release)
- MPO to breakout splitter cable with QSFP-40G-SR4 and 4 x SFP-10G-SR on the ends
- MPO to breakout splitter cable with QSFP-100G-SR4-S and 4 x SFP-25G-SR-S on the ends
- MPO to breakout splitter cable with QDD-400G-DR4-S, QDD-4X100G-FR-S, or QDD-4X100G-LR-S, and 4 x QSFP-100G-DR-S, 4 x QSFP-100G-FR-S, or 4 x QSFP-100G-LR-S on the ends

- MPO to breakout splitter cable with QDD-400G-SR4.2-BD and 4 x QSFP-100G-SR1.2 on the ends.



Note For the supported optics and cables, see the *Cisco Optics-to-Device Compatibility Matrix*:

https://tmgmatrix.cisco.com/

Observe the following guidelines and limitations:

- Breakout ports are supported on both down links and converted down links.
- The following switches support dynamic breakouts (both 100Gb and 40Gb) on profiled QSFP ports:
  - Cisco N9K-C93180YC-FX
  - Cisco N9K-C93216TC-FX2
  - Cisco N9K-C93360YC-FX2
  - Cisco N9K-C93600CD-GX

    This applies only to ports 1/25 to 34. Ports 1/29 to 34 can be used for dynamic breakouts if the ports are converted to down links.
  - Cisco N9K-C9336C-FX2

    You can configure up to 34 dynamic breakouts.
  - Cisco N9K-C9364C-GX

    You can configure up to 30 dynamic breakouts on odd-numbered profiled QSFP ports from 1/1 to 59.
  - Cisco N9K-93600CD-GX

    You can configure up to 12 dynamic breakouts out of 24 40/100G ports and up to 10 dynamic breakouts out of ports 25 to 34. Ports 29 to 34 can be breakouts if the ports are converted to down links. The last 2 ports (ports 35 and 36) are reserved for fabric links.
  - Cisco N9K-C9400-SW-GX2A with the Cisco N9K-X9400-16W linecard

    You can configure breakouts on odd-numbered profiled QSFP ports.
- The Cisco N9K-C9336C-FX2 switch supports LACP fast hello on the breakout sub-port.
- Breakout ports cannot be used for Cisco Application Policy Infrastructure Controller (APIC) connectivity.
- Dynamic breakout configuration on spine switch interfaces is not supported.
- Fast Link Failover policies are not supported on the same port with the dynamic breakout feature.
- Breakout subports can be used in the same way other port types in the policy model are used.
- When a port is enabled for dynamic breakout, other policies (except monitoring policies) on the parent port are no longer valid.
- When a port is enabled for dynamic breakout, other EPG deployments on the parent port are no longer valid.

- A breakout sub-port can not be further broken out using a breakout policy group.
- Dynamic breakout or breakout of a 400Gb port to four 100Gb ports configured using a Cisco APIC policy is supported on QDD-4X100G-FR-S and QDD-4X100G-LR-S optics.
- A breakout sub-port supports LACP. The LACP transmit rate configuration defined in the "default" port channel member policy is used by default. The LACP transmit rate can be changed by either changing the "default" port channel member policy or using an override policy group on each PC/vPC interface policy group.
- Deleting the port profile policy on an already broken out port will clean up the breakout configuration and bring the port back up in the native direction upon switch reload.
- If the LACP transmit rate on port channels that have breakout sub-ports need to be changed, then all the port channels that include breakout sub-ports need to use the same LACP transmit rate configuration. You can configure an override policy to set the transmit rate as follows:
  1. Configure/change the default port channel member policy to include Fast Transmit Rate (**Fabric** > **Access Policies** > **Policies** > **Interface** > **Port Channel Member**).
  2. Configure all the PC/vPC interface policy groups to include the above default port channel member policy under the override policy groups (**Fabric** > **Access Policies** > **Interfaces** > **Leaf Interfaces** > **Policy Groups** > **PC/vPC Interface**).
- The following guidelines and limitations apply to the Cisco N9K-C9364C-GX switch:
  - Odd numbered ports (on rows 1 and row 3) support breakout. Adjacent even numbered ports (on row 2 or row 4) will be disabled ("hw-disabled"). This is applicable to ports 1/1 to 60.
  - The last 2 ports (1/63 and 64) are reserved for fabric links.
  - Ports 1/61 and 62 can be converted to down links, but breakout is not supported. Breakout ports and 40/100G non-breakout ports cannot be mixed in a set of 4 ports starting from 1/1, such as 1/1 to 4 or 1/5 to 8.

    For example, if port 1/1 is breakout enabled, port 1/3 can be used with breakout enabled or native 10G. Port 1/3 will be error-disabled if it is 40/100G.
  - The maximum number of down links are 30 x 4 ports 10/25 (breakout) + 2 ports (1/61 and 62) = 122 ports. Ports 1/63 and 64 are reserved for fabric links and even numbered ports from 1/2 to 60 are error-disabled.
  - This switch supports 10G with QSA on all ports. Native 10G requires QSA.
- The following guidelines and limitations apply to the Cisco N9K-93600CD-GX switch:
  - Odd numbered ports (all ports on row 1) support breakout. Even numbered ports on row 2 will be disabled ("hw-disabled"). This is applicable only to ports 1 to 24.
  - Breakout and 40/100G non-breakout cannot be mixed in a set of 4 ports starting from 1/1 until 1/24, such as 1/1 to 4 or 1/5 to 8. For example:
    - For ports 1/1 to 24, you can have 4 ports per set.

      For example, if port 1/1 is breakout enabled, port 1/3 can be used with breakout enabled or native 10G. Port 1/3 will be error-disabled if it is 40/100G.
    - For ports 1/25 to 28, you can have 2 ports per set.

For example, even if port 1/25 is breakout enabled, port 1/27 can be used with 40/100G.

- The maximum number of down links are 12 x 4 ports 10/25G (breakout) + 10 x 4 ports 10/25G (breakout) = 88 ports. Ports 35 and 36 are reserved for fabric links and 12 ports are disabled.
- This switch supports 10G with QSA on all ports. Native 10G requires QSA.

Guidelines and Limitations for Auto Breakout Ports on Fabric Links

When you insert a transceiver into a line card on which breakout is supported, the ports become broken out automatically. You do not need to configure breakout manually.

The 400G to 100G breakout feature is supported on the fabric ports of the following line card:

- N9K-X9716D-GX with the QDD-4X100G-FR-S, QDD-4X100G-LR-S, or QDD-400G-SR4.2-BD transceiver

The 400G to 100G breakout feature is supported on the fabric ports of the following switches:

- N9K-C9348D-GX2A
- N9K-C9364D-GX2A
- N9K-C9332D-GX2B
- N9K-C93600CD-GX
- N9K-C9316D-GX
- N9K-C9408 (beginning in the 6.0(2) release)
- Beginning with the 6.0(2) release, the QDD-400G-SR4.2-BD optic is supported on 400G ports. The peer node with the 100G speed must use QSFP-100G-SR1.2 optics.
- Beginning with the 6.0(3) release, 400G-to-100G breakout using a QDD-4ZQ100 cable of any length is supported only with the following platforms:
  - N9K-C9332D-GX2B, N9K-C9348D-GX2A, N9K-C9364D-GX2A, N9K-C9316D-GX, and N9K-C93600CD-GX leaf and spine switches
  - N9K-X9716D-GX line card
- The N9K-C9400-SW-GX2A switch supports 400G breakout with the N9K-X9400-8D line card expansion module and the peer 100G-PCC is supported with the N9K-X9400-8D and N9K-X9400-16W line card expansion modules.
- QDD-4X100G-FR-S and QDD-4X100G-LR-S optics are supported on 400G ports. The peer node with the 100G speed can use the following optics:
  - QSFP-100G-DR-S
  - QSFP-100G-FR-S
  - QSFP-100G-LR-S

Connect the ports with one of the following cables:

- Cisco QDD-4ZQ100-CU (1M, 2M, 2.5M, and 3M) (beginning in the 6.0(3) release)
- MPO to 4xLC breakout splitter cable with QDD-4X100G-FR-S or QDD-4X100G-LR-S and 4 x QSFP-100G-DR-S, 4 x QSFP-100G-FR-S, or 4 x QSFP-100G-LR-S on the ends
- MPO to breakout splitter cable with QDD-400G-SR4.2-BD and 4 x QSFP-100G-SR1.2 on the ends.

Observe the following general guidelines and limitations:

- Deleting the port profile policy on an already broken out port will clean up the breakout configuration and bring the port back up in the native direction upon switch reload.

Observe the following guidelines and limitations for 400G to 4x100G breakout on fabric links:

- To change a fabric port with a QDD-400G-SR4.2-BD optic from non-breakout to breakout, you must remove the transceiver, connect the breakout cable to the transceiver, then reinsert the transceiver.
- GX2 switches support the following switch to switch connectivity:
  - Spine switch-to-leaf switch
  - Leaf switch-to-spine switch
  - Leaf switch-to-leaf switch (multi-tier)
- GX line cards support the following switch to switch connectivity:
  - Spine switch to leaf switch
- Spine switch-to-spine switch breakout is not supported.
- Spine switch-to-IPN or ISN breakout is only supported with Auto Breakout and only with the following 100G and 400G optics.
  - QDD-4X100G-FR-S or QDD-4X100G-LR-S optic on the spine switch and the compatible 100G optics such as QSFP-100G-FR-S and QSFP-100G-LR-S on the IPN or ISN side.
  - QDD-400G-SR4.2-BD optic on the spine switch and the QSFP-100G-SR1.2 optic on the IPN or ISN side.
  - The link does not come up when QDD-400G-SR4.2-BD optic on both the sides.
- In a Cisco ACI environment, the operational requirements for the QDD-400G-SR4.2-BD transceiver vary significantly depending on whether the interface is configured as a Fabric or Access port and the specific nature of the peer device.

  1. Fabric Port Behavior

  Fabric ports utilize automated discovery mechanisms that differ from standard access ports.

  - ACI-to-ACI Connections:

    In native 400G mode, the system automatically performs a breakout upon detecting the transceiver and, following the successful exchange and validation of LLDP TLVs across all four subports, updates the configuration to transition the interface to native 400G mode.

To operate in breakout mode, the local device must be equipped with a QDD-400G-SR4.2-BD transceiver, while the peer device must utilize 4x 100G-SR1.2 transceivers.

- ACI-to-Non-ACI Connections:

  In native 400G mode, the link will fail to establish because the ACI fabric port attempts an auto-breakout that the non-ACI peer does not support, resulting in a known configuration mismatch.

  To operate in breakout mode, compatibility is ensured by using the QDD-400G-SR4.2-BD transceiver on the ACI side and 4x 100G-SR1.2 transceivers on the peer device, mirroring the requirements for ACI-to-ACI connections.

2. Access Port Behavior

   Access ports do not rely on the same automated fabric discovery as fabric ports and require explicit configuration.

   - In native 400G mode, a standard QDD-400G-SR4.2-BD connection is required between both devices.

     In breakout mode, the ACI access port requires the manual application of a 4x 100G breakout policy to correctly initialize the interface.

- If you downgrade to a release that does not support breakout for the given hardware and ports, the breakout ports will become not broken out and the links will be down. If all connections between a spine and leaf switch are only breakout, then upon downgrading to a release that does not support breakout, the links will go down and the node will be out of the fabric.
- If you remove an SFP transceiver from a switch, you must wait at least 15 seconds before re-adding the transceiver.
- With the Cisco Nexus 9300 GX2 series or Cisco N9K-X9716D-GX line cards, ports will not come up if an optic is replaced while the line card is powered off. For example:

  1. You have a Cisco N9K-X9716D-GX line card in slot 4 with a 4x100-FR-S transceiver inserted in a port (for example, port 8). Port 8 is in broken out into four ports (Eth4/8/1-4) due to the auto-breakout feature that automatically activates when a 4x100-FR-S transceiver is inserted.
  2. You power off the line card in slot 4.
  3. While the line card is powered off, you remove the 4x100G-FR-S optic from port 8 and insert another optics other than the 4x100G-FR-S.
  4. You power on the line card in slot 4. Port Eth4/8 will not come up even after connecting it with any compatible port and transceiver combination on the peer end.

Configuring Breakout Ports With Profiles and Selectors Using the GUI

This procedure uses port profiles and selectors to configure breakout ports. You configure a breakout leaf port with an leaf interface profile, associate the profile with a switch, and configure the sub-ports.

Before you begin

- The Cisco Application Centric Infrastructure (ACI) fabric is installed, the Cisco Application Policy Infrastructure Controllers (APICs) are online, and the Cisco APIC cluster is formed and healthy.
- A Cisco APIC fabric administrator account is available that can configure the breakout ports.
- The target leaf switches are registered in the Cisco ACI fabric and available.
- The 40GE or 100GE leaf switch ports are connected with Cisco breakout cables to the downlink ports.

Procedure

Step 1 On the menu bar, choose **Fabric** > **Access Policies**.

Step 2 In the Navigation pane, expand **Interfaces** and **Leaf Interfaces** and **Profiles**.

Step 3 Right-click **Profiles** and choose **Create Leaf Interface Profile**.

Step 4 Type the name and optional description, click the + symbol on **Interface Selectors**

Step 5 Perform the following:

a) Type a name (and optional description) for the **Access Port Selector**.
b) In the **Interface IDs** field, type the slot and port for the breakout port.
c) In the **Interface Policy Group** field, click the down arrow and choose **Create Leaf Breakout Port Group**.
d) Type the name (and optional description) for the **Leaf Breakout Port Group**.
e) In the **Breakout Map** field, choose **10g-4x** or **25g-4x**.

For a list of the switches that support breakout, see Configuration of Breakout Ports, on page 1.

f) Click **Submit**.

Step 6 To assign a Breakout Port to an EPG, perform the following steps:

On the menu bar, choose **Tenant** > **Application Profiles** > **Application EPG**. Right-click on **Application EPG** to open **Create Application EPG**dialog box, and perform the following steps:

a) Select the **Statically Link with Leaves/Paths** check box to gain access to the **Leaves/Paths** tab in the dialog box.
b) Complete one of the following sets of steps:

Option	Description
If you want to deploy the EPG on...	Then
A node	**a.** Expand the **Leaves** area. **b.** From the **Node** drop-down list, choose a node. **c.** In the **Encap** field, enter the appropriate VLAN. **d.** (Optional) From the **Deployment Immediacy** drop-down list, accept the default **On Demand** or choose **Immediate**. **e.** (Optional) From the Mode drop-down list, accept the default **Trunk** or choose another mode.
A port on the node	**a.** Expand the **Paths** area.

Option	Description
	b. From the **Path** drop-down list, choose the appropriate node and port. **c.** (Optional) In the **Deployment Immediacy** field drop-down list, accept the default **On Demand** or choose **Immediate**. **d.** (Optional) From the Mode drop-down list, accept the default **Trunk** or choose another mode. **e.** In the **Port Encap** field, enter the secondary VLAN to be deployed. **f.** (Optional) In the **Primary Encap** field, enter the primary VLAN to be deployed.

Step 7 To associate the Leaf Interface Profile to a the leaf switch, perform the following steps:

a) Expand **Switches** and **Leaf Switches**, and **Profiles**.
b) Right-click **Profiles** and select **Create Leaf Profiles**.
c) Type the name and optional description of the Leaf Profile.
d) Click the **+** symbol on the **Leaf Selectors** area.
e) Type the leaf selector name and an optional description.
f) Click the down arrow on the **Blocks** field and choose the switch to be associated with the breakout leaf interface profile.
g) Click the down arrow on the **Policy Group** field and choose **Create Access Switch Policy Group**.
h) Type a name and optional description for the Access Switch Policy Group.
i) Optional. Enable other policies.
j) Click **Submit**.
k) Click **Update**.
l) Click **Next**.
m) In the **Associations Interface Selector Profiles** area, choose the Interface Selector Profile you previously created for the breakout port.
n) Click **Finish**.

Step 8 To verify the breakout port has been split into four sub ports, perform the following steps:

a) On the Menu bar, click **Fabric** > **Inventory**.
b) On the Navigation bar, click the Pod and Leaf where the breakout port is located.
c) Expand **Interfaces** and **Physical Interfaces**.
You should see four ports at the position where the breakout port was configured. For example, if you configured 1/10 as a breakout port, you should see the following:

- **eth1/10/1**
- **eth1/10/2**
- **eth1/10/3**
- **eth1/10/4**

Step 9 To configure the sub ports, perform the following steps:

a) On the Menu bar, click **Fabric** > **Access Policies**.
b) On the Navigation bar, expand **Interfaces**, **Leaf Interfaces**, **Profiles**, and the breakout leaf interface profile you previously created.

You will see a port selector with the breakout cable. Instead of defining a sub port block under already existing port selector, you need to define on a new access port selector.

c) On the Navigtion bar, right click the higher level interface profile and select **Create Access Port Selector**.
d) In the **Name** field, enter the sub port name.
e) In the **Interface IDs** field, enter the IDs for the four sub ports in a format such as `1/10/1-4`.
f) In the **Interface Policy Group** field, select **Create Leaf Access Port Policy Group**.
g) Click **Submit**.

Step 10 To apply the Policy Group to an individual interface that links the AAEP to the port, perform the following steps:

a) In the **Name** field, enter the name for the Leaf Access Port Group Policy.
b) In the **Link Level Policy** field, select **link-level_auto**.
c) In the **CDP Policy** field, select **cdp_enabled**.
d) In the **LLDP Policy** field, select **default**.
e) In the **Attached Entity Profile** field, select the AAEP profile to attach to the policy group.
f) Click **Submit**.

Configuring Breakout Ports With Profiles and Selectors Using the GUI

This procedure uses port profiles and selectors to configure breakout ports. You configure a breakout leaf port with an leaf interface profile, associate the profile with a switch, and configure the sub-ports.

Before you begin

- The Cisco Application Centric Infrastructure (ACI) fabric is installed, the Cisco Application Policy Infrastructure Controllers (APICs) are online, and the Cisco APIC cluster is formed and healthy.
- A Cisco APIC fabric administrator account is available that can configure the breakout ports.
- The target leaf switches are registered in the Cisco ACI fabric and available.
- The 40GE or 100GE leaf switch ports are connected with Cisco breakout cables to the downlink ports.

Procedure

Step 1 On the menu bar, choose **Fabric** > **Access Policies**.
Step 2 In the Navigation pane, expand **Interfaces** and **Leaf Interfaces** and **Profiles**.
Step 3 Right-click **Profiles** and choose **Create Leaf Interface Profile**.
Step 4 Type the name and optional description, click the + symbol on **Interface Selectors**
Step 5 Perform the following:

a) Type a name (and optional description) for the **Access Port Selector**.
b) In the **Interface IDs** field, type the slot and port for the breakout port.
c) In the **Interface Policy Group** field, click the down arrow and choose **Create Leaf Breakout Port Group**.

d) Type the name (and optional description) for the **Leaf Breakout Port Group**.
e) In the **Breakout Map** field, choose **10g-4x** or **25g-4x**.

For a list of the switches that support breakout, see Configuration of Breakout Ports, on page 1.

f) Click **Submit**.

Step 6 To assign a Breakout Port to an EPG, perform the following steps:

On the menu bar, choose **Tenant** > **Application Profiles** > **Application EPG**. Right-click on **Application EPG** to open **Create Application EPG**dialog box, and perform the following steps:

a) Select the **Statically Link with Leaves/Paths** check box to gain access to the **Leaves/Paths** tab in the dialog box.
b) Complete one of the following sets of steps:

Option	Description
If you want to deploy the EPG on...	Then
A node	**a.** Expand the **Leaves** area. **b.** From the **Node** drop-down list, choose a node. **c.** In the **Encap** field, enter the appropriate VLAN. **d.** (Optional) From the **Deployment Immediacy** drop-down list, accept the default **On Demand** or choose **Immediate**. **e.** (Optional) From the Mode drop-down list, accept the default **Trunk** or choose another mode.
A port on the node	**a.** Expand the **Paths** area. **b.** From the **Path** drop-down list, choose the appropriate node and port. **c.** (Optional) In the **Deployment Immediacy** field drop-down list, accept the default **On Demand** or choose **Immediate**. **d.** (Optional) From the Mode drop-down list, accept the default **Trunk** or choose another mode. **e.** In the **Port Encap** field, enter the secondary VLAN to be deployed. **f.** (Optional) In the **Primary Encap** field, enter the primary VLAN to be deployed.

Step 7 To associate the Leaf Interface Profile to a the leaf switch, perform the following steps:

a) Expand **Switches** and **Leaf Switches**, and **Profiles**.
b) Right-click **Profiles** and select **Create Leaf Profiles**.
c) Type the name and optional description of the Leaf Profile.
d) Click the + symbol on the **Leaf Selectors** area.
e) Type the leaf selector name and an optional description.
f) Click the down arrow on the **Blocks** field and choose the switch to be associated with the breakout leaf interface profile.
g) Click the down arrow on the **Policy Group** field and choose **Create Access Switch Policy Group**.
h) Type a name and optional description for the Access Switch Policy Group.

i) Optional. Enable other policies.
j) Click **Submit**.
k) Click **Update**.
l) Click **Next**.
m) In the **Associations Interface Selector Profiles** area, choose the Interface Selector Profile you previously created for the breakout port.
n) Click **Finish**.

Step 8 To verify the breakout port has been split into four sub ports, perform the following steps:

a) On the Menu bar, click **Fabric** > **Inventory**.
b) On the Navigation bar, click the Pod and Leaf where the breakout port is located.
c) Expand **Interfaces** and **Physical Interfaces**.
You should see four ports at the position where the breakout port was configured. For example, if you configured 1/10 as a breakout port, you should see the following:

- **eth1/10/1**
- **eth1/10/2**
- **eth1/10/3**
- **eth1/10/4**

Step 9 To configure the sub ports, perform the following steps:

a) On the Menu bar, click **Fabric** > **Access Policies**.
b) On the Navigation bar, expand **Interfaces**, **Leaf Interfaces**, **Profiles**, and the breakout leaf interface profile you previously created.

You will see a port selector with the breakout cable. Instead of defining a sub port block under already existing port selector, you need to define on a new access port selector.

c) On the Navigtion bar, right click the higher level interface profile and select **Create Access Port Selector**.
d) In the **Name** field, enter the sub port name.
e) In the **Interface IDs** field, enter the IDs for the four sub ports in a format such as `1/10/1-4`.
f) In the **Interface Policy Group** field, select **Create Leaf Access Port Policy Group**.
g) Click **Submit**.

Step 10 To apply the Policy Group to an individual interface that links the AAEP to the port, perform the following steps:

a) In the **Name** field, enter the name for the Leaf Access Port Group Policy.
b) In the **Link Level Policy** field, select **link-level_auto**.
c) In the **CDP Policy** field, select **cdp_enabled**.
d) In the **LLDP Policy** field, select **default**.
e) In the **Attached Entity Profile** field, select the AAEP profile to attach to the policy group.
f) Click **Submit**.

Configuring Breakout Ports With the Interface Configuration Using the GUI

Beginning with the Cisco Application Policy Infrastructure Controller (APIC) 6.0(1) release, you can use the **Fabric** > **Access Policies** > **Interface Configuration** workflow to configure breakout ports.

Before you begin

- The Cisco Application Centric Infrastructure (ACI) fabric is installed, the Cisco APICs are online, and the Cisco APIC cluster is formed and healthy.
- A Cisco APIC fabric administrator account is available that can configure the breakout ports.
- The target leaf switches are registered in the Cisco ACI fabric and available.
- The 40GE or 100GE leaf switch ports are connected with Cisco breakout cables to the downlink ports.

Procedure

Step 1 On the menu bar, choose **Fabric** > **Access Policies**.

Step 2 In the Navigation pane, choose **Interface Configuration**.

Step 3 In the Work pane, choose **Actions** > **Breakout**.

Step 4 In the **Breakout** page, perform the following substeps:

a) For **Node**, click **Select Node**, put a check in the box for the desired switch (node), then click **OK**. You can select multiple switches.
b) For **Interfaces For All Switches**, enter the range of desired interfaces.
c) For **Breakout Map**, choose the desired breakout type.
d) Click **Save**.

Configuring Dynamic Breakout Ports Using the NX-OS Style CLI

Use the following steps to configure a breakout port, verify the configuration, and configure an EPG on a sub port, using the NX-OS style CLI.

Before you begin

- The ACI fabric is installed, APIC controllers are online, and the APIC cluster is formed and healthy.
- An APIC fabric administrator account is available that will enable creating the necessary fabric infrastructure configurations.
- The target leaf switches are registered in the ACI fabric and available.
- The 40GE or 100GE leaf switch ports are connected with Cisco breakout cables to the downlink ports.

SUMMARY STEPS

1. **configure**
2. **leaf** *ID*
3. **interface ethernet***slot/port*
4. **breakout10g-4x** | **25g-4x**
5. **show run**
6. **tenant** *tenant-name*
7. **vrf context** *vrf-name*
8. **bridge-domain** *bridge-domain-name*
9. **vrf member** *vrf-name*
10. **application** *application-profile-name*
11. **epg** *epg-name*
12. **bridge-domain member** *bridge-domain-name*
13. **leaf** *leaf-name*
14. **speed** *interface-speed*
15. **show run**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**configure** **Example:** `apic1# configure`	Enters configuration mode.
Step 2	**leaf** *ID* **Example:** `apic1(config)# leaf 101`	Selects the leaf switch where the breakout port will be located and enters leaf configuration mode.
Step 3	**interface ethernet***slot/port* **Example:** `apic1(config-leaf)# interface ethernet 1/16`	Identifies the interface to be enabled as a 40 Gigabit Ethernet (GE) breakout port.
Step 4	**breakout10g-4x**	**25g-4x** **Example:** `apic1(config-leaf-if)# breakout 10g-4x`	Enables the selected interface for breakout. **Note** For switch support for the Dynamic Breakout Port feature, see Configuration of Breakout Ports, on page 1.

Step 5	**show run**
	Example:
	`apic1(config-leaf-if)# show run`
	`# Command: show running-config leaf 101 interface ethernet 1 / 16`
	`# Time: Fri Dec 2 18:13:39 2016`
		Verifies the configuration by showing the running configuration of the interface and returns to global configuration mode.

	Command or Action	Purpose

	`leaf 101`
	`interface ethernet 1/16`
	`breakout 10g-4x`
	`apic1(config-leaf-if)# exit`
	`apic1(config-leaf)# exit`	

Step 6	**tenant** *tenant-name* **Example:** `apic1(config)# tenant tenant64`	Selects or creates the tenant that will consume the breakout ports and enters tenant configuration mode.

Step 7	**vrf context** *vrf-name*
	Example:
	`apic1(config-tenant)# vrf context vrf64`
	`apic1(config-tenant-vrf)# exit`
		Creates or identifies the Virtual Routing and Forwarding (VRF) instance associated with the tenant and exits the configuration mode.

Step 8	**bridge-domain** *bridge-domain-name* **Example:** `apic1(config-tenant)# bridge-domain bd64`	Creates or identifies the bridge-domain associated with the tenant and enters BD configuration mode.

Step 9	**vrf member** *vrf-name*
	Example:
	`apic1(config-tenant-bd)# vrf member vrf64`
	`apic1(config-tenant-bd)# exit`
		Associates the VRF with the bridge-domain and exits the configuration mode.

Step 10	**application** *application-profile-name* **Example:** `apic1(config-tenant)# application app64`	Creates or identifies the application profile associated with the tenant and the EPG.
Step 11	**epg** *epg-name* **Example:** `apic1(config-tenant)# epg epg64`	Creates or identifies the EPG and enters into EPG configuration mode.

Step 12	**bridge-domain member** *bridge-domain-name*
	Example:
	`apic1(config-tenant-app-epg)# bridge-domain member bd64`
	`apic1(config-tenant-app-epg)# exit`
	`apic1(config-tenant-app)# exit`
	`apic1(config-tenant)# exit`
		Associates the EPG with the bridge domain and returns to global configuration mode.
		Configure the sub ports as desired, for example, use the speed command in leaf interface mode to configure a sub port.

Step 13	**leaf** *leaf-name*
	Example:
	`apic1(config)# leaf 1017`
	`apic1(config-leaf)# interface ethernet 1/13`
	`apic1(config-leaf-if)# vlan-domain member dom1`
	`apic1(config-leaf-if)# switchport trunk allowed vlan 20 tenant t1 application AP1 epg EPG1`
	Note
		Associates the EPG with a break-out port.

	Command or Action	Purpose
	The vlan-domain and vlan-domain member commands mentioned in the above example are a pre-requisite for deploying an EPG on a port.	

Step 14	**speed** *interface-speed*
	Example:
	`apic1(config)# leaf 101`
	`apic1(config-leaf)# interface ethernet 1/16/1`
	`apic1(config-leaf-if)# speed 10G`
	`apic1(config-leaf-if)# exit`
		Enters leaf interface mode, sets the speed of an interface, and exits the configuration mode.

Step 15	**show run** **Example:** `apic1(config-leaf)# show run`	After you have configured the sub ports, entering this command in leaf configuration mode displays the sub port details.

The port on leaf 101 at interface 1/16 is confirmed enabled for breakout with sub ports 1/16/1, 1/16/2, 1/16/3, and 1/16/4.

Example

This example configures the port for breakout:

```
apic1# configure
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/16
apic1(config-leaf-if)# breakout 10g-4x
```

This example configures the EPG for the sub ports.

```
apic1(config)# tenant tenant64
apic1(config-tenant)# vrf context vrf64
apic1(config-tenant-vrf)# exit
apic1(config-tenant)# bridge-domain bd64
apic1(config-tenant-bd)# vrf member vrf64
apic1(config-tenant-bd)# exit
apic1(config-tenant)# application app64
apic1(config-tenant-app)# epg epg64
apic1(config-tenant-app-epg)# bridge-domain member bd64
apic1(config-tenant-app-epg)# end
```

This example sets the speed for the breakout sub ports to 10G.

```
apic1(config)# leaf 101
apic1(config-leaf)# interface ethernet 1/16/1
apic1(config-leaf-if)# speed 10G
apic1(config-leaf-if)# exit

apic1(config-leaf)# interface ethernet 1/16/2
apic1(config-leaf-if)# speed 10G
apic1(config-leaf-if)# exit
apic1(config-leaf)# interface ethernet 1/16/3
apic1(config-leaf-if)# speed 10G
apic1(config-leaf-if)# exit
apic1(config-leaf)# interface ethernet 1/16/4
apic1(config-leaf-if)# speed 10G
apic1(config-leaf-if)# exit
```

This example shows the four sub ports connected to leaf 101, interface 1/16.

```
apic1#(config-leaf)# show run
# Command: show running-config leaf 101
# Time: Fri Dec  2 00:51:08 2016
  leaf 101
    interface ethernet 1/16/1
      speed 10G
      negotiate auto
      link debounce time 100
      exit
    interface ethernet 1/16/2
      speed 10G
      negotiate auto
      link debounce time 100
      exit
    interface ethernet 1/16/3
      speed 10G
      negotiate auto
      link debounce time 100
      exit
    interface ethernet 1/16/4
      speed 10G
      negotiate auto
      link debounce time 100
      exit
    interface ethernet 1/16
      breakout 10g-4x
      exit
    interface vfc 1/16
```