



Dynamic Breakout Ports

This chapter contains the following sections:

- [Configuration of Dynamic Breakout Ports, on page 1](#)
- [Configuring Dynamic Breakout Ports Using the APIC GUI, on page 4](#)
- [Configuring Dynamic Breakout Ports Using the NX-OS Style CLI, on page 7](#)
- [Configuring Dynamic Breakout Ports Using the REST API, on page 10](#)

Configuration of Dynamic 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 or a 100Gb port to be split into four independent and logical 25Gb ports.

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

- Cisco QSFP-4SFP10G
- Cisco QSFP-4SFP25G
- Cisco QSFP-4X10G-AOC
- 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



Note For the supported optics and cables, see the *Cisco Optics-to-Device Compatibility Matrix*:
<https://tmgmatrix.cisco.com/>

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 (beginning in the 5.1(3) release)
- N9K-C93180YC-FX3 (beginning in the 5.1(3) release)
- N9K-C93600CD-GX (beginning in the 5.1(3) release)
- N9K-C9364C-GX (beginning in the 5.1(3) release)

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

- N9K-C93180LC-EX
- N9K-C9336C-FX2
- N9K-C93180YC-FX
- N9K-C93360YC-FX2
- N9K-C93216TC-FX2
- N9K-C93108TC-FX3P (beginning in the 5.1(3) release)
- N9K-C93180YC-FX3 (beginning in the 5.1(3) release)
- N9K-C93600CD-GX (beginning in the 5.1(3) release)
- N9K-C9364C-GX (beginning in the 5.1(3) release)

Observe the following guidelines and limitations:

- Breakout ports are supported only on downlinks and converted downlinks.
- 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 (beginning in the 5.1(3) release)

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

- Cisco N9K-93600CD-GX (beginning in the 5.1(3) release)

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.

- 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.
- 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.
- 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.
- 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**).
- 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 max number of downlinks 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.
- 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 downlinks 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.

Configuring Dynamic Breakout Ports Using the APIC GUI

Configure a Breakout Leaf Port with an Leaf Interface Profile, associate the profile with a switch, and configure the sub ports with the following steps.



Note You can also configure ports for breakout in the APIC GUI by navigating to **Fabric > Inventory**, and clicking **Topology** or **Pod**, or expanding **Pod** and clicking **Leaf**. Then, enable configuration and click the **Interface** tab.

Procedure

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.

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 switches supporting breakout, see [Configuration of Dynamic 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	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 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 Navigation 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 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.

Procedure

	Command or Action	Purpose
Step 1	configure Example: <code>apic1# configure</code>	Enters configuration mode.
Step 2	leaf ID Example: <code>apic1(config)# leaf 101</code>	Selects the leaf switch where the breakout port will be located and enters leaf configuration mode.
Step 3	interface ethernetslot/port Example: <code>apic1(config-leaf)# interface ethernet 1/16</code>	Identifies the interface to be enabled as a 40 Gigabit Ethernet (GE) breakout port.
Step 4	breakout10g-4x 25g-4x Example: <code>apic1(config-leaf-if)# breakout 10g-4x</code>	Enables the selected interface for breakout. Note For switch support for the Dynamic Breakout Port feature, see Configuration of Dynamic Breakout Ports, on page 1 .
Step 5	show run Example: <code>apic1(config-leaf-if)# show run</code> # Command: show running-config leaf 101 interface ethernet 1 / 16 # Time: Fri Dec 2 18:13:39 2016 leaf 101	Verifies the configuration by showing the running configuration of the interface and returns to global configuration mode.

	Command or Action	Purpose
	<pre> interface ethernet 1/16 breakout 10g-4x apicl(config-leaf-if) # exit apicl(config-leaf) # exit </pre>	
Step 6	tenant <i>tenant-name</i> Example: <pre> apicl(config) # tenant tenant64 </pre>	Selects or creates the tenant that will consume the breakout ports and enters tenant configuration mode.
Step 7	vrf context <i>vrf-name</i> Example: <pre> apicl(config-tenant) # vrf context vrf64 apicl(config-tenant-vrf) # exit </pre>	Creates or identifies the Virtual Routing and Forwarding (VRF) instance associated with the tenant and exits the configuration mode.
Step 8	bridge-domain <i>bridge-domain-name</i> Example: <pre> apicl(config-tenant) # bridge-domain bd64 </pre>	Creates or identifies the bridge-domain associated with the tenant and enters BD configuration mode.
Step 9	vrf member <i>vrf-name</i> Example: <pre> apicl(config-tenant-bd) # vrf member vrf64 apicl(config-tenant-bd) # exit </pre>	Associates the VRF with the bridge-domain and exits the configuration mode.
Step 10	application <i>application-profile-name</i> Example: <pre> apicl(config-tenant) # application app64 </pre>	Creates or identifies the application profile associated with the tenant and the EPG.
Step 11	epg <i>epg-name</i> Example: <pre> apicl(config-tenant) # epg epg64 </pre>	Creates or identifies the EPG and enters into EPG configuration mode.
Step 12	bridge-domain member <i>bridge-domain-name</i> Example: <pre> apicl(config-tenant-app-epg) # bridge-domain member bd64 apicl(config-tenant-app-epg) # exit apicl(config-tenant-app) # exit apicl(config-tenant) # exit </pre>	<p>Associates the EPG with the bridge domain and returns to global configuration mode.</p> <p>Configure the sub ports as desired, for example, use the speed command in leaf interface mode to configure a sub port.</p>
Step 13	leaf <i>leaf-name</i> Example: <pre> apicl(config) # leaf 1017 apicl(config-leaf) # interface ethernet 1/13 apicl(config-leaf-if) # vlan-domain member dom1 apicl(config-leaf-if) # switchport trunk </pre>	Associates the EPG with a break-out port.

	Command or Action	Purpose
	<pre>allowed vlan 20 tenant t1 application AP1 epg EPG1</pre> <p>Note The vlan-domain and vlan-domain member commands mentioned in the above example are a pre-requisite for deploying an EPG on a port.</p>	
Step 14	<p>speed <i>interface-speed</i></p> <p>Example:</p> <pre>apic1(config)# leaf 101 apic1(config-leaf)# interface ethernet 1/16/1 apic1(config-leaf-if)# speed 10G apic1(config-leaf-if)# exit</pre>	Enters leaf interface mode, sets the speed of an interface, and exits the configuration mode.
Step 15	<p>show run</p> <p>Example:</p> <pre>apic1(config-leaf)# show run</pre>	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

```

Configuring Dynamic Breakout Ports Using the REST API

Configure a Breakout Leaf Port with an Leaf Interface Profile, associate the profile with a switch, and configure the sub ports with the following steps.

For switch support for the breakout feature, see [Configuration of Dynamic Breakout Ports, on page 1](#).

Procedure

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.

Procedure

Step 1 Configure a breakout policy group for the breakout port with JSON, such as the following example:

Example:

In this example, we create an interface profile 'brkout44' with the only port 44 underneath its port selector. The port selector is pointed to a breakout policy group 'new-brkoutPol'.

```
{
  "infraAccPortP": {
    "attributes": {
      "dn": "uni/infra/accportprof-brkout44",
      "name": "brkout44",
      "rn": "accportprof-brkout44",
      "status": "created,modified"
    },
    "children": [ {
      "infraHPortS": {
        "attributes": {
          "dn": "uni/infra/accportprof-brkout44/hports-new-brkoutPol-typ-range",
          "name": "new-brkoutPol",
          "rn": "hports-new-brkoutPol-typ-range",
          "status": "created,modified"
        },
        "children": [ {
          "infraPortBlk": {
            "attributes": {
              "dn": "uni/infra/accportprof-brkout44/hports-new-brkoutPol-typ-range/portblk-block2",
              "fromPort": "44",
              "toPort": "44",
              "name": "block2",
              "rn": "portblk-block2",
              "status": "created,modified"
            },
            "children": []
          }, {
            "infraRsAccBaseGrp": {
              "attributes": {
                "tDn": "uni/infra/funcprof/brkoutportgrp-new-brkoutPol",
                "status": "created,modified"
              },
              "children": []
            }
          }
        ]
      }
    }
  ]
}
```

Step 2 Create a new switch profile and associate it with the port profile, previously created, with JSON such as the following example:

Example:

In this example, we create a new switch profile 'leaf1017' with switch 1017 as the only node. We associate this new switch profile with the port profile 'brkout44' created above. After this, the port 44 on switch 1017 will have 4 sub ports.

Example:

```
{
  "infraNodeP": {
    "attributes": {
      "dn": "uni/infra/nprof-leaf1017",
      "name": "leaf1017", "rn": "nprof-leaf1017",
      "status": "created,modified"
    },
    "children": [ {
      "infraLeafS": {
        "attributes": {
          "dn": "uni/infra/nprof-leaf1017/leaves-1017-typ-range",
          "type": "range",
          "name": "1017",
          "rn": "leaves-1017-typ-range",
          "status": "created"
        },
        "children": [ {
          "infraNodeBlk": {
            "attributes": {
              "dn": "uni/infra/nprof-leaf1017/leaves-1017-typ-range/nodeblk-102bf7dc60e63f7e",
              "from_": "1017", "to_": "1017",
              "name": "102bf7dc60e63f7e",
              "rn": "nodeblk-102bf7dc60e63f7e",
              "status": "created"
            },
            "children": [ ] }
          }
        ]
      }
    }, {
      "infraRsAccPortP": {
        "attributes": {
          "tDn": "uni/infra/accportprof-brkout44",
          "status": "created,modified"
        },
        "children": [ ] }
      }
    ]
  }
}
```

Step 3 Configure the subports.**Example:**

This example configures subports 1/44/1, 1/44/2, 1/44/3, 1/44/4 on switch 1017, for instance, in the example below, we configure interface 1/44/3. It also creates the `infraSubPortBlk` object instead of the `infraPortBlk` object.

```
{
  "infraAccPortP": {
    "attributes": {
      "dn": "uni/infra/accportprof-brkout44",
      "name": "brkouttest1",
      "rn": "accportprof-brkout44",
      "status": "created,modified"
    },
    "children": [ {
      "infraHPortS": {
        "attributes": {
          "dn": "uni/infra/accportprof-brkout44/hports-sell-typ-range",
          "name": "sell",

```

Step 4

Deploy an EPG on a specific port.

```
<fvTenant name="<tenant_name>" dn="uni/tn-test1" >
  <fvCtx name="<network_name>" pcEnfPref="enforced" knwMcastAct="permit"/>
  <fvBD name="<bridge_domain_name>" unkMcastAct="flood" >
    <fvRsCtx tnFvCtxName="<network_name>"/>
  </fvBD>
  <fvAp name="<application_profile>" >
    <fvAEPg name="<epg_name>" >
      <fvRsPathAtt tDn="topology/pod-1/paths-1017/pathep-[eth1/13]" mode="regular"
instrImedcy="immediate" encap="vlan-20"/>
    </fvAEPg>
  </fvAp>
</fvTenant>
```

