



Overview

- [Licensing Requirements, on page 1](#)
- [Supported Platforms, on page 1](#)
- [About Interfaces, on page 1](#)
- [Virtual Device Contexts, on page 12](#)
- [High Availability for Interfaces, on page 12](#)

Licensing Requirements

For a complete explanation of Cisco NX-OS licensing recommendations and how to obtain and apply licenses, see the [Cisco NX-OS Licensing Guide](#) and the [Cisco NX-OS Licensing Options Guide](#).

Supported Platforms

Starting with Cisco NX-OS release 7.0(3)I7(1), use the [Nexus Switch Platform Support Matrix](#) to know from which Cisco NX-OS releases various Cisco Nexus 9000 and 3000 switches support a selected feature.

About Interfaces

Cisco NX-OS supports multiple configuration parameters for each of the interface types supported. Most of these parameters are covered in this guide but some are described in other documents.

The following table shows where to get further information on the parameters you can configure for an interface.

Table 1: Interface Parameters

Feature	Parameters	Further Information
Basic parameters	Description, duplex, error disable, flow control, MTU, beacon	“Configuring Basic Interface Parameters”
Layer 3	Medium, IPv4 and IPv6 addresses	“Configuring Layer 3 Interfaces”

Feature	Parameters	Further Information
Layer 3	Bandwidth, delay, IP routing, VRFs	<i>Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide</i> <i>Cisco Nexus 9000 Series NX-OS Multicast Routing Configuration Guide</i>
Port Channels	Channel group, LACP	“Configuring Port Channels”
Security	EOU	<i>Cisco Nexus 9000 Series NX-OS Security Configuration Guide</i>

Ethernet Interfaces

- Ethernet interfaces include routed ports.
- For N9K-C9316D-GX, Ports 1-16 supports 400G/100G/40G, and 10G with QSA.

Cisco Nexus N9K-C9364C-GX and N9K-C93600CD-GX Guidelines and Limitations

Cisco Nexus N9K-C9364C-GX and N9K-C93600CD-GX switches have the following guidelines and limitations:

- Consecutive groups of four interfaces (1-4, 5-8, 9-12, and so on, are referred to as a quad group). Attempting to use a mix of link speeds within a quad group is not supported. This applies to ports 1-24 of the N9K-C93600CD-GX and all ports of the N9K-C9364C-GX.
- Only one speed will be up in a quad group at a time. The first link up in a quad group determines the speed of the quad group. Ports with any other speed will be down with the reason of "Link not connected".
- The speed which remains functional when mixing within a quad group is not stored. When a mismatched speed transceiver is inserted into a quad group and brought up, all ports in the quad group will be reset. The first link which comes up after the reset determines the speed of the quad group. It is possible that pre-existing links may be shut down. You can remove the mismatched speed transceiver to recover from this state.
- FC-FEC is not supported on the second lane of the 50Gx2 breakout port. The second breakout port will not come up when 50Gx2 breakout is configured, ensure to configure RS-FEC with 50Gx2 breakout.

For Cisco NX-OS Release 10.1(1), Port-group Support - GX provides the following support:

- The following points are applicable for N9K-C93600CD-GX:
 - For ports 1-24, every four ports (1-4, 5-8, 9-12, and so on, referred to as a "quad") operate at the same speed.
 - All the ports in a quad operate in 10G with QSA, or 40G or 100G.
 - Mixed speed is not supported within the same quad.
 - With QSA, all ports in a quad can operate at 10G speed.

- The speed of the quad is determined by the first Transceiver Type plugged in that quad and not by what Transceiver is plugged in first port of the quad.

For example, if port 14 (belonging to quad containing ports 13-16) is a 10G Transceiver with QSA plugged in as the first in the quad, then all other ports in the quad should be of speed 10G.

Similarly if Port 23 (belonging to quad 21-24) if plugged in with 40G Transceiver as the first port in the quad then all other ports should be plugged in with 40G Transceiver.

- If a 100G Transceiver is plugged in port 24 while port 21 is already plugged-in with 40G Transceivers, then the 100G Interface will be in “XCVR speed mismatch” state and link will not come up.
- In order for Port 24 with 100G Transceiver to be able to link up then all other non 100G transceivers in that quad needs to be removed and port 24 needs to be flapped.
- This is same with the case when QSA+10G is inserted in a port where there is a non QSA+10G Transceiver plugged already in that quad.
- When a Mismatch Transceiver is plugged into a quad, syslog will be generated as “Interface Ethernet1/X is down (Reason: Inserted transceiver speed mismatch with quad speed Y).
- If Port 4 is plugged in with a 100G Transceiver followed by a 40G Transceiver in Port 1, then 100G will bring up the link but 40G transceiver will not link up and be in “XCVR mismatch speed”. If this configuration is copied as startup configuration and switch is reloaded, after the switch comes up, 100G transceiver will be up and 40G Transceiver will be in “XCVR mismatch speed” state.
- The port number will not decide the speed of the quad, only the first plugged-in Transceiver will determine the speed even when all interfaces are in “Admin shut” state in the quad. Same state will be preserved when “copy running-config startup-config” is performed and the switch is reloaded. But when the switch is reloaded with “Reload ascii” option, after the switch comes up, first port (with a plugged in transceiver) in the quad will determine the port group. Other mismatching transceivers will be in “XCVR mismatch Speed”.
- If 40G transceiver needs to be up then 100G Transceiver needs to be removed. All other ports in that quad can be empty or plugged in only with 40G transceiver for all the transceivers to link up in that particular quad.
- Only exception to this is when a quad has 40G Transceivers and is plugged in with a Dual Rate Transceiver (40G/100G Capable) configured with 40G Speed then dual rate transceiver with 40G Speed will link up. The port number will not decide the speed of the quad, only the first plugged-in Transceiver will determine the speed of the quad even when all interfaces are in “Admin shut” state in the quad. Same state will be preserved when “copy running-config startup-config” is performed and switch is reloaded but when the switch is reloaded with “Reload ascii” option, after the switch comes up, first port (with a plugged in transceiver) in the quad will determine the speed of the quad, other mismatching transceivers will be in “XCVR mismatch Speed”.
- It is highly recommended to use only Transceivers of same speed in a quad to avoid any disruption or indeterministic state. Same logic is extended for quad ports 25-26 and other quad ports 27-28.
- For N9K-C9316D-GX: Ports 1-16 support 400G/100G/40G and 10G with QSA and there are no port group limitations.
- Beginning with Cisco Nexus NX-OS Release 10.1(2) Auto negotiation is supported for Speed 40G and 100G on NX-OS N9K-C93600CD-GX, N9K-C9316D-GX and N9K-C9364C-GX .
- Mixed speed is not supported within the same quad.

- With QSA all ports in a quad can operate at 10G speed.

Cisco Nexus N9K-X9400-16W Guidelines and Limitations

Beginning with Cisco Nexus NX-OS Release 10.5(1), the following guidelines and limitations are applicable:

Access Ports

An access port carries traffic for one VLAN. This type of port is a Layer 2 interface only.

For more information on access ports, see the “Information About Access and Trunk Interfaces” section.

Routed Ports

A routed port is a physical port that can route IP traffic to another device. A routed port is a Layer 3 interface only.

For more information on routed ports, see the “Routed Interfaces” section.

Management Interface

You can use the management Ethernet interface to connect the device to a network for remote management using a Telnet client, the Simple Network Management Protocol (SNMP), or other management agents. The management port (mgmt0) is autosensing and operates in full-duplex mode at a speed of 10/100/1000 Mb/s.

For more information on the management interface, see the [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide](#).

Port-Channel Interfaces

A port channel is a logical interface that is an aggregation of multiple physical interfaces. You can bundle up to 32 individual links (physical ports) into a port channel to improve bandwidth and redundancy. For more information about port-channel interfaces, see the “Configuring Port Channels” section.

Subinterfaces

You can create virtual subinterfaces using a parent interface configured as a Layer 3 interface. A parent interface can be either a physical port or a port-channel. A parent interface can be a physical port. Subinterfaces divide the parent interface into two or more virtual interfaces on which you can assign unique Layer 3 parameters such as IP addresses and dynamic routing protocols.

Loopback Interfaces

A loopback interface is a virtual interface with a single endpoint that is always up. Any packet that is transmitted over a virtual loopback interface is immediately received by that interface. Loopback interfaces emulate a physical interface. For more information about subinterfaces, see the “Loopback Interfaces” section.

Breakout Interfaces

Cisco NX-OS supports the breakout of a high bandwidth interface into one or more low bandwidth interfaces at the module level or at the per-port level.

Module Level Breakout

Module Level Breakout allows certain high-density ports on a switch module to be split into multiple lower-bandwidth ports, providing increased flexibility and customization in network configurations.

You can configure the **interface breakout** command to split a high bandwidth interface of a module into multiple lower speed ports.

Some modules breakdowns the all the ports to 4x10G. 4x25G, 4x50G, 4x100G, 2x50G, 2x100G.

For example, a module level breakout 4X10G implies that high bandwidth 40G interface is broken down into four 10G interfaces. The module is reloaded and the configuration for the interface is removed when the command is executed.

The following is an example of the command:

```
switch# configure terminal
switch(config)# interface breakout module 1
Module will be reloaded. Are you sure you want to continue(yes/no)? yes
```

The **no interface breakout module *module_number*** command undoes the breakout configuration. In the above example of 4x10G interface, all interfaces of the module in 40G mode and deletes the configuration for the previous 10G interfaces.

About the Lane Selector

The lane selector is a push button switch and 4 LEDs located on the Cisco Nexus switch (left side of front panel, labeled "LS"). The push button switch and LEDs are used to indicate the status of the ports. The lane selector is supported on Cisco Nexus Series 9000 series switches and the Cisco Nexus 3164 and 3232 switches.

By default, the LEDs indicate the link/activity status of a 1 x 40G configuration. When the ports are configured as 4 x 10G, you can access the link status of each individual 10G port with the lane selector.

By pressing the lane selector push button, the port LED shows the selected lane's link/activity status. The 1st time the push button is pressed, the first LED displays the status of the first port. Pressing the push button a 2nd time displays the status of the second port, and so on. You can display the status of each of the four ports by pressing the push button in this manner.

For example, if port 60 is configured as 4 x 10G, pressing the lane selector push button once displays the link status of 60/1/1. Pressing the push button a second time displays the link status of 60/1/2.

When you press the push button after displaying the status of the last port, all four of the LEDs should extinguish to indicate that the lane selector has returned to display the status for the default 1 x 40G configuration.



Note A 10G breakout port's LED blinks when the beacon feature has been configured for it.



Note When a port is configured to be in 10G breakout mode and no lane is selected, the 40G port's LED illuminates as green even though only one of the 10G breakout ports is up.

Support for Breakout Interfaces

The table provides detailed information of the supported or not supported breakout modes. For more information, see [Cisco Nexus Data Sheets](#).

Table 2: Breakout Modes Support Matrix

Switches	4x10G	4x25G	2x50G	2x100G	2x200G	2x400G	4x50G	4x100G	8x100G
Nexus 9300-FX3 Platform Switches	Yes	Yes	Yes	No	No	No	No	No	No
N9K-C93108TC-FX3P									
N9K-C93180YC-FX3									
N9K-X9636C-RX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-X9636C-R	Yes	Yes	Yes	No	No	No	No	No	No
N9K-X9636Q-R	Yes	No	No	No	No	No	No	No	No
N9K-X96136YC-R	No	No	No	No	No	No	No	No	No
N3K-C3636C-R	Yes	Yes	Yes	No	No	No	No	No	No
N3K-C36180YC-R	Yes	Yes	Yes	No	No	No	No	No	No
N9K-93108TC-FX3P	Yes	Yes	Yes	No	No	No	No	No	No
N9K-93108TC-EX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-93180YC-EX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-93108TC-FX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-93180YC-FX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-9348GC-FXP	Yes	Yes	Yes	No	No	No	No	No	No
N9K-X9732C-EX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-X9736C-EX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-X9732C-EXM	Yes	Yes	Yes	No	No	No	No	No	No
N9K-X9736C-FX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-X9736Q-FX	Yes	No	No	No	No	No	No	No	No
N9K-X9788TC-FX	Yes	Yes	Yes	No	No	No	No	No	No

Switches	4x10G	4x25G	2x50G	2x100G	2x200G	2x400G	4x50G	4x100G	8x100G
N9K-X9732C-FX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-C9348GC-FXP	Yes	Yes	Yes	No	No	No	No	No	No
N9K-C9336C-FX2	Yes	Yes	Yes	No	No	No	No	No	No
N9K-C93216TC-FX2	Yes	Yes	Yes	No	No	No	No	No	No
N9K-C93360YC-FX2	Yes	Yes	Yes	No	No	No	No	No	No
N9K-C9364C-GX	Yes	Yes	Yes	No	No	No	No	No	No
N9K-C9316D-GX	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No
N9K-C93600CD-GX	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No
N9K-X9716D-GX	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No
N9K-C9332D-GX2B	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No
N9K-C9348D-GX2A	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No
N9K-X9400-16W	Yes	Yes	Yes	Yes	No	No	Yes	No	No
N9K-X9400-8D	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No
N9K-X98900CD-A	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No

Guidelines and Limitations for Breakout

- Cisco Nexus 9516 switch does not support breakout on Modules 8 to 16.
- Beginning with Cisco NX-OS Release 7.0(3)F2(1), the 36-port 100-Gigabit Ethernet QSFP28 line cards (N9K-X9636C-R) and 36-port 40-Gigabit Ethernet QSFP+ line cards (N9K-X9636Q-R) support breakout of 4x10-Gigabit.
- Beginning with Cisco NX-OS Release 9.2(1), N9K-9636C-R, N9K-X9636Q-R and N9K-X9636C-RX line cards support break out of 40G ports into 4x10 Gigabits.
- Beginning with Cisco NX-OS Release 9.2(2), N9K-X9636C-R and N9K-X9636C-RX line cards support break out of 100G ports into 4x25 Gigabits. The N9K-C9636C-R does not support RS-FEC.
Beginning with Cisco NX-OS Release 9.3(3), the default FEC mode on N9K-X9636C-R and N9K-X9636C-RX is FC-FEC for 25Gx4 and 50Gx2. When you connect N9K-X9636C-RX to N9K-X9636C-R, you must configure FC-FEC on N9K-X9636C-RX, because RS-FEC is not supported.
The N9K-X96136YC-R line card does not support breakout.
- Beginning with Cisco NX-OS Release 9.3(3), these switches support breakout.
 - Cisco Nexus 93600CD-GX switch
 - Cisco Nexus 9500 R-Series switches support break out of 100G ports into 2x50 Gigabits.

On Nexus 9500 R-Series switches with N9K-X9636C-R and N9K-X9636C-RX line cards, limited optics (QSFP-100G-PSM4-S, QSFP-100G-AOC, QSFP-100G-CU1M -CU3M) support 2x50G and 4x25G breakout. For more information see *Cisco Optics-to-Device Compatibility Matrix*.

Caveats

- As of Cisco NX-OS Release 7.0(3)I7(2), manual breakout of QSA ports is not supported.

Manual breakout is supported on the following platforms because auto-breakout does not happen successfully on them—N9K-C93128TX, N9K-9332, N9K-C9396PX, N9K-C9396TX, N9K-C9372PX, N9K-C9372TX, N9K-C9332PQ, N9K-C93120TX, N9K-9432PQ, N9K-9536PQ, N9K-9636PQ, N9K-X9632PC-QSFP100, N9K-X9432C-S, N3K-C3132Q-V, N3K-C3164Q, N3K-C3132C, N3K-C3232C, N3K-C3264Q, N3K-C3264C, N3K-3064Q, N3K-3016, N3K-3172.

You need to perform manual breakout using "interface breakout module <module number> port <port range> map <breakout mapping>" command.

- When a break-out port is configured as a part of a port-channel, you need to apply the configuration twice (after write-erase/reload), to ensure the effectiveness of the port-channel.
- Forward error correction (FEC) is mandatory for all cable types except for 1- and 2-meter passive copper cables. The Cisco's default mode FC-FEC CL74, and also supports CONS16-RS-FEC and IEEE-RS-FEC. There are two primary FEC algorithms used in 25G Ethernet:
 - **FC-FEC** (also known as "FireCode," "BASE-R," or "Clause 74") provides low-latency error protection (under 100 nanoseconds) optimized for bursty error correction. It is used on 3- and 5-meter passive copper cables, as well as on active optical 25G cables up to 10 meters in length. This FEC type is also utilized across all 100G interfaces.
 - **RS-FEC** (also referred to as "Reed Solomon," "Clause 91," or "Clause 108") offers better error protection. It is required for 25G multimode fiber (MMF) transceivers, such as Cisco SFP-25G-SR-S, supporting distances up to 100 meters. RS-FEC may also be necessary for active optical cables exceeding 10 meters.

All 25G devices support FC-FEC by default. The Cisco Nexus 9300-FX series is the first Cisco switch to support RS-FEC.

Beginning with Cisco NX-OS Release 7.0(3)I7(3) you see two additional options to configure FEC such as **rs-cons16** and **rs-ieee** as per IEEE standards.

The RS FEC IEEE (25G), enabled by the **fec rs-ieee** command on Cisco Nexus 9000 switches, implements RS-FEC based for enhanced error correction on high-speed Ethernet interfaces.

```
switch# (config-if)# fec ?
auto FEC auto
fc-fec CL74(25/50G)
off Turn FEC off
rs-cons16 RS FEC Consortium 1.6 (25G)
rs-fec CL91(100G) or Consortium 1.5 (25/50G)
rs-ieee RS FEC IEEE (25G)
```

- Beginning with Cisco NX-OS Release 7.0(3)I7(7) you can display the admin and oper status of FEC interface information with the **show interface fec** command.

Example:

```
switch# show interface fec
```



```

-----
Name   Ifindex Admin-fec Oper-fec   Status  Speed  Type
-----
Eth1/1   0x1a000000 auto   auto connected    10G  SFP-H10GB-AOC2M
Eth1/2   0x1a000200      Rs-fec notconnected    auto QSFP-100G-AOC3M
Eth1/3/1 0x38014000 auto   auto disabled auto QSFP-H40G-AOC3M
Eth1/3/2 0x38015000 auto   auto disabled auto QSFP-H40G-AOC3M
Eth1/3/3 0x38016000 auto   auto disabled auto QSFP-H40G-AOC3M
Eth1/3/4 0x38017000 auto   auto disabled auto QSFP-H40G-AOC3M
-----

```



Note Auto-FEC is not supported in Cisco NX-OS Release 7.0(3)I7(x)

Cisco Nexus 9000 C93180LC-EX Switch

For 7.0(3)I7(1) and later, Cisco Nexus 9000 C93180LC-EX switch provides three different modes of operation:

- Mode 1: 28 x 40G + 4 x 40G/100G (Default configuration)
 - Hardware profile portmode 4x100g + 28x40g.
 - 10x4 breakout is supported on the top ports from 1 to 27 (ports 1,3,5, 7...27). If any of the top port is broken out, the corresponding bottom port becomes non-operational. For example, if port 1 is broken out port 2 becomes non-operational.
 - 1 Gigabit and 10 Gigabit QSA is supported on ports 29, 30, 31, and 32. However, QSAs on the top and bottom front panel ports must be of same speed.
 - Ports 29, 30, 31, and 32 support 10x4, 25x4, and 50x2 breakout.
- Mode 2: 24 x 40G + 6 x 40G/100G
 - Hardware profile portmode 6x100g + 24x40g.
 - 10x4 breakout is supported on the top ports from 1 to 23 (ports 1,3,5, 7...23). If any of the top port is broken out the corresponding bottom port becomes non-operational.
 - Ports 25, 27, 29, 30, 31, and 32 support 10x4, 25x4, and 50x2 breakout.
 - 1 Gigabit and 10 Gigabit QSA is supported on ports 29, 30, 31, and 32. However, QSAs on the top and bottom front panel ports must be of same speed.
- Mode 3: 18 x 40G/100G
 - Hardware profile portmode 18x100g.
 - 10x4, 25x4, and 50x2 breakout is supported on top ports from 1 to 27 (ports 1,3,5, 7...27) and on ports 29,30,31,32.
 - 1 Gigabit and 10 Gigabit QSA is supported on all the 18 ports.

Changing Mode 3 to any other mode or vice versa requires **copy running-config startup-config** command followed by **reload** command to take effect. However, moving between Modes 1 and 2 is dynamic and requires only **copy running-config startup-config** command.

Use the **show running-config | grep portmode** command to display the current operation mode.

Example:

```
switch(config-if-range)# show running-config | grep portmode  
hardware profile portmode 4x100G+28x40G
```

With the Cisco Nexus C93180LC-EX switch, there are three breakout modes:

- 40G to 4x10G breakout ports
 - Enables the breakout of 40G ports into 4 X 10G ports.
 - Use the **interface breakout module 1 port x map 10g-4x** command.
- 100G to 4x25G breakout ports
 - Enables the breakout of 100G ports into 4 X 25G ports.
 - Use the **interface breakout module 1 port x map 25g-4x** command.
- 100G to 2x50G breakout ports
 - Enables the breakout of 100G ports into 2 X 50G ports.
 - Use the **interface breakout module 1 port x map 50g-2x** command.

Cisco Nexus 9000 C9364C-GX Switch

Cisco Nexus N9K-C9364C-GX breakout considerations:

- For ports 1-64, 2 x 50G, 4 x 25G and 4 x 10G breakout is supported only on odd numbered ports.
- When an odd numbered port in a quad is broken out, the even ports in that quad are removed and the other odd port in the same quad is broken out automatically to the same speed. For example, if port 1 or port 3 is broken out into 2 x 50, 4 x 25G or 4 x 10G, then the other odd port in that quad is automatically broken out to same speed and ports 2 and 4 in that quad are removed. When the above breakout configuration is removed, all ports in that quad revert to default.
- QSFP28 (100G) transceivers support the 4 x 25G breakout feature. Beginning Cisco NX-OS Release 9.3(5), the 2 x 50G breakout feature is supported.
- QSFP+ (40G) transceivers support the 4 x 10G breakout feature.
- 100G to 2x50G breakout ports
 - Enables the breakout of 100G ports into 2 X 50G ports on all odd ports.
 - Use the **interface breakout module 1 port x map 50g-2x** command.
- 40G to 4x10G breakout ports
 - Enables the breakout of 40G ports into 4 X 10G ports.
 - Use the **interface breakout module 1 port x map 10g-4x** command.

Cisco Nexus 9000 C93600CD-GX Switch

Cisco Nexus N9K-C93600CD-GX breakout considerations:

- In Cisco Nexus N9K-C93600CD-GX, every 4 ports from 1 through 24 are referred to as a quad. The breakout configuration and the speed must be same within a quad. The breakout feature may not function as expected if there is a mismatch of speed or breakout configuration within a quad. The six quads are made of ports 1-4, 5-8, 9-12, 13-16, 17-20 and 21-24.
- Beginning Cisco NX-OS Release 9.3(5), 2x50G breakout is supported on ports 1-36.
- 4x25G and 4x10G breakout is supported only on odd ports, between ports 1 through 24. The even ports will be purged within a quad (4 ports).
- When an odd-numbered port in a quad is broken out, the even ports in that quad are removed and the other odd ports within the quad is broken out automatically to the same speed. For example, if port 1 is broken out into 4x25G or 4x10G, then the other odd port in that quad is automatically broken out to same speed; and ports 2 and 4 in that quad are removed. When this breakout configuration is removed, all ports in that quad reverts to the default configuration.
- 2x50G breakout is supported on all ports from 1 through 24. All ports in a quad are broken out automatically to same speed when one port in a quad is broken out to 2x50G. For example when Port 2 is broken out into 2x50G, ports 1,3, and 4 are automatically broken out into 2x50G.



Note Only RS-FEC is supported on both lanes for 50G speed on ports 1 through 24

- Beginning with Cisco NX-OS Release 9.3(3) ports 25-28 support 4x10G, 4x25G, and 2x50G breakout features. These breakout feature are supported in port pairs. - for example 25-26 and 27-28.



Note Lane 2 of 2x50G should be configured with RS-FEC for link to be up.

- Beginning with Cisco NX-OS Release 9.3(3), consider the following breakout configuration for ports 29-36:
 - QSFP-DD-400G-DR4 transceivers support only the 4 x 100G breakout feature.
 - QSFP-DD-400G-FR4 and QSFP-DD-400G-LR8 transceivers do not support the breakout feature.
 - QSFP28 (100G) transceivers support the 2 x 50G and 4 x 25G breakout features.
 - QSFP+ (40G) transceivers support the 4 x 10G breakout feature.

Cisco Nexus 9000 C9316D-GX Switch

Cisco Nexus N9K-C9316D-GX breakout considerations:

- Port 1-16 breakout consideration:
 - QSFP-DD-400G-DR4 transceivers support only the 4 x 100G and 4x10G breakout feature.
 - QSFP-DD-400G-FR4 and QSFP-DD-400G-LR8 transceivers do not support the breakout feature.
 - QSFP28 (100G) transceivers support the 2 x 50G, 4 x 25G, and 4x10G breakout feature.

Virtual Device Contexts

Cisco NX-OS can segment operating system and hardware resources into virtual device contexts (VDCs) that emulate virtual devices. The Cisco Nexus 9000 Series switch does not support multiple VDCs. All switch resources are managed in the default VDC.

High Availability for Interfaces

Interfaces support stateful and stateless restarts. A stateful restart occurs on a supervisor switchover. After the switchover, Cisco NX-OS applies the runtime configuration.