



Overview

- [Licensing requirements](#) , on page 1
- [Supported platforms](#), on page 1
- [Interface parameters](#) , on page 1
- [Virtual device contexts](#), on page 16
- [High availability for interfaces](#), on page 16

Licensing requirements

To operate Cisco NX-OS, you must obtain and install appropriate licenses according to the features and platform requirements.

- Base (Essential) and add-on licenses are available for different feature sets.
- Licenses may be permanent, temporary, or evaluation, depending on product and purchase option.
- Advanced features require additional feature licenses beyond the base license.
- Licenses are applied and managed through the device command-line interface (CLI).

For detailed information on license types and installation instructions, see the [Cisco NX-OS Licensing Guide](#) and the [Cisco NX-OS Licensing Options Guide](#).

Supported platforms

The Nexus Switch Platform Support Matrix lists:

- Supported Cisco Nexus 9000 and 3000 switch models
- NX-OS software release versions

For the full platform-feature mapping, see the [Nexus Switch Platform Support Matrix](#).

Interface parameters

Interface parameters are configuration settings that

- define the operational characteristics of network interfaces,
- enable administrators to tailor interface behavior for specific roles, and
- support enhancements to performance, security, and connectivity.

Cisco NX-OS supports multiple configuration parameters for each supported interface type. Most of these parameters are described in this guide. Some parameters are described in other documents

The table provides sources for more information about configurable interface parameters.

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”
Layer 3	Bandwidth, delay, IP routing, Virtual Routing and Forwarding (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, Link Aggregation Control Protocol (LACP)	“Configuring Port Channels”
Security	Ethernet OAM Unidirectional (EOU)	<i>Cisco Nexus 9000 Series NX-OS Security Configuration Guide</i>

Best practice for Ethernet interfaces

Ethernet interfaces have these characteristics.

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

Best practices for quad group configuration on Cisco Nexus N9K-C9364C-GX and N9K-C93600CD-GX

Use these guidelines to configure quad groups on Cisco Nexus N9K-C9364C-GX and N9K-C93600CD-GX switches.

- Consecutive groups of four interfaces (1-4, 5-8, 9-12, and so on), form 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 can be active in a quad group at a time. The first link that comes up in a quad group sets its speed. Ports with other speeds are down and show **Link not connected**.

- If you mix different speeds in a quad group, the working speed is not recorded. If you insert and bring up a mismatched transceiver, all ports in the group reset. The first link that comes up after the reset sets the quad group speed. Pre-existing links may shut down. Remove the mismatched transceiver to recover.
- FC-FEC is not supported on the second lane of the 50G x 2 breakout port. The second breakout port does not come up when 50G x 2 breakout is configured. You must configure RS-FEC with 50G x 2 breakout.
- 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 in NRZ mode.
- Beginning with Cisco Nexus NX-OS Release 10.4(3)F, on N9K-C93600CD-GX and N9K-C9316D-GX, auto negotiation is not supported on 100G copper PAM4 links. You must configure **speed 100000** on peer side to bring the link up.
- Beginning with Cisco Nexus NX-OS Release 10.4(3)F, on N9K-C93600CD-GX, 100G PAM4 links is supported on ports 29-36 only.

Breakout port considerations on Nexus N9K-X9400-16W

These are the limitations for breakout ports are on the Cisco Nexus 9408 chassis and the Cisco N9K-X9400-16W (16x200G line-Ccrd expansion module (LEM)).

- Native port supports 100G, 40G, 10G on all ports.
- Breakout ports support 4 x 10G, 4 x 25G with these limitations:
 1. The 4 x 10G, 4 x 25G breakout ports are supported only on odd ports.
 2. If you configure breakout x 4 on an odd port, the corresponding even port is purged automatically.
- Breakout ports support 2 x 50G with these limitations.
 1. The 2 x 50G breakout is supported on odd and even ports.
 2. When the 2 x 50G breakout is configured on an odd or even port, the corresponding even or odd port is broken out 2 x 50G automatically.
- 10G using QSA is supported on all ports with these limitations:
 1. When a 10G, 40G, or 100G transceiver is present on an odd or even port in the linked up state, it does not allow any other speed on the corresponding even or odd port.

A warning or syslog is printed for the mismatched XCVR, and the port status changes to **speed mismatch** state for the XCVR port that was inserted later.

The port status is indicated in the **show interface brief** and **show interface status** command outputs.
 2. When odd port has 40G or 100G and corresponding even port has 10G transceiver or vice versa, and are in **admin shut** state, then these conditions hold.
 - No precedence is decided as long as the ports remain admin shut, whichever port is configured as **no shutdown** gets the first precedence.
 - If both ports are configured as **no shutdown** at the same time, then the first detected port by software gets precedence, and the other gets **xcvr mismatch** state.

If switch is reloaded, then during boot-up the port which is detected first by software takes precedence and rest gets **speed mismatch** state.

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

- For ports 1-16, each pair of ports (1,2 | 3,4 | 5,6 | 7,8 | 9,10 | 11,12 | 13,14 | 15,16) forms a quad group.
- All the ports in a quad operate at 10G with QSA, or 40G or 100G or 200G.
- Mixed speed is not supported within the same quad with these exceptions:
 - Mixed speed of 40G and 100G can be supported in quad
 - However, 100G-CR2 cannot be mixed with either 40G or other types of 100G optics in quad

- The quad speed mismatch check runs on optics insertion and removal sequence. The first inserted transceiver in a quad group determines the speed of the quad group.

The ports with unsupported speed is down with the reason of **XCVR speed mismatch**. With unsupported mixed speeds, only one speed is up in a quad group at a time.

- For a particular port to be up and functional, ensure to remove all the optics or cables from all the ports in that quad and plug in the optics or cables first in the port that needs to be up and then plug in the other optics or cables.
- For a particular speed mismatch port to be up and functional, ensure to remove the optics or cables from all other ports in that quad, flap the needed port, and then plug in the other ports.
- Save the configuration (copy running start-up) to preserve the port state.
- When a mismatch transceiver is plugged into a quad, syslog is generated.

```
Interface Ethernet1/X is down (Reason: Inserted transceiver speed mismatch with quad speed Y)
```

- After reload ascii, port states might change depending on the order in which interfaces are detected
- Ensure you use only transceivers of the same speed in a quad to avoid disruption or entering a non-deterministic state.



Note Ensure LEM is powered on and online before removing or inserting optics. Removing or inserting optics when it is powered off or offline can cause software to miss the optics and result in inconsistent port states.

Port considerations Cisco Nexus N9K-X9400-22L

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

- For ports 1-22, each group of consecutive four ports 1-4, 5-8, 11-14, 15-18, 19-22 and two ports 9-10 is referred to as a quad group.
- All the ports in a quad operate in 10G, or 25G or 50G.
- Mixed speed is not supported within the same quad with these exceptions:
 - Mixed speed of 10G and 25G can be supported in quad.

- Quad speed mismatch check runs on optics insertion and removal sequence. The first inserted transceiver in a quad group determines the speed of the quad group.

The ports with unsupported speed is down with the reason of **XCVR speed mismatch**. With unsupported mixed speeds, only one speed is up in a quad group at a time.

- To make a particular port functional, remove all optics or cables from every port in that quad. Plug the optics or cables into the desired port first, then connect the others.
- To make a port with a speed mismatch functional, remove optics or cables from the other ports in that quad. Flap the required port and then connect the other ports.
- Save the configuration (copy running start-up) to preserve port state persistence.
- When a mismatch transceiver is plugged into a quad, syslog is generated as **Interface Ethernet1/X is down (Reason: Inserted transceiver speed mismatch with quad speed Y)**.
- The port states may not be persistent on reload `ascii`. port states depends on the order of interface detected sequence on reload `ascii`.
- Ensure to only use the transceivers of same speed in a quad to avoid any disruption or indeterministic state.
- If all dual speed optics in one quad are set to non-default speeds and **reload `ascii`** is performed, some ports may be down with reason **xcvr speed mismatch**.

Use **shut** and **no shut** commands on these ports to bring them up and make them functional.



Note

Do not remove or insert optics when LEM is either powered off or offline. Otherwise, the software cannot detect the optics and might lead to inconsistent port-state.

Access ports

An access port is a Layer 2 switchport that carries traffic for only a single 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 Layer 3 interface that you configure on a physical switch port (not a virtual interface). It routes IP traffic to another device.

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

Management interface

A management interface is a network interface that

- provides dedicated connectivity for device administration,
- operates independently from data traffic interfaces, and
- supports remote access protocols such as Telnet and SNMP.

You use the management interface (commonly labeled as mgmt0) to detect connection types automatically. It supports full-duplex mode and operates at speeds of 10, 100, or 1000 Megabits per second.

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

Port-Channel interfaces

A port-channel interface is a logical network interface that

- aggregates multiple physical interfaces into a single channel,
- increases bandwidth and enhances redundancy, and
- supports up to 32 bundled Ethernet links.

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

A subinterface is a virtual interface that

- operates under a parent physical or port-channel interface,
- allows assignment of unique Layer 3 parameters such as IP addresses and routing protocols, and
- enables division of a single physical interface into multiple, independently configured virtual interfaces.

You can create virtual subinterfaces by configuring a parent interface as a Layer 3 interface.

Loopback interfaces

A loopback interface is a virtual network interface that

- has a single endpoint and is always operational,
- immediately receives any packet it transmits, and
- emulates the behavior of a physical interface without connecting to external devices.

Loopback interfaces are often used for testing, diagnostics, or internal routing purposes, as they guarantee the interface remains active regardless of hardware state. For more information about subinterfaces, see the *Loopback Interfaces* section.

Breakout interfaces

A breakout interface is a high-speed network port feature that

- splits a single high-bandwidth physical port into multiple lower-speed logical interfaces,
- enables a switch or router to connect to several lower-speed devices simultaneously, and

- maximizes port utilization by allowing flexibility in network configuration.

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 breakouts on ports

A module-level breakout is a port configuration technique that

- enables splitting of certain high-density ports into multiple lower-bandwidth ports,
- provides increased network configuration flexibility, and
- supports a range of port breakdown options such as 4x10G, 4x25G, 4x50G, etc.

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

Some modules break down all the ports into 4x10G, 4x25G, 4x50G, 4x100G, 2x50G, or 2x100G configurations.

Example: Module level breakout

For example, a module level breakout of 4x10G splits a 40G interface into four 10G interfaces. When you execute the command, the module reloads and removes the existing interface configurations.

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

To undo a breakout, use the **no interface breakout module module_number** command. This restores ports to their original configuration and deletes previous breakout configurations.

Lane Selectors

A lane selector is a control panel feature that

- consists of a push-button switch and four LEDs,
- enables users to view the link or activity status of switch ports, and
- supports switching between 1 x 40G and 4 x 10G configurations on compatible Cisco Nexus 9000 Series switches and the Cisco Nexus 3164 and 3232 switches.

Additional information

Lane selectors are located on the left side of the Cisco Nexus switch front panel and are labeled 'LS'.

When used in a 1 x 40G configuration, LEDs indicate the link/activity status of the main port. When configured for 4 x 10G, pressing the push button cycles the LEDs through the status of each 10G port. On the last press, all LEDs extinguish, and the display resets to the default mode.

By pressing the lane selector push button, the port LED shows the selected lane's link/activity status.

The first time the push button is pressed, the first LED displays the status of the first port. Pressing the push button a second time displays the status of the second port, and so on. To display the status of each of the four ports, press the push button as described.

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

Examples

If port 60 is configured as 4 x 10G, pressing the lane selector once displays the link status of 60/1/1, twice for 60/1/2, and so on.



Note The lane selector does not manage ports not configured for link/activity monitoring.

Guidelines

When a port is in 10G breakout mode and no lane is selected, the 40G port's LED lights green, even if only one of the 10G breakout ports is up.

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

Breakout port support on Cisco Nexus switches

The matrix provides detailed information about supported breakout modes (for example, 4x10G, 4x25G, 2x50G, etc.) for Cisco Nexus switches and line card platforms. 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-FX3									
N9K-C93108TC-FX3P									
N9K-C93180YC-FX3									
N9K-C9348GC-FX3									
N9K-C9348GC-FX3P									
N9K-C9364C-H1	Yes	Yes	Yes	No	No	No	No	No	No
N9K-C93400LD-H1	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No
N9K-C9332D-H2R	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No
N9K-X9736C-FX3	Yes	Yes	Yes	No	No	No	No	No	No
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

Switches	4x10G	4x25G	2x50G	2x100G	2x200G	2x400G	4x50G	4x100G	8x100G
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-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
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-C9364D-GX2A	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

Switches	4x10G	4x25G	2x50G	2x100G	2x200G	2x400G	4x50G	4x100G	8x100G
N9K-X9836DM-A	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No
N9364E-SG2-Q	No	No	No	No	No	Yes	No	Yes	Yes

Guidelines and limitations for breakout ports

- The Cisco Nexus 9516 switch does not support breakout on Modules 8 to 16.
- Starting with Cisco NX-OS Release 7.0(3)F2(1), the 36-port 100-Gigabit Ethernet QSFP28 line cards (N9K-X9636C-R) and the 36-port 40-Gigabit Ethernet QSFP+ line cards (N9K-X9636Q-R) support breakout to 4 x 10G.
- Starting with Cisco NX-OS Release 9.2(1), the N9K-9636C-R, N9K-X9636Q-R, and N9K-X9636C-RX line cards support breaking out 40G ports into 4 x 10G.
- Starting with Cisco NX-OS Release 9.2(2), N9K-X9636C-R and N9K-X9636C-RX line cards support break out of 100G ports into 4 x 25G. The N9K-C9636C-R does not support RS-FEC.

Starting with Cisco NX-OS Release 9.3(3), the default FEC mode on N9K-X9636C-R and N9K-X9636C-RX is FC-FEC for 25G x 4 and 50G x 2.

When connecting N9K-X9636C-RX to N9K-X9636C-R, configure FC-FEC on N9K-X9636C-RX because RS-FEC is not supported.

The N9K-X96136YC-R line card does not support breakout.

- Starting with Cisco NX-OS Release 9.3(3), these switches support breakout.

The Cisco Nexus 93600CD-GX switch and the Cisco Nexus 9500 R-Series switches support breaking out 100G ports into 2 x 50G.

On Nexus 9500 R-Series switches with N9K-X9636C-R and N9K-X9636C-RX line cards, only specific optics (QSFP-100G-PSM4-S, QSFP-100G-AOC, QSFP-100G-CU1M, and CU3M) support 2 x 50G and 4 x 25G breakout.

For more information see *Cisco Optics-to-Device Compatibility Matrix*.

- Starting with Cisco NX-OS Release 10.4(3), the Cisco N9K-X98900CD-A switch supports breakout on 4 x 25G port.

In releases prior to Cisco NX-OS Release 10.4(3), breakout is not supported on 4 x 25G port.

Best practices for manual breakout configuration

You must use the **interface breakout module module number port port range map breakout mapping** command when performing manual breakout on Cisco Nexus devices.

- When you upgrade a Cisco Nexus 9000 device to Cisco NX-OS Release 7.0(3)I7(2) or later, interfaces configured with manual breakout using a QSA are no longer supported. You must remove the configuration and manually reconfigure the breakout settings for the affected interface.



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



Note This restriction does not apply to the following platforms, where manual breakout remains fully supported—N9K-C93128TX, N9K-9332, N9K-C9396PX, N9K-C9396TX, N9K-C9372PX, N9K-C9372TX, N9K-C9332PQ, 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—because manual breakout is supported on these platforms.

- Manual breakout is supported on the following platforms because auto-breakout does not occur 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.

Forward error correction (FEC) settings for breakout ports

FEC is required on all cable types except for 1-meter and 2-meter passive copper cables. Cisco switches use FC-FEC CL74 by default. You can configure RS-FEC Consortium 1.6, RS-FEC IEEE, and other FEC algorithms.



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

When configuring a break-out port, ensure that the FEC is matching for the link to be up.

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 supports RS-FEC.

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

Enable the RS FEC IEEE (25G) using the **fec rs-ieee** command on Cisco Nexus 9000 switches to implement RS-FEC 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 operational 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

Breakout modes on Cisco Nexus C9364C-H1 switch

Starting with Cisco NX-OS Release 10.5(3)F, the Cisco Nexus C9364C-H1 switch supports breakout mode.

Breakout modes are port configuration settings on the Cisco Nexus C9364C-H1 switch that

- allow a single port to be split into multiple logical interfaces (such as 2x50G, 4x25G, or 4x10G),
- are available only on the first port of every front port quad grouping (e.g., ports 1, 5, 9, ...).



Note During breakout of the interface, the three adjacent front ports are removed, and are not visible in the interface verification or configuration commands.

Cisco Nexus 9000 C93180LC-EX switch - Operation and breakout modes

Operation and breakout modes are switch configuration profiles. These profiles let you group and set ports, split high-speed physical ports into multiple lower-speed logical ports, and find out which types of equipment and cabling you can use for each mode.

Cisco Nexus 9000 C93180LC-EX modes

Operation modes are switch configuration profiles that

- determine available bandwidth and port groupings
- enable different breakout capabilities, and
- require you to use distinct configuration procedures to switch between modes.

The Cisco Nexus 9000 C93180LC-EX switch supports three operation modes (7.0(3)I7(1) and later):

- Mode 1: 28 x 40G + 4 x 40G/100G (Default configuration)**

This is a hardware profile port mode 4x100G + 28x40G ports. It supports:

- Breakout support of 10 x 4 on top ports from 1 to 27 (ports 1,3,5, 7...27).

If you break out any of the top ports, 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 support on ports 29, 30, 31, and 32. However, QSAs on the top and bottom front panel ports must be of same speed.
- Breakout support of 10 x 4, 25 x 4, and 50 x 2 on ports 29, 30, 31, and 32.

- **Mode 2: 24 x 40G + 6 x 40G/100G**

This hardware profile port mode 6 x 100G + 24 x 40G ports. It supports:

- Breakout support of 10 x 4 on 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.
- Breakout support of 10 x 4, 25 x 4, and 50 x 2 on ports 25, 27, 29, 30, 31, and 32.
- 1 Gigabit and 10 Gigabit QSA support 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**

This hardware profile port mode 18 x 100G that ports. It supports:

- Breakout support of 10 x 4, 25 x 4, and 50 x 2 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 on all the 18 ports.

To change from Mode 3 to another mode, enter the **copy running-config startup-config** command followed by **reload** command to take effect. However, to move between Modes 1 and 2, you only need to enter the **copy running-config startup-config** command.

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

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

Breakout modes

The Cisco Nexus C93180LC-EX switch has three breakout modes.

- Support for 40G to 4 x 10G breakout ports
 - This mode enables the breakout of 40G ports into 4 x 10G ports.
 - To configure this mode, use the **interface breakout module 1 port x map 10g-4x** command.
- Support for 100G to 4 x 25G breakout ports
 - This mode enables the breakout of 100G ports into 4 x 25G ports.
 - To configure this mode, use the **interface breakout module 1 port x map 25g-4x** command.
- Support for 100G to 2 x 50G breakout ports
 - This mode enables the breakout of 100G ports into 2 x 50G ports.

- To configure this mode, use the **interface breakout module 1 port x map 50g-2x** command.

Breakout considerations for Cisco Nexus 9000 C9364C-GX switch

These are breakout considerations for Cisco Nexus N9K-C9364C-GX switches.

- Configure breakout modes—1-64, 2 x 50G, 4 x 25G, and 4 x 10G—only on odd-numbered ports.



Note Do not attempt break out on even-numbered ports.

- When you break out an odd-numbered port, even-numbered ports in that quad are automatically removed, and the other odd port is configured to the same breakout 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 set to the same speed and ports 2 and 4 in that quad are removed. All ports in that quad revert to default when this breakout configuration is removed.

- To revert a quad to default port status, remove the breakout configuration from both odd ports in the quad.
- 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.
- Use the interface **breakout module 1 port x map 50g-2x** command to enable the breakout of 100G ports into 2 x 50G ports on all odd ports.
- Use the interface **breakout module 1 port x map 10g-4x** command to enable the breakout of 40G ports into 4 x 10G ports.

Breakout features on Cisco Nexus 9000 C93600CD-GX switches

Use the breakout considerations on the Cisco Nexus N9K-C93600CD-GX.

- In Cisco Nexus N9K-C93600CD-GX, every four ports from 1 through 24 are referred to as a quad.



Note The breakout configuration and the speed must be the 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 consist of ports 1–4, 5–8, 9–12, 13–16, 17–20, and 21–24.

- Beginning Cisco NX-OS Release 9.3(5), 2 x 50G breakout feature is supported on ports 1-36.
- 4 x 25G and 4 x 10G breakout features are supported only on odd ports, between ports 1 through 24. The even ports within a quad are removed (four 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 are broken out automatically broken out to the same speed.

For example, if port 1 is broken out into 4 x 25G or 4 x 10G, the other odd ports within the quad are automatically broken out to the same speed, and ports 2 and 4 in that quad are removed. When this breakout configuration is removed, all ports in that quad revert to the default configuration.

- 2 x 50G breakout is supported on all ports from 1 through 24. All ports in a quad are broken out automatically to the same speed when one port in the quad is broken out to 2 x 50G.

For example, when Port 2 is broken out into 2 x 50G, ports 1, 3, and 4 are automatically broken out into 2 x 50G.



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 4 x 10G, 4 x 25G, and 2 x 50G breakout features. These breakout features are supported in port pairs, for example 25-26 and 27-28.



Note Lane 2 of 2 x 50G should be configured with RS-FEC for the link to be up.

- Beginning with Cisco NX-OS Release 9.3(3), ports 29-36 support these breakout configurations.
 - 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 features.
 - QSFP28 (100G) transceivers support 2 x 50G and 4 x 25G breakout features.
 - QSFP+ (40G) transceivers support 4 x 10G breakout features.

Breakout considerations on Cisco Nexus C9316D-GX switches

Use these breakout considerations for ports 1 through 16 on the Cisco Nexus N9K-C9316D-GX switch.

- QSFP-DD-400G-DR4 transceivers support only the 4 x 100G and 4 x 10G breakout features.



Note QSFP-DD-400G-FR4 and QSFP-DD-400G-LR8 transceivers *do not* support the breakout features.

- QSFP28 (100G) transceivers support the 2 x 50G, 4 x 25G, and 4 x 10G breakout feature.

Breakout considerations on Cisco Nexus 93C64E-SG2-Q switch

You can use the breakout feature, a switch port capability that allows a single high-speed port to be divided into multiple lower-speed ports for flexible connectivity.

Starting with Cisco NX-OS Release 10.5(3)F, the Cisco Nexus 93C64E-SG2-Q switch provides

- breakout configurations of 2x400G and 8x100G,
- compatibility with supported optics, and
- flexible port configuration.

The supported breakout modes include:

- **2x400G breakout:** Divide a single port into two 400G ports.
- **8x100G breakout:** Divide a single port into eight 100G ports.

Optics

Beginning with Cisco NX-OS Release 10.5(3)F, the Cisco Nexus 93C64E-SG2-Q switch supports these optics.

- QDD-8X100G-FR
- QDD-8x100G-LR
- QDD-2X400G-FR4
- QDD-2x400G-LR4

The Cisco Nexus 93C64E-SG2-Q switch also supports 64 QSFP-DD800 ports. This enables high-density and high-speed connectivity.

Virtual device contexts

A virtual device context (VDC) is a network virtualization technology that

- segments operating system and hardware resources,
- emulates independent logical switches within a physical switch, and
- allows separate configuration, administration, and management for each context.

The Cisco Nexus 9000 Series switch does not support multiple VDCs. All switch resources are managed in the default VDC.

High availability for interfaces

High availability for interfaces is a network feature that

- enables interfaces to continue operating during supervisor switchovers, and
- supports both stateful and stateless restart mechanisms.

A stateful restart occurs on a supervisor switchover. After the switchover, Cisco NX-OS applies the runtime configuration.