Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 15.3(2)

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Introduction

The Cisco NX-OS software for the Cisco Nexus 9000 series switches is a data center, purpose-built operating system designed with performance, resiliency, scalability, manageability, and programmability at its foundation. It provides a robust and comprehensive feature set that meets the requirements of virtualization and automation in data centers.

This release works only on Cisco Nexus 9000 Series switches in ACI mode.

This document describes the features, issues, and limitations for the Cisco NX-OS software. For the features, issues, and limitations for the Cisco Application Policy Infrastructure Controller (APIC), see the Cisco Application Policy Infrastructure Controller Release Notes, Release 5.3(2).

For more information about this product, see " Related Content."

Date	Description
May 4, 2026	Defect CSCwt09384 was added to the list of Open issues.
October 10, 2025	Defect CSCwn63197 was added to the list of Open Issues.
February 10, 2025	Release 15.3(2f) became available. Added the resolved bugs for this release.
December 4, 2024	Release 15.3(2e) became available.
September 27, 2024	Release 15.3(2d) became available. Added the resolved bugs for this release.
June 27, 2024	Release 15.3(2c) became available. Added the resolved bugs for this release.
April 2, 2024	Release 15.3(2b) became available. Added the resolved bugs for this release.
February 1, 2024	Release 15.3(2a) became available.

Supported Hardware

Table 1. Modular Spine Switches

Product ID	Description
N9K-C9504	Cisco Nexus 9504 switch chassis
N9K-C9508	Cisco Nexus 9508 switch chassis
N9K-C9508-B1	Cisco Nexus 9508 chassis bundle with 1 supervisor module, 3 power supplies, 2 system controllers, 3 fan trays, and 3 fabric modules
N9K-C9508-B2	Cisco Nexus 9508 chassis bundle with 1 supervisor module, 3 power supplies, 2 system controllers, 3 fan trays, and 6 fabric modules
N9K-C9516	Cisco Nexus 9516 switch chassis

Table 2. Modular Spine Switch Line Cards

© 2025 Cisco and/or its affiliates. All rights reserved.

Product ID	Description	Maximum Quantity		
		Cisco Nexus 9504	Cisco Nexus 9508	Cisco Nexus 9516
N9K-X9716D-GX	Cisco Nexus 9500 16-port 400 Gigabit Ethernet QSFP line card	4	8	16
N9K-X9736C-FX	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet Cloud Scale line card	4	8	16
N9K-X9736Q-FX	Cisco Nexus 9500 36-port 40 Gigabit Ethernet Cloud Scale line card	4	8	16
N9K-X9732C-EX	Cisco Nexus 9500 32-port, 40/100 Gigabit Ethernet Cloud Scale line card Note: The N9K-X9732C-EX line card cannot be used when a fabric module is installed in FM slot 25.	4	8	16

Table 3. Modular Spine Switch Fabric Modules

Product ID	Description	Minimum	Maximum
N9K-C9504-FM-G	Cisco Nexus 9508 cloud scale fabric module (400G capable)	4	5
N9K-C9508-FM-G	Cisco Nexus 9508 cloud scale fabric module (400G capable)	4	5
N9K-C9504-FM-E	Cisco Nexus 9504 cloud scale fabric module	4	5
N9K-C9508-FM-E	Cisco Nexus 9508 cloud scale fabric module	4	5
N9K-C9508-FM-E2	Cisco Nexus 9508 cloud scale fabric module	4	5
N9K-C9516-FM-E2	Cisco Nexus 9516 cloud scale fabric module	4	5

Table 4. Modular Spine Switch Fans

Product ID	Description
N9K-C9504-FAN2	Nexus 9500 4-slot Fan Tray (Gen 2)
N9K-C9504-FAN-PWR	Nexus 9500 4-slot Fan Tray Power Card Blank
N9K-C9504-FAN	Fan tray for Cisco Nexus 9504 chassis
N9K-C9508-FAN2	Nexus 9500 8-slot Fan Tray (Gen 2)
N9K-C9508-FAN-PWR	Nexus 9500 8-slot Fan Tray Power Card Blank
N9K-C9508-FAN	Fan tray for Cisco Nexus 9508 chassis
N9K-C9516-FAN	Fan tray for Cisco Nexus 9516 chassis

Table 5. Modular Spine Switch Supervisor and System Controller Modules

© 2025 Cisco and/or its affiliates. All rights reserved.

Product ID	Description
N9K-SUP-A+	Cisco Nexus 9500 Series supervisor module
N9K-SUP-B+	Cisco Nexus 9500 Series supervisor module
N9K-SUP-A	Cisco Nexus 9500 Series supervisor module
N9K-SUP-B	Cisco Nexus 9500 Series supervisor module
N9K-SC-A	Cisco Nexus 9500 Series system controller

Table 6. Fixed Spine Switches

Product ID	Description
N9K-C9364D-GX2A	Cisco Nexus 9300 platform switch with 64 400/100-Gbps QSFP-DD ports and 2 1/10 SFP+ ports.
N9K-C9348D-GX2A	Cisco Nexus 9300 platform switch with 48 400/100-Gbps QSFP-DD ports and 2 1/10 SFP+ ports.
N9K-C9332D-GX2B	Cisco Nexus 9300 platform switch with 32 400/100-Gbps QSFP-DD ports and 2 1/10 SFP+ ports.
N9K-C93600CD-GX	Cisco Nexus 9300 platform switch with 28 10/40/100-Gigabit Ethernet QSFP28 ports (ports 1-28) and 8 10/40/100/400-Gigabit QSFP-DD ports (ports 29-36).
N9K-C9316D-GX	Cisco Nexus 9300 platform switch with 16 10/40/100/400-Gigabit QSFP-DD ports (ports 1-16).
N9K-C9332C	Cisco Nexus 9300 platform switch with 32 40/100-Gigabit QSFP28 ports and 2 SFP ports. Ports 25-32 offer hardware support for MACsec encryption.
N9K-C9364C-GX	Cisco Nexus 9300 platform switch with 64 100-Gigabit Ethernet QSFP28 ports, two management ports (one RJ-45 port and one SFP port), one console port (RS-232), and one USB port.
N9K-C9364C	Cisco Nexus 9300 platform switch with 64 40/100-Gigabit QSFP28 ports and two 1/10-Gigabit SFP+ ports. The last 16 of the QSFP28 ports are colored green to indicate that they support wire-rate MACsec encryption.

Table 7. Fixed Spine Switch Power Supply Units

Product ID	Description
N9K-PAC-1200W	1200W AC power supply, port side intake pluggable Note: This power supply is supported only by the Cisco Nexus 93120TX and 9336PQ ACI-mode switches
N9K-PAC-1200W-B	1200W AC power supply, port side exhaust pluggable Note: This power supply is supported only by the Cisco Nexus 93120TX and 9336PQ ACI-mode switches
NXA-PAC-1200W-PE	1200W AC power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance

© 2025 Cisco and/or its affiliates. All rights reserved.

Product ID	Description
NXA-PAC-1200W-PI	1200W AC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance
NXA-PAC-1100W-PE2	1100W AC power supply, port side exhaust pluggable
NXA-PAC-1100W-PI2	1100W AC power supply, port side intake pluggable
NXA-PAC-750W-PE	750W AC power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance Note: This power supply is supported only on release 14.2(1) and later.
NXA-PAC-750W-PI	750W AC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance Note: This power supply is supported only on release 14.2(1) and later.
NXA-PDC-1100W-PE	1100W AC power supply, port side exhaust pluggable
NXA-PDC-1100W-PI	1100W AC power supply, port side intake pluggable
NXA-PDC-930W-PE	930W AC power supply, port side exhaust pluggable
NXA-PDC-930W-PI	930W AC power supply, port side intake pluggable
NXA-PHV-1100W-PE	1100W HVAC/HVDC power supply, port-side exhaust
NXA-PHV-1100W-PI	1100W HVAC/HVDC power supply, port-side intake
N9K-PUV-1200W	1200W HVAC/HVDC dual-direction airflow power supply

Table 8. Fixed Spine Switch Fans

Product ID	Description
N9K-C9300-FAN3	Burgundy port side intake fan
N9K-C9300-FAN3-B	Blue port side exhaust fan
NXA-FAN-160CFM-PE	Blue port side exhaust fan
NXA-FAN-160CFM-PI	Burgundy port side intake fan
NXA-FAN-35CFM-PE	Blue port side exhaust fan
NXA-FAN-35CFM-PI	Burgundy port side intake fan

Table 9. Fixed Leaf Switches

Product ID	Description
N9K-C9364D-GX2A	Cisco Nexus 9300 platform switch with 64 400/100-Gbps QSFP-DD ports and 2 1/10 SFP+ ports.
N9K-C9348D-GX2A	Cisco Nexus 9300 platform switch with 48 400/100-Gbps QSFP-DD ports and 2 1/10 SFP+

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	Description
	ports.
N9K-C9332D-GX2B	Cisco Nexus 9300 platform switch with 32 400/100-Gbps QSFP-DD ports and 2p 1/10 SFP+ ports.
N9K-C9316D-GX	Cisco Nexus 9300 platform switch with 16 10/40/100/400-Gigabit QSFP-DD ports (ports 1-16).
N9K-C9364C-GX	Cisco Nexus 9300 platform switch with 64 100-Gigabit Ethernet QSFP28 ports, two management ports (one RJ-45 port and one SFP port), one console port (RS-232), and one USB port.
N9K-C93600CD-GX	Cisco Nexus 9300 platform switch with 28 10/40/100-Gigabit Ethernet QSFP28 ports (ports 1-28) and 8 10/40/100/400-Gigabit QSFP-DD ports (ports 29-36).
N9K-C93180YC-FX3	Cisco Nexus 9300 platform switch with 48 100M/1/10/25-Gigabit Ethernet SFP28 ports, 6 40/100-Gigabit QSFP28 ports, two management ports (one 10/100/1000BASE-T and one SFP+), one console port (RS-232), and one USB port.
N9K-C93108TC-FX3P	Cisco Nexus 9300 platform switch with 48 100M/1/10-GBASE-T (copper) ports, 6 40/100-Gigabit QSFP28 ports, two management ports (one 10/100/1000BASE-T port and one SFP port), one console port (RS-232), and one USB port.
N9K-C93240YC-FX2	Cisco Nexus 9300 platform switch with 48 1/10/25-Gigabit Ethernet SFP28 ports and 12 40/100-Gigabit Ethernet QSFP28 ports. The N9K-C93240YC-FX2 is a 1.2-RU switch. Note: 10/25G-LR-S with QSA is not supported.
N9K-C93216TC-FX2	Cisco Nexus 9300 platform switch with 96 1/10GBASE-T (copper) front panel ports and 12 40 /100-Gigabit Ethernet QSFP28 spine-facing ports
N9K-C93360YC-FX2	Cisco Nexus 9300 platform switch with 96 1/10/25-Gigabit front panel ports and 12 40 /100-Gigabit Ethernet QSFP spine-facing ports. Note: The supported total number of fabric ports and port profile converted fabric links is 56.
N9K-C9336C-FX2-E	Cisco Nexus C9336C-FX2 Top-of-rack (ToR) switch with 36 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: 1-Gigabit QSA is not supported on ports 1/1-6 and 1/33-36. The port profile feature supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.
N9K-C9336C-FX2	Cisco Nexus C9336C-FX2 Top-of-rack (ToR) switch with 36 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: 1-Gigabit QSA is not supported on ports 1/1-6 and 1/33-36. The port profile feature supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.
N9K-C93108TC-FX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.
N9K-C93108TC-FX-24	Cisco Nexus 9300 platform switch with 24 1/10GBASE-T (copper) front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	Description
	module drops.
N9K-C93180YC-FX	Cisco Nexus 9300 platform switch with 48 1/10/25-Gigabit Ethernet SFP28 front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. The SFP28 ports support 1-, 10-, and 25-Gigabit Ethernet connections and 8-, 16-, and 32-Gigabit Fibre Channel connections. Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.
N9K-C93180YC-FX-24	Cisco Nexus 9300 platform switch with 24 1/10/25-Gigabit Ethernet SFP28 front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. The SFP28 ports support 1-, 10-, and 25-Gigabit Ethernet connections and 8-, 16-, and 32-Gigabit Fibre Channel connections. Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.
N9K-C9348GC-FXP	Cisco Nexus 9348GC-FXP switch with 48 100/1000-Megabit 1GBASE-T downlink ports, 4 10-/25-Gigabit SFP28 downlink ports, and 2 40-/100-Gigabit QSFP28 uplink ports.
N9K-C93108TC-EX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) front panel ports and 6 40/100-Gigabit QSFP28 spine facing ports.
N9K-C93108TC-EX-24	Cisco Nexus 9300 platform switch with 24 1/10GBASE-T (copper) front panel ports and 6 40/100-Gigabit QSFP28 spine facing ports.
N9K-C93180LC-EX	Cisco Nexus 9300 platform switch with 24 40-Gigabit front panel ports and 6 40/100-Gigabit QSFP28 spine-facing ports. The switch can be used as either a 24 40G port switch or a 12 100G port switch. If 100G is connected the Port1, Port 2 will be HW disabled.
N9K-C93180YC-EX	Cisco Nexus 9300 platform switch with 48 1/10/25-Gigabit front panel ports and 6-port 40/100 Gigabit QSFP28 spine-facing ports.
N9K-C93180YC-EX-24	Cisco Nexus 9300 platform switch with 24 1/10/25-Gigabit front panel ports and 6-port 40/100 Gigabit QSFP28 spine-facing ports.
N9K-C93120TX	Cisco Nexus 9300 platform switch with 96 1/10GBASE-T (copper) front panel ports and 6-port 40-Gigabit Ethernet QSFP spine-facing ports.

Table 10. Fixed Leaf Switch Power Supply Units

Product ID	Description
NXA-PAC-2KW-PE	Nexus 9000 2KW AC power supply, port-side exhaust Note: This power supply is supported only by the Cisco Nexus 9364C-GX ACI-mode switch.
NXA-PAC-2KW-PI	Nexus 9000 2KW AC power supply, port-side intake Note: This power supply is supported only by the Cisco Nexus 9364C-GX ACI-mode switch.
N9K-PAC-1200W	1200W AC power supply, port side intake pluggable Note: This power supply is supported only by the Cisco Nexus 93120TX and 9336PQ ACI-mode switches

© 2025 Cisco and/or its affiliates. All rights reserved.

Product ID	Description
N9K-PAC-1200W-B	1200W AC power supply, port side exhaust pluggable Note: This power supply is supported only by the Cisco Nexus 93120TX and 9336PQ ACI-mode switches
N9k-PAC-3000W-B	3000W AC power supply, port side intake
N9K-PAC-650W	650W AC power supply, port side intake pluggable
N9K-PAC-650W-B	650W AC power supply, port side exhaust pluggable
NXA-PAC-1200W-PE	1200W AC power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance
NXA-PAC-1200W-PI	1200W AC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance
NXA-PAC-1100W-PE2	1100W AC power supply, port side exhaust pluggable
NXA-PAC-1100W-PI2	1100W AC power supply, port side intake pluggable
NXA-PAC-750W-PE	750W AC power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance Note: This power supply is supported only on release 14.2(1) and later.
NXA-PAC-750W-PI	750W AC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance Note: This power supply is supported only on release 14.2(1) and later.
NXA-PAC-650W-PE	650W AC power supply, port side exhaust pluggable
NXA-PAC-650W-PI	650W AC power supply, port side intake pluggable
NXA-PAC-500W-PE	500W AC Power supply, port side exhaust pluggable
NXA-PAC-500W-PI	500W AC Power supply, port side intake pluggable
NXA-PAC-350W-PE	350W AC power supply, port side exhaust pluggable
NXA-PAC-350W-PI	350W AC power supply, port side intake pluggable
NXA-PDC-2KW-PE	Nexus 9000 2KW DC power supply, port-side exhaust Note: This power supply is supported only by the Cisco Nexus 9364C-GX ACI-mode switch.
NXA-PDC-2KW-PI	Nexus 9000 2KW DC power supply, port-side intake Note: This power supply is supported only by the Cisco Nexus 9364C-GX ACI-mode switch.
NXA-PDC-1100W-PE	1100W AC power supply, port side exhaust pluggable
NXA-PDC-1100W-PI	1100W AC power supply, port side intake pluggable
NXA-PDC-930W-PE	930W AC power supply, port side exhaust pluggable
NXA-PDC-930W-PI	930W AC power supply, port side intake pluggable

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	Description
NXA-PDC-440W-PE	440W DC power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance Note: This power supply is supported only by the Cisco Nexus 9348GC-FXP ACI-mode switch.
NXA-PDC-440W-PI	440W DC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance Note: This power supply is supported only by the Cisco Nexus 9348GC-FXP ACI-mode switch.
NXA-PHV-2KW-PE	Nexus 9000 2KW AC power supply, port-side exhaust Note: This power supply is supported only by the Cisco Nexus 9364C-GX ACI-mode switch.
NXA-PHV-2KW-PI	Nexus 9000 2KW AC power supply, port-side intake Note: This power supply is supported only by the Cisco Nexus 9364C-GX ACI-mode switch.
NXA-PHV-1100W-PE	1100W HVAC/HVDC power supply, port-side exhaust
NXA-PHV-1100W-PI	1100W HVAC/HVDC power supply, port-side intake
NXA-PHV-350W-PE	350W HVAC/HVDC power supply, port-side exhaust
NXA-PHV-350W-PI	350W HVAC/HVDC power supply, port-side intake
N9K-PUV-1200W	1200W HVAC/HVDC dual-direction airflow power supply
N9K-PUV-3000W-B	3000W AC power supply, port side exhaust pluggable
UCSC-PSU-930WDC V01	Port side exhaust DC power supply compatible with all leaf switches
UCS-PSU-6332-DC	930W DC power supply, reversed airflow (port side exhaust)

Table 11. Fixed Leaf Switch Fans

Product ID	Description
N9K-C9300-FAN2	Burgundy port side intake fan
N9K-C9300-FAN2-B	Blue port side exhaust fan
N9K-C9300-FAN3	Burgundy port side intake fan
N9K-C9300-FAN3-B	Blue port side exhaust fan
NXA-FAN-160CFM2-PE	Blue port side exhaust fan
NXA-FAN-160CFM2-PI	Burgundy port side intake fan
NXA-FAN-160CFM-PE	Blue port side exhaust fan
NXA-FAN-160CFM-PI	Burgundy port side intake fan

© 2025 Cisco and/or its affiliates. All rights reserved.

Product ID	Description
NXA-FAN-30CFM-B	Burgundy port side intake fan
NXA-FAN-30CFM-F	Blue port side exhaust fan
NXA-FAN-35CFM-PE	Blue port side exhaust fan
NXA-FAN-35CFM-PI	Burgundy port side intake fan
NXA-FAN-65CFM-PE	Blue port side exhaust fan
NXA-SFAN-65CFM-PE	Blue port side exhaust fan
NXA-FAN-65CFM-PI	Burgundy port side intake fan
NXA-SFAN-65CFM-PI	Burgundy port side intake fan

No Longer Supported Hardware

Starting in the 15.0(1) release, the following hardware is no longer supported:

Product Type	Product ID
Spine switch	N9K-C9336PQ
Modular spine switch line card	N9K-X9736PQ
Modular spine switch fabric module	N9K-C9504-FM N9K-C9508-FM N9K-C9516-FM
Leaf Switch	N9K-C9372PX-E N9K-C9372TX-E N9K-C9332PQ N9K-C9372PX N9K-C9372TX N9K-C9396PX N9K-C9396TX N9K-C93128TX
Expansion Modules	N9K-M12PQ N9K-M6PQ N9K-M6PQ-E

Prior to upgrading your fabric to release 15.0(1) or later, replace these hardware elements in your fabric with other supported hardware. For modular spine switches, replace all unsupported modular line cards

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

and fabric modules because these old generation line cards and fabric modules cannot be operated with newer line cards and fabric modules in the same chassis.

If you attempt to upgrade one of the unsupported hardware to the 15.0(1) release or later, the hardware will unsuccessfully attempt to boot three times, after which the switch will be reverted to the release that was previously installed on it. Therefore, the unsupported hardware will not upgrade to release 15.0(1) or later and the Cisco ACI fabric will operate with inconsistent firmware releases in each switch, which is why we recommend that you replace the unsupported hardware prior to performing the upgrade.

Supported FEX Models

For tables of the FEX models that the Cisco Nexus 9000 Series ACI Mode switches support, see the following webpage:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/hw/interoperability/fexmatrix/fextables.html

For more information on the FEX models, see the Cisco Nexus 2000 Series Fabric Extenders Data Sheet at the following location:

https://www.cisco.com/c/en/us/products/switches/nexus-2000-series-fabric-extenders/datasheet-listing.html

New Hardware Features

There are no new hardware features in this release.

New Software Features

For new software features, see the Cisco Application Policy Infrastructure Controller Release Notes, Release 5.3(2).

Changes in Behavior

For the changes in behavior, see the Cisco ACI Releases Changes in Behavior document.

Open Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug. The " Exists In" column of the table specifies the 15.3(2) releases in which the bug exists. A bug might also exist in releases other than the 15.3(2) releases.

Bug ID	Description	Exists in
CSCwn63197	BUM Traffic on a specific Leaf is not functioning properly. On the Bridge Domain, BUM traffic is transferred via a route called the FTAG Tree, which has a particular spine as its root. In this bug, ISIS process does not register with MRIB process and the link between the spine and the leaf is not included in the FTAG route calculation, where the BUM traffic related with the leaf is impacted.	15.3(2a) and later
CSCwh35818	Remote leaf has coop crash in idle state.	15.3(2a) through

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
		15.3(2e)
CSCwm97890	iStack Core seen after delete/readd of VRF(s).	15.3(2a) through 15.3(2e)
CSCwc25621	Leaf/Spine Switches in ACI fabric showing minor temperature alarms after threshold exceeded.	15.3(2a) through 15.3(2e)
CSCwk44519	N9K-C93180YC-FX3 fabric interfaces are flapping on eth 1/51 and 52.	15.3(2a) through 15.3(2e)
CSCwm70093	On -FX3 platform, link using Cisco-INNOLIGHT branded QSFP type QSFP-100G-SM-SR are experiencing random link flaps.	15.3(2a) through 15.3(2e)
CSCwn00902	rx_amp tunning change to 11 from 33 for CISCO-INNOLIGHT QSFP-100G-SR4 Program the rx_amp(One of the transceiver parameters) to 0x11 for 100G non-adaptive optics. 0x11 is the expected value as per the Signal integrity team so we are going ahead and programming this rx_amp field to 0x11	15.3(2a) through 15.3(2e)
CSCwh46885	When the fabric nodes are using ACI release 15.2(7g), the N9K-C9348 switch fails during the POAP DHCP discover phase.	15.3(2a) through 15.3(2e)
CSCwf01300	Bloggerd or any other process crash.	15.3(2c)
CSCwf50879	ARP reply packet dropped on ingress TOR due to ACL_DROP when it ingresseson GX2 boxes via any of the two special inband management ports.	15.3(2c)
CSCwh14873	The memory usage of PLTM process continues to increase, especially, leak speed increases when OSPF session flaps frequently occur (ex. 2k OSPF sessions flap in 20 loops). This leak can be seen in libcore in its library, but core file of PLTM is needed to confirm whether memory leak occurs in libcore.	15.3(2c)
CSCwi59385	Flow: Pod1---ND1-----9332D-GX2B--------ND2---- Pod2 Selective flow drop of flow between pod1- pod2 due to ACI leaf. Many pods stuck in crashloopbackoff due to communication break.	15.3(2c)
CSCwi75324	Multisite EVPN routes stopped being advertised to remote sites after GOLF label was removed from the tenant.	15.3(2c)
CSCwe36335	Upon Active Sup kernel reboot, the switchover to Standby (now Active following switchover) may lead to NFM HAP reset/crash. This results in both Supervisors reloading affecting chassis Supervisor redundancy.	15.3(2a) through 15.3(2c)
CSCwf51388	PIM crashed after deleting SSM range (224/4) from GUI and running " show ip pim route x.x.x.x vrf all" after that.	15.3(2a) through 15.3(2c)

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
CSCwi32475	The memory usage of PLTM process continues to increase.Especially, leak speed increases when OSPF session flaps frequently occur (ex. 2k OSPF sessions flap in 20 loops).This leak can be seen in libcore in its library, but core file of PLTM is needed to confirm whether memory leak occurs in libcore.	15.3(2a) through 15.3(2c)
CSCwi71543	The link is not coming up when using the SFP 'SFP-10/25G-CSR-S' to Dell servers (XC7525-24 CORE), Broadcom adapter " BCM57414 NetXtreme-E 10/25Gbps " and some other SFPs.	15.3(2a) through 15.3(2c)
CSCwj32102	ARP resolution failed on hosts/VMs due to COPP drops on leafs with FIE configuration.	15.3(2a) through 15.3(2c)
CSCwj40433	ARP resolution failed on hosts/VMs due to COPP drops on leafs with FIE configuration.	15.3(2a) through 15.3(2c)
CSCwj78565	ACI fabric modules N9K-C9508-FM-E2 and fixed spine switches N9K-C9364C and N9K-C9332C may report HAL process crash and reboot when a specific internal cli (show platform internal bky table roc_bky_xbc_acc) was issued.	15.3(2a) through 15.3(2c)
CSCwk36249	An EPG is configured with shared services and microsegmentation and the microsegmented traffic is misclassified with the pctag of the parent EPG.	15.3(2a) through 15.3(2c)
CSCwk41796	All the traffic flows work on the destination leaf switch except an unknown unicast packet (flood) received on the working vPC switch. An ELAM capture for the flood packet shows that the vPC_df bit set as 0x0 instead of 0x1, which dropped the flood traffic.	15.3(2a) through 15.3(2c)
CSCwk87152	PBR Device marked down by health monitoring.	15.3(2a) through 15.3(2c)
CSCvg85886	When an ARP request is generated from one endpoint to another endpoint in an isolated EPG, an ARP glean request is generated for the first endpoint.	15.3(2a) and later
CSCvw89840	Traffic originating from a vPC TEP is dropped for Layer 2 multicast and unknown unicast traffic when pod redundancy is triggered.	15.3(2a) and later
CSCvx31008	When a Cisco ACI Multi-Pod infra B2B OSPF link goes down, any faults for the multiPodDirect instance that would normally be raised will not be raised. Also, the operational state for the multiPodDirect instance will not be updated in the DME database.	15.3(2a) and later
CSCvy31805	The PBR destination group for bypass action is not properly programmed with PBR service graph for service devices behind l3out and with " bypass" action enabled to redirect to another service node in the graph.Now on bypass switchover, the traffic doesn't get redirected to the next service node in the chain.	15.3(2a) and later
CSCvy94715	PTP is not supported on breakout ports.	15.3(2a) and later

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
CSCwa45189	A static MAC address configuration fails if the same endpoint is already learned as a dynamic MAC address in the fabric with VIP addresses attached to it. Static MAC address deployment from the Cisco APIC fails and the operational state of the static MAC address managed object on the leaf switch shows as " down." The static MAC address configuration will be missing in the EPM/EPMC database. After deploying a static MAC address configuration on the Cisco APIC, the following command does not show the static MAC: show system internal epm endpoint mac <MAC>	15.3(2a) and later
CSCwc56728	Layer 2 unicast traffic drop is seen for 15 minutes as the destination MAC goes rogue on the source vPC switch pair when the destination remote leaf switch vPC pair is upgraded. This issue is seen with rogue enabled with a count of 4 moves in 60 seconds.	15.3(2a) and later
CSCwd21973	An L3Out interface is not deployed after a certain configuration sequence.	15.3(2a) and later
CSCwd34452	When the leaf switches of a vPC pair are upgraded with the 'Graceful Upgrade' mode, unicast traffic to the connected devices will have 5+ minutes traffic loss.	15.3(2a) and later
CSCwe69786	A switch does not update the exporter configuration when a user deletes or removes VRF instance and EPG information from the exporter.	15.3(2a) and later
CSCwe85338	This issue is seen when there is an overlap between an endpoint IP address that has been marked rogue due to excessive moves. While the endpoint is marked rogue, if a user configures that IP address as a Layer 4 to Layer 7 services vIP address, there is an impact to the DSR functionality.	15.3(2a) and later
CSCwe90254	When a TechSupport file for a 9500 chassis with FM-E2 fabric modules is collected, the CLI commands needed to be run for the NX-OS TechSupport are wrong.	15.3(2a) and later
CSCwf15461	Whenever an EPG is configured with multiple physical domains using overlapping VLAN pools, following some configuration such as deleting/adding a pool to a domain or adding/removing a domain in an EPG, the fabric encapsulation (VXLAN ID used for VLAN encapsulation) might be mismatched.F3274 will be seen and traffic to a vPC server might be impacted in that EPG.	15.3(2a) and later
CSCwf38424	100G-CR4 copper links with either 3 meter or 5 meter cables do not come up when a Cisco N9K-C9332D-GX2B switch is connected to a Mellanox NIC and auto-negotiation is enabled on both.	15.3(2a) and later
CSCwf57396	The 30 second input rate and 30 second output rate show values beyond 30 seconds for an interface that is disabled	15.3(2a) and later
CSCwf58246	In the case of large network instability with a lot of flaps, the APIC may disable hardware learning and disable COOP to endpoint notification on a leaf switch. This can lead to a COOP entry on a spine switch pointing to a " wrong" location. This is a very rare scenario.	15.3(2a) and later
CSCwf88948	After a system controller switchover, there is no ping/ssh response from the spine switch in-band management for several minutes. It seems that there is an issue with path between SUP and linecard.	15.3(2a) and later
CSCwh03684	HAL has high CPU utilization.	15.3(2a) and later
CSCwh07391	Traffic coming from ISN or IPN may get misclassified as iTraceroute or will not	15.3(2a) and

© 2025 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Exists in
	preserve cos correctly.On any FM that was reloaded, dot1p preserve may have not been set correctly post reload.	later
CSCwh15691	fvL3EpDef is not removed after adjacency gets updated.	15.3(2a) and later
CSCwp13198	There is a delay in bringing up DAC cable links in N9K-93180YC-FX3 when connected to Mellanox NIC cards with AN-OFF. This is observed when there is a link flap from either end of the link.	15.3(2a) and later
CSCwh21417	A switch's power supply is functioning properly even though the following error message persists:LOG_LOCAL0-2-SYSTEM_MSG [E4204936][transition][critical][sys] %PLATFORM-2-PS_UNSUPPORTED: Detected an unsupported power supply 2 Unknown for CISCO Multilayer Switch (Serial number LIT233023Z5)	15.3(2a) and later
CSCwh29782	A Cisco Nexus 9000 switch in the ACI-mode cannot negotiate on a 1G link due to the " Remote Fault seen" error.	15.3(2a) and later
CSCwh48737	Bounce entry for an endpoint may point to wrong TEP address, leading to connectivity failures.	15.3(2a) and later
CSCwh60203	There is unexpected behavior with the DHCP relay when using DHCP relay with the " DHCP server preference" feature. The issue is triggered by deleting one of the bridge domains that use the DHCP relay label. This causes the Cisco APIC to remove the DHCP server's SVI from all switches involved in the change, which means that all other bridge domains on those switches can no longer do DHCP relay.	15.3(2a) and later
CSCwh64732	The hardware is DOM-capable (Y), but DOM information is not showing up for the command: " show int ethernet 1/X transceiver details"	15.3(2a) and later
CSCwh67412	ACI displays fan speed percentage incorrectly.	15.3(2a) and later
CSCwh75559	tcpdump on the tahoe0 interface randomly has incorrect time stamps (sometimes an old time stamp).	15.3(2a) and later
CSCwh76977	The device reloaded unexpectedly because of " sdkhal hap reset" after the " show platform internal hal l3 intfdb" command was executed in command-line interface in vsh_lc mode.	15.3(2a) and later
CSCwh76996	While inserting or reloading a leaf switch, its vPC peer will try to bring up the vPC when the peer IP is 0.0.0.0.	15.3(2a) and later
CSCwh78987	Breakout ports configured as port channel members are no longer part of the port channel post clean reload.	15.3(2a) and later
CSCwh92659	An endpoint may become out of sync between spine switches in different pods. Each spine switch may point to a local pod TEP as the tunnel next hop for the leaf switches. The issue does not get cleared until the incorrect pod spine switches age/delete their COOP entry.	15.3(2a) and later
CSCwi01017	This bug is only applicable when scale of the switch config is high. Repeated tech-support collection causes the log files to be completely deleted from under /mnt/ifc/log/.	15.3(2a) and later
CSCwi87437	MFDM MTS buffer is stuck, causing the MFDM process to not process updates coming from other processes. This can impact broadcast programming on the downstream component, which causes a BUM issue.	15.3(2a) and later

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
CSCwj19730	A port gets stuck on an old VLAN.	15.3(2a) and later
CSCwj23667	- Enabling MACsec between ACI and Nexus devices. - The MACsec session needs to be established. - CDP neighbors are not available for the interface, but the LLDP neighbors are correct.	15.3(2a) and later
CSCwj51461	A host has connectivity loss when the destination host fail overed using UCS Fabric Failover. Issue happens when PVLAN and Micro segmentation are enabled on EPG with Proxy-ARP flag.	15.3(2a) and later
CSCwj67989	When configuring a /32 subnet for " IP Data-plane Learning disabled" under an EPG, the subnet is deployed to a leaf switch with the EPG as well as the EPGs' consumed source EPG by the contract. As a result, it adds a remote endpoint (XR EP) to the hardware and causes any received iVXLAN packet with an inner destination to hit the /32 endpoint, which gets dropped by the loop prevention ACL. This /32 endpoint may become stale even after the /32 EPG subnet has been deleted from the Cisco APIC.	15.3(2a) and later
CSCwk39458	A leaf switch crashes due fcpc hap reset.	15.3(2a) and later
CSCwk47025	When querying for OID 1.3.6.1.2.1.1.2.0 for N9K-C9332D-GX2B, the incorrect value is returned by the switch if role is spine: - If role is spine: it will return 1.3.6.1.4.1.9.12.3.1.3.1570 (Nexus 9336PQ) - If role is leaf: it will return 1.3.6.1.4.1.9.12.3.1.3.2403 (Nexus 9332D-GX2B)	15.3(2a) and later
CSCwf80642	The mosquitto process' memory usage consistently increases.	15.3(2a) and 15.3(2b)
CSCwh21375	When an SNMP GET is sent with the OID " iso.3.6.1.2.1.1.2.0" on a leaf/spine switch, the leaf/spine switch responds with the faulty value 1.3.6.1.4.1.9.12.3.1.3.1570.	15.3(2a) and 15.3(2b)
CSCwi80642	After initiating a leaf switch reload, a traffic blackhole occurs for a brief duration.	15.3(2a) and 15.3(2b)
CSCwi89535	A leaf switch crashes due to the following reason: Reason: Reset triggered due to HA policy of Reset Service: sdkhal hap reset Version : xxx	15.3(2a) and 15.3(2b)
CSCwj30440	The supervisor of a leaf switch silently reloads.	15.3(2a) and 15.3(2b)
CSCwj51767	There is a kernel panic resulting in a switch reset or crash: Reason: Kernel Panic System version: 15.2(8e) Service: system crash	15.3(2a) and 15.3(2b)
CSCwk37834	A leaf switch will crash due to aclqos hap reset.	15.3(2a) and

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
		15.3(2b)
CSCwa34808	Non-related flows that should not be captured are cached and exported, including infra, inb mgmt and user tenant flows that are not ingressing from the flow interface.	15.3(2a)
CSCwb46927	* DHCP Relay configuration is correct. * It works for some clients, but not for others. * The server packet capture shows that it sent the offer packet. * The client packet capture shows that it did not send the offer packet. * SPAN on the port where the traffic flow of the server enters the network shows the correct information for option 82 of DHCP.	15.3(2a)
CSCwb81672	* DHCP Relay configuration is correct. * It works for some clients, but not for others. * The server packet capture shows that it sent the offer packet. * The client packet capture shows that it did not send the offer packet. * SPAN on the port where the traffic flow of the server enters the network shows the correct information for option 82 of DHCP.	15.3(2a)
CSCwf25837	Transit BFD packets are dropped in the egress direction of a Layer 2 bridge domain if the egress interface is also forwarding the VLAN that is used.	15.3(2a)
CSCwf53105	" vsh" process generates multiple core files on switches after starting OnDemand Techsupport collection for leaf switches.	15.3(2a)
CSCwh77467	SDKHAL crashes are seen on ACI spine N9K-C9364D-GX2A switch running 15.2(7g). A core file is also generated.	15.3(2a)
CSCwh77567	If there is a single endpoint move local-to-remote or remote-to-local within the detection interval, then the move count is not reset. This may lead to an endpoint being incorrectly marked as rogue on the node.	15.3(2a)
CSCwi18214	A leaf switch repeatedly reloads due to policyelem abnormal exit and HAP reset.	15.3(2a)
CSCwi42105	QSFP-100G-SM-SR links do not come up on -GX2A and -GX2B leaf switch re-timer ports.	15.3(2a)
CSCwt09384	SDK HAL core, leaf crash during Ts collection on GX2, H2, H1 model leaf.	15.3(2a) and later

Resolved Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug. The " Fixed In" column of the table specifies whether the bug was resolved in the base release or a patch release.

Bug ID	Description	Fixed in
CSCwh35818	Remote leaf has coop crash in idle state.	15.3(2f)
CSCwm97890	iStack Core seen after delete/readd of VRF(s).	15.3(2f)

© 2025 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Fixed in
CSCwc25621	Leaf/Spine Switches in ACI fabric showing minor temperature alarms after threshold exceeded.	15.3(2f)
CSCwk44519	N9K-C93180YC-FX3 fabric interfaces are flapping on eth 1/51 and 52.	15.3(2f)
CSCwm70093	On -FX3 platform, link using Cisco-INNOLIGHT branded QSFP type QSFP-100G-SM-SR are experiencing random link flaps.	15.3(2f)
CSCwn00902	rx_amp tunning change to 11 from 33 for CISCO-INNOLIGHT QSFP-100G-SR4 Program the rx_amp(One of the transceiver parameters) to 0x11 for 100G non-adaptive optics. 0x11 is the expected value as per the Signal integrity team so we are going ahead and programming this rx_amp field to 0x11	15.3(2f)
CSCwh46885	When the fabric nodes are using ACI release 15.2(7g), the N9K-C9348 switch fails during the POAP DHCP discover phase.	15.3(2f)
CSCwe36335	Upon Active Sup kernel reboot, the switchover to Standby (now Active following switchover) may lead to NFM HAP reset/crash. This results in both the Supervisors reloading, affecting chassis Supervisor redundancy.	15.3(2d)
CSCwf01300	Bloggerd or any other process crash.	15.3(2d)
CSCwf50879	ARP reply packet dropped on ingress TOR due to ACL_DROP when it ingresses on GX2 boxes via any of the two special inband management ports.	15.3(2d)
CSCwf51388	PIM crashed after deleting SSM range (224/4) from GUI and running " show ip pim route x.x.x.x vrf all" after that.	15.3(2d)
CSCwh14873	The memory usage of PLTM process continues to increase, especially, leak speed increases when OSPF session flaps frequently occur (ex. 2k OSPF sessions flap in 20 loops).This leak can be seen in libcore in its library, but core file of PLTM is needed to confirm whether memory leak occurs in libcore.	15.3(2d)
CSCwi32475	The memory usage of PLTM process continues to increase, especially, leak speed increases when OSPF session flaps frequently occur (ex. 2k OSPF sessions flap in 20 loops). This leak can be seen in libcore in its library, but core file of PLTM is needed to confirm whether memory leak occurs in libcore.	15.3(2d)
CSCwi59385	Flow: Pod1---ND1-----9332D-GX2B--------ND2---- Pod2 Selective flow drop of flow between pod1- pod2 due to ACI leaf. Many pods stuck in crashloopbackoff due to communication break.	15.3(2d)
CSCwi71543	The link is not coming up when using the SFP 'SFP-10/25G-CSR-S' to Dell servers (XC7525-24 CORE), Broadcom adapter " BCM57414 NetXtreme-E 10/25Gbps " and some other SFPs.	15.3(2d)
CSCwi75324	Multisite EVPN routes stopped being advertised to remote sites after GOLF label was removed from the tenant.	15.3(2d)

© 2025 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Fixed in
CSCwj32102	ARP resolution failed on hosts/VMs due to COPP drops on leafs with FIE configuration.	15.3(2d)
CSCwj40433	ARP resolution failed on hosts/VMs due to COPP drops on leafs with FIE configuration.	15.3(2d)
CSCwj78565	ACI fabric modules N9K-C9508-FM-E2 and fixed spine switches N9K-C9364C and N9K-C9332C may report HAL process crash and reboot when a specific internal cli (show platform internal bky table roc_bky_xbc_acc) was issued.	15.3(2d)
CSCwk36249	An EPG is configured with shared services and microsegmentation and the microsegmented traffic is misclassified with the pctag of the parent EPG.	15.3(2d)
CSCwk41796	All the traffic flows work on the destination leaf switch except an unknown unicast packet (flood) received on the working vPC switch. An ELAM capture for the flood packet shows that the vPC_df bit set as 0x0 instead of 0x1, which dropped the flood traffic.	15.3(2d)
CSCwk87152	PBR device marked down by health monitoring.	15.3(2d)
CSCwf80642	The mosquitto process' memory usage consistently increases.	15.3(2c)
CSCwh21375	When an SNMP GET is sent with the OID " iso.3.6.1.2.1.1.2.0" on a leaf/spine switch, the leaf/spine switch responds with the faulty value 1.3.6.1.4.1.9.12.3.1.3.1570.	15.3(2c)
CSCwi80642	After initiating a leaf switch reload, a traffic blackhole occurs for a brief duration.	15.3(2c)
CSCwi89535	A leaf switch crashes due to the following reason: Reason: Reset triggered due to HA policy of Reset Service: sdkhal hap reset Version : xxx	15.3(2c)
CSCwj30440	The supervisor of a leaf switch silently reloads.	15.3(2c)
CSCwj51767	There is a kernel panic resulting in a switch reset or crash: Reason: Kernel Panic System version: 15.2(8e) Service: system crash	15.3(2c)
CSCwk37834	A leaf switch will crash due to aclqos hap reset.	15.3(2c)
CSCwa34808	Non-related flows that should not be captured are cached and exported, including infra, inb mgmt and user tenant flows that are not ingressing from the flow interface.	15.3(2b)

© 2025 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Fixed in
CSCwb46927	* DHCP Relay configuration is correct. * It works for some clients, but not for others. * The server packet capture shows that it sent the offer packet. * The client packet capture shows that it did not send the offer packet. * SPAN on the port where the traffic flow of the server enters the network shows the correct information for option 82 of DHCP.	15.3(2b)
CSCwb81672	* DHCP Relay configuration is correct. * It works for some clients, but not for others. * The server packet capture shows that it sent the offer packet. * The client packet capture shows that it did not send the offer packet. * SPAN on the port where the traffic flow of the server enters the network shows the correct information for option 82 of DHCP.	15.3(2b)
CSCwf25837	Transit BFD packets are dropped in the egress direction of a Layer 2 bridge domain if the egress interface is also forwarding the VLAN that is used.	15.3(2b)
CSCwf53105	" vsh" process generates multiple core files on switches after starting OnDemand Techsupport collection for leaf switches.	15.3(2b)
CSCwh77467	SDKHAL crashes are seen on ACI spine N9K-C9364D-GX2A switch running 15.2(7g). A core file is also generated.	15.3(2b)
CSCwh77567	If there is a single endpoint move local-to-remote or remote-to-local within the detection interval, then the move count is not reset. This may lead to an endpoint being incorrectly marked as rogue on the node.	15.3(2b)
CSCwi18214	A leaf switch repeatedly reloads due to policyelem abnormal exit and HAP reset.	15.3(2b)
CSCwi42105	QSFP-100G-SM-SR links do not come up on -GX2A and -GX2B leaf switch re-timer ports.	15.3(2b)
CSCwf43672	After IPv6 configuration change to bridge domain or L3Out, some IPv6 addresses are unreachable.	15.3(2a)
CSCwf49076	The APIC connectivity file is removed after a port profile stateless reload is done, which causes the front panel ports come UP before all configurations are completed. This causes packet drops for more than 10 minutes.	15.3(2a)
CSCwh13845	Post a spine switch stateful reload/upgrade, you may observe traffic drops in the ACI fabric. If there are remote leaf switches, you may see remote leaf switches not forming a tunnel with the spine switches. On upgraded/stateful reloaded spine switches, check the output of " show coop internal info global" . If you observe the " Local Adjacency" as " Citizen" , then you have hit this defect. Example: System info:Local Adjacency (Citizen).	15.3(2a)
CSCwh72876	The EPM process crashed when there was no disk space was available at /var/sysmgr/tmp_logs/.	15.3(2a)
CSCwh73346	After removing SG association from a shared L3Out contract, traffic will be dropped on a border leaf switch.	15.3(2a)

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Fixed in
CSCwh73782	Traffic that is forwarded by a spine switch toward a leaf switch is dropped by one of the spine switch's fabric modules.On this fabric module where packets are dropped, the TEP of the destination leaf switch is not programmed in FIB and HAL.	15.3(2a)
CSCwh79632	Uplink ports flap frequently on leaf switch. Eth1/51 may be observed to flap more frequently than other ports.	15.3(2a)
CSCwh81430	After being upgraded, a Cisco 93108TC-FX3P leaf switch reboots and random RJ45 ports remain link down after the switch reloads.	15.3(2a)
CSCwh84746	An endpoint moves between POD1 and POD2. On POD3, the endpoint changes to the FREEZE state, but on POD1 and POD2 the endpoint does not change to the FREEZE state after it stops moving.	15.3(2a)
CSCwh91351	With FX3 switches in a vPC pair and the endpoints in different VRF instances, traffic gets dropped from the second switch in the pair.	15.3(2a)
CSCwi05613	This issue can be seen when BGP L3Outs have import or export route maps configured using match statements based on regular expressions (for example matching of regex: (65[2-3]01:102..). If BGP regex communities are configured, both the deny and permit statements are not always honored and there are unexpected results. From the BGP and routing table perspective, routes come in with an extended community value that are being permitted incorrectly or denied incorrectly based on the ACI route map that is configured.	15.3(2a)
CSCwi21299	When there is an FCOE interface flap or speed change due to inserting an SFP, this affects the dataplane of other FCOE interfaces that share the same MAC address.	15.3(2a)
CSCwi31656	SPAN traffic does not go out from the destination SPAN port after the peer interface flaps. MAC credit goes to zero for the SPAN destination port after the peer interface flaps. You also might see the native interface that is part of same MAC address hardware in which the SPAN destination port is configured stop sending control plane packets because the CPU buffer is exhausted by the SPAN destination port.	15.3(2a)

Known Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug. The " Exists In" column of the table specifies the 15.3(2) releases in which the bug exists. A bug might also exist in releases other than the 15.3(2) releases.

Bug ID	Description	Exists in
CSCuo37016	When configuring the output span on a FEX Hif interface, all the layer 3 switched packets going out of that FEX Hif interface are not spanned. Only layer 2 switched packets going out of that FEX Hif are spanned.	15.3(2a) and later
CSCuo50533	When output SPAN is enabled on a port where the filter is VLAN, multicast traffic in the VLAN that goes out of that port is not spanned.	15.3(2a) and later
CSCup65586	The **show interface** command shows the tunnel's Rx/Tx counters as 0.	15.3(2a) and later

© 2025 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Exists in
CSCup82908	The **show vpc brief** command displays the wire-encap VLAN Ids and the **show interface .. trunk** command displays the internal/hardware VLAN IDs. Both VLAN IDs are allocated and used differently, so there is no correlation between them.	15.3(2a) and later
CSCup92534	Continuous " threshold exceeded" messages are generated from the fabric.	15.3(2a) and later
CSCuq39829	Switch rescue user (" admin") can log into fabric switches even when TACACS is selected as the default login realm.	15.3(2a) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	15.3(2a) and later
CSCuq77095	When the command **show ip ospf vrf <vrf_name>** is run from bash on the border leaf switch, the checksum field in the output always shows a zero value.	15.3(2a) and later
CSCuq83910	When an IP address moves from one MAC behind one leaf switch to another MAC behind another leaf switch, even though the VM sends a GARP packet, in ARP unicast mode, this GARP packet is not flooded. As a result, any other host with the original MAC to IP binding sending an L2 packet will send to the original leaf switch where the IP was in the beginning (based on MAC lookup), and the packet will be sent out on the old port (location). Without flooding the GARP packet in the network, all hosts will not update the MAC-to-IP binding.	15.3(2a) and later
CSCuq92447	When modifying the L2Unknown Unicast parameter on a Bridge Domain (BD), interfaces on externally connected devices may bounce. Additionally, the endpoint cache for the BD is flushed and all endpoints will have to be re-learned.	15.3(2a) and later
CSCuq93389	If an endpoint has multiple IPs, the endpoint will not be aged until all IPs go silent. If one of the IP addresses is reassigned to another server/host, the fabric detects it as an IP address move and forwarding will work as expected.	15.3(2a) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	15.3(2a) and later
CSCur81822	The access-port operational status is always " trunk" .	15.3(2a) and later
CSCus18541	An MSTP topology change notification (TCN) on a flood domain (FD) VLAN may not flush endpoints learned as remote where the FD is not deployed.	15.3(2a) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	15.3(2a) and later
CSCus43167	Any TCAM that is full, or nearly full, will raise the usage threshold fault. Because the faults for all TCAMs on leaf switches are grouped together, the fault will appear even on those with low usage. Workaround: Review the leaf switch scale and reduce the TCAM usage. Contact TAC to isolate further which TCAM is full.	15.3(2a) and later
CSCus54135	The default route is not leaked by BGP when the scope is set to context. The scope should be set to **Outside** for default route leaking.	15.3(2a) and later

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
CSCus61748	If the leaf switch 1RU system is configured with the RED fan (the reverse airflow), the air will flow from front to back. The temperature sensor in the back will be defined as an inlet temperature sensor, and the temperature sensor in the front will be defined as an outlet temperature sensor. If the leaf switch 1RU system is configured with the BLUE fan (normal airflow), the air will flow from back to front. The temperature sensor in the front will be defined as an inlet temperature sensor, and the temperature sensor in the back will be defined as outlet temperature sensor. From the airflow perspective, the inlet sensor reading should always be less than the outlet sensor reading. However, in the leaf switch 1RU family, the front panel temperature sensor has some inaccurate readings due to the front panel utilization and configuration, which causes the inlet temperature sensor reading to be very close, equal, or even greater than the outlet temperature reading.	15.3(2a) and later
CSCut59020	If Backbone and NSSA areas are on the same leaf switch, and default route leak is enabled, Type-5 LSAs cannot be redistributed to the Backbone area.	15.3(2a) and later
CSCuu11347	Traffic from the orphan port to the vPC pair is not recorded against the tunnel stats. Traffic from the vPC pair to the orphan port is recorded against the tunnel stats.	15.3(2a) and later
CSCuu11351	Traffic from the orphan port to the vPC pair is only updated on the destination node, so the traffic count shows as excess.	15.3(2a) and later
CSCuu66310	If a bridge domain "Multi Destination Flood" mode is configured as "Drop", the ISIS PDU from the tenant space will get dropped in the fabric.	15.3(2a) and later
CSCuv57302	Atomic counters on the border leaf switch do not increment for traffic from an endpoint group going to the Layer 3 out interface.	15.3(2a) and later
CSCuv57315	Atomic counters on the border leaf switch do not increment for traffic from the Layer 3 out interface to an internal remote endpoint group.	15.3(2a) and later
CSCuv57316	TEP counters from the border leaf switch to remote leaf switch nodes do not increment.	15.3(2a) and later
CSCuw09389	For direct server return operations, if the client is behind the Layer 3 out, the server-to-client response will not be forwarded through the fabric.	15.3(2a) and later
CSCux97329	With the common pervasive gateway, only the packet destination to the virtual MAC is being properly Layer 3 forwarded. The packet destination to the bridge domain custom MAC fails to be forwarded. This is causing issues with certain appliances that rely on the incoming packets' source MAC to set the return packet destination MAC.	15.3(2a) and later
CSCuy00084	BCM does not have a stats option for yellow packets/bytes, and so BCM does not show in the switch or APIC GUI stats/observer.	15.3(2a) and later
CSCuy02543	Bidirectional Forwarding Detection (BFD) echo mode is not supported on IPv6 BFD sessions carrying link-local as the source and destination IP address. BFD echo mode also is not supported on IPv4 BFD sessions over multihop or VPC peer links.	15.3(2a) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	15.3(2a) and later
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	15.3(2a) and later

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
CSCuy25780	An overlapping or duplicate prefix/subnet could cause the valid prefixes not to be installed because of batching behavior on a switch. This can happen during an upgrade to the 1.2(2) release.	15.3(2a) and later
CSCuy47634	EPG statistics only count total bytes and packets. The breakdown of statistics into multicast/unicast/broadcast is not available on new hardware.	15.3(2a) and later
CSCuy56975	You must configure different router MACs for SVI on each border leaf switch if L3out is deployed over port-channels/ports with STP and OSPF/OSPFv3/eBGP protocols are used. There is no need to configure different router MACs if you use VPC.	15.3(2a) and later
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to " 0" , and so traffic will still flow.	15.3(2a) and later
CSCuy96912	The debounce timer is not supported on 25G links.	15.3(2a) and later
CSCuz13529	With the N9K-C93180YC-EX switch, drop packets, such as MTU or storm control drops, are not accounted for in the input rate calculation.	15.3(2a) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	15.3(2a) and later
CSCuz13810	When subnet check is enabled, a leaf switch does not learn IP addresses locally that are outside of the bridge domain subnets. However, the packet itself is not dropped and will be forwarded to the fabric. This will result in such IP addresses getting learned as remote endpoints on other leaf switches.	15.3(2a) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	15.3(2a) and later
CSCva98767	The front port of the QSA and GLC-T 1G module has a 10 to 15-second delay as it comes up from the insertion process.	15.3(2a) and later
CSCvb36823	If you have only one spine switch that is part of the infra WAN and you reload that switch, there can be drops in traffic. You should deploy the infra WAN on more than one spine switch to avoid this issue.	15.3(2a) and later
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	15.3(2a) and later
CSCvb49451	In the case of endpoints in two different leaf switch pairs across a spine switch that are trying to communicate, an endpoint does not get relearned after being deleted on the local leaf switch pair. However, the endpoint still has its entries on the remote leaf switch pair.	15.3(2a) and later
CSCvd11146	Bridge domain subnet routes advertised out of the Cisco ACI fabric through an OSPF L3Out can be relearned in another node belonging to another OSPF L3Out on a different area.	15.3(2a) and later
CSCvd63567	After upgrading a switch, Layer 2 multicast traffic flowing across PODs gets affected for some of the bridge domain Global IP Outsides.	15.3(2a) and later
CSCvh18100	If Cisco ACI Virtual Edge or AVS is operating in VxLAN non-switching mode behind a FEX, the traffic between endpoints in the same EPG will fail when the bridge domain has ARP flooding enabled.	15.3(2a) and later

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
CSCvn94400	There is a traffic blackhole that lasts anywhere from a few seconds to a few mins after a border leaf switch is restored.	15.3(2a) and later
CSCvp04772	During an upgrade on a dual-SUP system, the standby SUP may go into a failed state.	15.3(2a) and later
CSCvq71034	There is a policy drop that occurs with L3Out transit cases.	15.3(2a) and later
CSCvr12912	A switch reloads due to a sysmgr heartbeat failure and sysmgr HAP reset.	15.3(2a) and later
CSCvr61096	In a port group that has ports of mixed speeds, the first port in the port group that has valid optics present and is not in the admin down state is processed. The ports that come up later are brought up if they are using the same speed; otherwise, they are put in the hw-disabled state. For example, if ports 14 and 15 are up and are using the 100G speed, then if ports 13 and 16 are using the 40G speed, these ports will be put in the hw-disabled state. After reloading or upgrading, you might not have the same interfaces in the port group in the UP state and in the hw-disabled state as you did before the reload or upgrade.	15.3(2a) and later
CSCvs75598	A link flap policy can be configured on Cisco switches that were released prior to the -EX switches. However, this feature does not work with such switches and no fault is raised in the Cisco APIC.	15.3(2a) and later
CSCvt61851	When MPLS VRF stats (egress) is compared with Layer 2 interface egress stats, we can find that the packet count matches for both while there could be a discrepancy with the bytes count.	15.3(2a) and later
CSCvu02371	The DEI value in a Layer 2 header of spanned Tx packets from an MPLS interface might not have the same value as the actual data path packet.	15.3(2a) and later
CSCvu42069	The event log shows VTEP tunnel down and up events. The down time and up time are the same, and there is no fault message.	15.3(2a) and later
CSCvx62362	When a service device is connected behind an L3Out in 2-arm mode with both legs on the same leaf switch, tracking packets get dropped.	15.3(2a) and later
CSCvx67074	Redirected traffic to a service device behind an L3Out is dropped by the service device due to the lack of proper routing information.	15.3(2a) and later
CSCvx75128	Redirected traffic coming back from a service device behind an L3Out is dropped on the second leg.	15.3(2a) and later
CSCvy06135	The leaf switch techsupport with a specified time range fails when the space "/mnt/ifc/log" gets filled up by more than 80%.	15.3(2a) and later
CSCvz84284	Upon deletion of a VRF instance that has a micro-BFD port channel in the "up" state, all the member ports of the port channel that were in the "up" state prior to the VRF instance deletion go to the "down" state. The micro-BFD port channels never transition back to the "up" state.	15.3(2a) and later
CSCwa78857	Cisco APIC allows you to configure any number of DHCP relay addresses. However, the maximum number of relay address that can be supported is 16 from a switch. If a 17th DHCP provider is added to the DHCP label, it will not be used even if one of first 16 DHCP providers is removed.	15.3(2a) and later

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
N/A	Load balancers and servers must be Layer 2 adjacent. Layer 3 direct server return is not supported. If a load balancer and servers are Layer 3 adjacent, then they have to be placed behind the Layer 3 out, which works without a specific direct server return virtual IP address configuration.	15.3(2a) and later
N/A	IPN should preserve the CoS and DSCP values of a packet that enters IPN from the ACI spine switches. If there is a default policy on these nodes that change the CoS value based on the DSCP value or by any other mechanism, you must apply a policy to prevent the CoS value from being changed. At the minimum, the remarked CoS value should not be 4, 5, 6, or 7. If CoS is changed in the IPN, you must configure a DSCP-CoS translation policy in the APIC for the pod that translates queuing class information of the packet into the DSCP value in the outer header of the iVXLAN packet. You can also embed CoS by enabling CoS preservation. For more information, see the Cisco APIC and QoS KB article.	15.3(2a) and later
N/A	The following properties within a QoS class under " Global QoS Class policies" should not be changed from their default value and is only used for debugging purposes: MTU (default – 9216 bytes) Queue Control Method (default – Dynamic) Queue Limit (default – 1522 bytes) Minimum Buffers (default – 0)	15.3(2a) and later
N/A	The modular chassis Cisco ACI spine nodes, such as the Cisco Nexus 9508, support warm (stateless) standby where the state is not synched between the active and the standby supervisor modules. For an online insertion and removal (OIR) or reload of the active supervisor module, the standby supervisor module becomes active, but all modules in the switch are reset because the switchover is stateless. In the output of the **show system redundancy status** command, warm standby indicates stateless mode.	15.3(2a) and later
N/A	When a recommissioned APIC controller rejoins the cluster, GUI and CLI commands can time out while the cluster expands to include the recommissioned APIC controller.	15.3(2a) and later
N/A	If connectivity to the APIC cluster is lost while a switch is being decommissioned, the decommissioned switch may not complete a clean reboot. In this case, the fabric administrator should manually complete a clean reboot of the decommissioned switch.	15.3(2a) and later
N/A	Before expanding the APIC cluster with a recommissioned controller, remove any decommissioned switches from the fabric by powering down and disconnecting them. Doing so will ensure that the recommissioned APIC controller will not attempt to discover and recommission the switch.	15.3(2a) and later
N/A	Multicast router functionality is not supported when IGMP queries are received with VxLAN encapsulation.	15.3(2a) and later
N/A	IGMP Querier election across multiple Endpoint Groups (EPGs) or Layer 2 outsides (External Bridged Network) in a given bridge domain is not supported. Only one EPG or Layer 2 outside for a given bridge domain should be extended to multiple multicast routers if any.	15.3(2a) and later
N/A	The rate of the number of IGMP reports sent to a leaf switch should be limited to 1000 reports per second.	15.3(2a) and later

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
N/A	Unknown IP multicast packets are flooded on ingress leaf switches and border leaf switches, unless " unknown multicast flooding" is set to " Optimized Flood" in a bridge domain. This knob can be set to " Optimized Flood" only for a maximum of 50 bridge domains per leaf switch. If " Optimized Flood" is enabled for more than the supported number of bridge domains on a leaf switch, follow these configuration steps to recover: Set " unknown multicast flooding" to " Flood" for all bridge domains mapped to a leaf switch. Set " unknown multicast flooding" to " Optimized Flood" on needed bridge domains.	15.3(2a) and later
N/A	Traffic destined to Static Route EP VIPs sourced from N9000 switches (switches with names that end in -EX) might not function properly because proxy route is not programmed.	15.3(2a) and later
N/A	An iVXLAN header of 50 bytes is added for traffic ingressing into the fabric. A bandwidth allowance of (50/50 + ingress_packet_size) needs to be made to prevent oversubscription from happening. If the allowance is not made, oversubscription might happen resulting in buffer drops.	15.3(2a) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported in combination with static endpoint configurations.	15.3(2a) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported with Layer 2–only bridge domains. Such a configuration will not be blocked, but the configuration will not take effect as there is no Layer 3 learning in these bridge domains.	15.3(2a) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported with external and infra bridge domains because there is no Layer 3 learning in these bridge domains.	15.3(2a) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported with a shared services provider configuration. The same or overlapping prefix cannot be used for a shared services provider and IP Ckt endpoint. However, this configuration can be applied in bridge domains having shared services consumer endpoint groups.	15.3(2a) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported with dynamic endpoint groups. Only static endpoint groups are supported.	15.3(2a) and later
N/A	No fault will be raised if the IP/MAC Ckt endpoint prefix configured is outside of the bridge domain subnet range. This is because a user can configure bridge domain subnet and IP/MAC Ckt endpoint in any order and so this is not error condition. If the final configuration is such that a configured IP/MAC Ckt endpoint prefix is outside all bridge domain subnets, the configuration has no impact and is not an error condition.	15.3(2a) and later
N/A	Dynamic deployment of contracts based on instrImmedcy set to onDemand/lazy not supported; only immediate mode is supported.	15.3(2a) and later
N/A	When a server and load balancer are on the same endpoint group, make sure that the Server does not generate ARP/GARP/ND request/response/solicits. This will lead to learning of LB virtual IP (VIP) towards the Server and defeat the purpose of DSR support.	15.3(2a) and later
N/A	Direct server return is not supported for shared services. Direct server return endpoints cannot be spread around different virtual routing and forwarding (VRF) contexts.	15.3(2a) and later
N/A	Configurations for a virtual IP address can only be /32 or /128 prefix.	15.3(2a) and later

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
N/A	Client to virtual IP address (load balancer) traffic always will go through proxy-spine because fabric data-path learning of a virtual IP address does not occur.	15.3(2a) and later
N/A	GARP learning of a virtual IP address must be explicitly enabled. A load balancer can send GARP when it switches over from active-to-standby (MAC changes).	15.3(2a) and later
N/A	Learning through GARP will work only in ARP Flood Mode.	15.3(2a) and later
N/A	A fault is not raised when a node that has an expired SUDI certificate joins the fabric.	15.3(2a) and later

Compatibility Information

- For the supported optics per device, see the Cisco Optics-to-Device Compatibility Matrix.
- 100mb optics, such as the GLC-TE, are supported in 100mb speed only on -EX, -FX, -FX2, and -FX3 switches, such as the N9K-C93180YC-EX and N9K-C93180YC-FX, and only on front panel ports 1/1-48. 100mb optics are not supported any other switches. 100mb optics cannot be used on EX or FX leaf switches on port profile converted downlink ports (1/49-52) using QSA.
- This release supports the hardware and software listed on the ACI Ecosystem Compatibility List, and supports the Cisco AVS, release 5.2(2)SV3(3.10).
- To connect the N2348UPQ to ACI leaf switches, the following options are available:
  - Directly connect the 40G FEX ports on the N2348UPQ to the 40G switch ports on the ACI leaf switches
  - Break out the 40G FEX ports on the N2348UPQ to 4x10G ports and connect to the 10G ports on all other ACI leaf switches

  Note: A fabric uplink port cannot be used as a FEX fabric port.
- To connect the Cisco APIC (the controller cluster) to the Cisco ACI fabric, it is required to have a 10G interface on the ACI leaf switch.
- We do not qualify third party optics in Cisco ACI. When using third party optics, the behavior across releases is not guaranteed, meaning that the optics might not work in some NX-OS releases. Use third party optics at your own risk. We recommend that you use Cisco SFPs, which have been fully tested in each release to ensure consistent behavior.
- On Cisco ACI platforms, 25G copper optics do not honor auto-negotiation, and therefore auto-negotiation on the peer device (ESX or standalone) must be disabled to bring up the links.
- 10G GLC-T transceivers cannot be used for the initial bring up between the Cisco APIC and a leaf switch. The fabric discovery process cannot occur because the transceiver needs the SFP media type to be pushed from the Cisco APIC to bring up the link.
- You cannot use the 100 megabit speed of a switch's QSFP28 ports.
- Cisco NX-OS uses an SSL library called CiscoSSL, which is a customized version of the OpenSSL library to support CVE fixes and FIPS compliance. Cisco maintains an extended support contract

© 2025 Cisco and/or its affiliates. All rights reserved.

with OpenSSL. CVE fixes from OpenSSL upstream is regularly incorporated in the older versions of CiscoSSL library as well.

- If you are using 10G copper cables, when you configure a link level policy, you must set the **Physical Media Type** to " SFP 10G TX."
- Do not use dual-rate optics (SFP-10/25G-LR-S and SFP-10/25G-CSR-S) for a Cisco APIC-to-switch fabric connection.

Table 12. Modular Spine Switch Fabric Module Compatibility Information

Product ID	N9K-C9504-FM-G	N9K-C9508-FM-G	N9K-C9504-FM-E	N9K-C9508-FM-E	N9K-C9508-FM-E2	N9K-C9516-FM-E2
N9K-X9716D-GX	4	4	No	No	4	4
N9K-X9736C-FX	5	5	5	5	5	5
N9K-X9736Q-FX	5	5	5	5	5	5
N9K-X9732C-EX	No	No	4	4	4	4

Table 13. Modular Spine Switch Line Card Compatibility Information

Product ID	Compatibility Information
N9K-X9716D-GX	If you connect a Cisco N9K-X9716D-GX breakout port to a non-Cisco ACI peer, such as a standalone switch capable of 100G, the link comes up and LLDP is detected. However, this is an unsupported scenario, but no fault is generated.

Table 14. Fixed Spine Switches Compatibility Information

Product ID	Compatibility Information
N9K-C9364C	You can deploy multipod or Cisco ACI Multi-Site separately (but not together) on the Cisco N9K-9364C switch starting in the 3.1 release. You can deploy multipod and Cisco ACI Multi-Site together on the Cisco N9K-9364C switch starting in the 3.2 release. A 930W-DC PSU (NXA-PDC-930W-PE or NXA-PDC-930W-PI) is supported in redundancy mode if 3.5W QSFP+ modules or passive QSFP cables are used and the system is used in 40C ambient temperature or less; for other optics or a higher ambient temperature, a 930W-DC PSU is supported only with 2 PSUs in non-redundancy mode. 1-Gigabit QSA is not supported on ports 1/49-64. This switch supports the following PSUs: • NXA-PAC-1200W-PE • NXA-PAC-1200W-PI • N9K-PUV-1200W • NXA-PDC-930W-PE • NXA-PDC-930W-PI
N9K-C9364D-GX2A	Ports 65 and 66 do not support flow telemetry nor NetFlow.
N9K-C9348D-GX2A	Ports 65 and 66 do not support flow telemetry nor NetFlow.
N9K-C9332D-GX2B	The following information applies to this switch:

© 2025 Cisco and/or its affiliates. All rights reserved.

Product ID	Compatibility Information
	• Ports 33 and 34 do not support the following things: ◦ 10G GLC-T optics ◦ 100M speed ◦ Flow telemetry ◦ NetFlow ◦ MACsec, PTP, and SyncE ◦ PFC and no-drop classes ◦ FC and FCoE mode • Port-side exhaust (PE) fans are not supported.

Table 15. Fixed Leaf Switches Compatibility Information

Product ID	Compatibility Information
N9K-C9364C-GX	This switch has the following limitations: • For ports 1-64, every 4 ports 1-4,5-8...60-64 is referred as a quad. Each quad can be operated only with a fixed speed. For example: Ports 1-4 can operate only on 10G or 40/100G. Similarly, ports 60-64 can operate only on 10G or 40/100G. • You cannot use mixed speeds of 10G and 40G, 10G and 100G, or 40G and 100G in a quad (1-4,5-8...21-24). Based on the port bring up sequence, the port in the quad where a speed mismatch is detected will be HW disabled. • If there is a speed mismatch in a quad even when the ports are configured in the disabled state, the working links in that quad might get into the HW disabled state upon upgrading and reloading because the mixed speed is brought up first before the admin down configuration is pushed. As a result, you must manually perform the shut and no shut commands on the ports to bring up the links. • Breakout of 4x25G or 4x10G ports is not supported. • There is a lane selector button on the hardware. The button is used for the breakout port LED status. Because breakout is not supported, this button does nothing. • The maximum number of downlinks is 30 x 4 ports 10/25G (breakout) + 2 ports (1/61-62) = 122 ports. Ports 1/63 and 1/64 are reserved for fabric links and even numbers from 1/1 to 1/60 are error-disabled. • 1G and 100MB speeds are not supported.
N9K-C93600CD-GX	This switch has the following limitations: • Auto-negotiation is not supported with 10G speed on ports 1 through 24. • For ports 1 through 24, every 4 ports (1-4, 5-8, 9-12, and so on, referred to as a "quad") will operate at a fixed speed. That is, all 4 ports will operate in 10G or 40/100G; you cannot mix the speeds. • Mixed speeds of 10G and 40G or 10G and 100G in a quad is not supported. Based on the port bring up sequence, the port in the quad where the speed mismatch is detected will be HW disabled. • If there is a speed mismatch in a quad even though the ports are configured in the disabled state, the working links in that quad might get into the HW disabled state upon upgrading or reloading, as the mixed speed is brought up first before admin down config is pushed. To avoid this issue, you must manually use the shut and no shut commands on the working ports to bring up the links. For more information, see bug CSCvr61096. • Ports 25-26 and ports 27-28 (port groups of 2 ports each) will operate in a fixed speed within the respective group, and you cannot mismatch the speed. • Uplink ports 29 to 36 do not have a mixed speed restriction; you can toggle the speed for the bidirectional ports. • For ports 1 to 28, even if you convert any ports to uplink with bidirectional optics, you cannot toggle the speed, as it will introduce mixed speeds and will disturb the neighboring ports. • For ports 1 to 28, if any of the ports are converted to uplink with bidirectional optics, the ports will stay in the not connected state if the peer is a 40G link. • 4x10 and 4x25 breakout is supported on ports 25-28 and 29-34 (port profile converted downlinks). • Ports 25-26 and 27-28 form respective port pairs, and each pair can operate with 4x10, 10G, or

© 2025 Cisco and/or its affiliates. All rights reserved.

Product ID	Compatibility Information
	4x25G speed. • This switch does not support 4x100 breakout in this release. • The Hardware Abstraction Layer (HAL) will spike and the console can hang if a port channel or vPC exists when overlying breakout ports are deleted. To avoid this issue, delete the PC or vPC before deleting the overlying breakout policy. • The maximum number of downlinks is 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 error-disabled. • 1G and 100M speeds are not supported.
N9K-C9364D-GX2A	Ports 65 and 66 do not support flow telemetry nor NetFlow.
N9K-C9348D-GX2A	Ports 65 and 66 do not support flow telemetry nor NetFlow.
N9K-C9332D-GX2B	The following information applies to this switch: • Ports 33 and 34 do not support the following things: ◦ 10G GLC-T optics ◦ 100M speed ◦ Flow telemetry ◦ NetFlow ◦ MACsec, PTP, and SyncE ◦ PFC and no-drop classes ◦ FC and FCoE mode • Port-side exhaust (PE) fans are not supported.
N9K-C9316D-GX	Auto-negotiation and forward error correction are not supported when you use this switch is as a leaf switch.
N9K-C93180YC-FX3	The following information applies to this switch: • The following ports are not supported: ◦ Antenna ◦ GNSS ◦ GPS ◦ PPS ◦ PTP GM • When using the SFP-10G-T-X optic on a port, the you must either leave the physically adjacent ports empty or only deploy direct attach cables (DACs) to those ports. • If you insert a non-DAC optic in a port that is physically adjacent to a port that is capable of supporting a 10G GLC-T optic, and later you insert a GLC-T optic into the GLC-T-capable port, the GLC-T optic will be hw-disabled. To bring up the GLC-T port, you must shut down the non-DAC port, then run the shut and no shut commands on the GLC-T port.
N9K-C9336C-FX2	The following information applies to this switch: • On older N9K-C9336C-FX2 switches, auto-negotiation does not work on port eth1/4. You can check whether your switch is older by using the following command: ifav124-leaf5# cat /sys/kernel/cisco_board_info/hw_change_bits 0x0 The output of "0x0" indicates an older switch that has this limitation. • You can apply a breakout configuration on ports 1 through 34, which can give up to 136 (34*4) server or downlink ports. • Port profiles and breakouts are not supported on the same port. However, you can apply a port profile to convert a fabric port to a downlink, and then apply a breakout configuration. • If you apply a breakout configuration on 34 ports, you must configure a port profile on the ports first,

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	Compatibility Information
	which requires you to reboot the leaf switch. • If you apply a breakout configuration to a leaf switch for multiple ports at the same time, it can take up to 10 minutes for the hardware of 34 ports to be programmed. The ports remain down until the programming completes. The delay can occur for a new configuration, after a clean reboot, or during switch discovery. • Ports 7 through 32 have a link bring up time of less than 2 seconds with QSFP-100G-LR4 and QSFP-40/100G-SRBD optics. For all other ports, the link up time for these optics is between 5 to 14 seconds. In the following situations, the link bring up time will also be greater than 2 seconds: ◦ After reloading the leaf switch switch ◦ When using port optical insertion and removal (OIR) ◦ When performing bulk flaps of ports on the leaf switch
N9K-C93240YC-FX2	The following information applies when this switch is configured with port-side intake airflow: • Ports 2, 6, 8, 12, 14, 18, 20, 24, 26, 30, 32, 36, 38, 42, 44, and 48 are capable of supporting the 10G GLC-T optic. After you configure these ports to use 10G GLC-T, these ports will be the only ports on the switch that can support 10G GLC-T. Without being configured for 10G GLC-T, these ports behave as normal switch ports. • If you configure port 12 for 10G GLC-T, then ports 9 and 15 must either be left empty or must deploy only DACs. • Ports 49 through 60 can be configured to use 10G GLC-T or can be normal ports, regardless of the configuration of the other ports. The following information applies when this switch is configured with port-side exhaust airflow: • Ports 6, 12, 18, 24, 30, 36, 42, and 48 are capable of supporting the 10G GLC-T optic. After you configure these ports to use 10G GLC-T, these ports will be the only ports on the switch that can support 10G GLC-T. Without being configured for 10G GLC-T, these ports behave as normal switch ports. • If you configure port 12 for 10G GLC-T, then ports 9, 11, and 15 must either be left empty or must deploy only DACs. • Ports 49 through 60 can be configured to use 10G GLC-T or can be normal ports, regardless of the configuration of the other ports. The following information applies regardless of the airflow direction: • When using the SFP-10G-T-X optic on a port, the you must either leave the physically adjacent ports empty or only deploy direct attach cables (DACs) to those ports. • If you insert a non-DAC optic in a port that is physically adjacent to a port that is capable of supporting a 10G GLC-T optic, and later you insert a GLC-T optic into the GLC-T-capable port, the GLC-T optic will be hw-disabled. To bring up the GLC-T port, you must shut down the non-DAC port, then run the **shut** and **no shut** commands on the GLC-T port.
N9K-C93360YC-FX2	The following information applies to this switch: • Ports 1, 4, 5, 8, 41, 44, 45, 48, 49, 52, 53, 56, 57, 60, 61, 64, 65, 68, 69, 72, 73, 76, 77, 80, 81, 84, 85, 88, 89, 92, 93, and 96 are capable of supporting the 10G GLC-T optic. After you configure these ports to use 10G GLC-T, these ports will be the only ports on the switch that can support 10G GLC-T. Without being configured for 10G GLC-T, these ports behave as normal switch ports. • If you configure port 60 for 10G GLC-T, then ports 58, 59, and 62 must either be left empty or must deploy only DACs. • Ports 97 through 108 can be configured to use 10G GLC-T or can be normal ports, regardless of the configuration of the other ports. • When using the SFP-10G-T-X optic on a port, the you must either leave the physically adjacent ports empty or only deploy direct attach cables (DACs) to those ports. • If you insert a non-DAC optic in a port that is physically adjacent to a port that is capable of supporting a 10G GLC-T optic, and later you insert a GLC-T optic into the GLC-T-capable port, the GLC-T optic will be hw-disabled. To bring up the GLC-T port, you must shut down the non-DAC port, then run the **shut** and **no shut** commands on the GLC-T port.
N9K-C9348GC-FXP	This switch supports the following PSUs: • NXA-PAC-350W-PI

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	Compatibility Information
	• NXA-PAC-350W-PE • NXA-PAC-1100W-PI • NXA-PAC-1100W-PE The following information applies to this switch: • Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops. • This switch does not support the 10G GLC-T optic. • The PSU SPROM is not readable when the PSU is not connected. The model displays as " UNKNOWN" and status of the module displays as " shutdown."
N9K-C93180YC-FX-24	This switch does not support the 10G GLC-T optic.
N9K-C93180YC-FX	The following information applies to this switch: • Auto-negotiation is not supported if you convert port 51 or 52 to a downlink and you have 40/100G copper cables connected. • Ports 1, 4, 5, 8, 9, 12, 13, 16, 37, 40, 41, 44, 45, and 48 are capable of supporting the 10G GLC-T and SFP-10G-T-X optics. After you configure these ports to use 10G GLC-T or SFP-10G-T-X, these ports will be the only ports on the switch that can support 10G GLC-T or SFP-10G-T-X. Without being configured for 10G GLC-T nor SFP-10G-T-X, these ports behave as normal switch ports. • If you configure port 12 for 10G GLC-T, then ports 10, 11, and 14 must either be left empty or must deploy only DACs. • Ports 49 through 54 can be configured to use 10G GLC-T or can be normal ports, regardless of the configuration of the other ports. • When using the SFP-10G-T-X optic on a port, the you must either leave the physically adjacent ports empty or only deploy direct attach cables (DACs) to those ports. • If you insert a non-DAC optic in a port that is physically adjacent to a port that is capable of supporting a 10G GLC-T optic, and later you insert a GLC-T optic into the GLC-T-capable port, the GLC-T optic will be hw-disabled. To bring up the GLC-T port, you must shut down the non-DAC port, then run the **shut** and **no shut** commands on the GLC-T port.
N9K-C93180YC-EX-24	This switch does not support the 10G GLC-T optic.
N9K-C93180YC-EX	The following information applies to this switch: • The following FEC modes are not supported on N9K-C93180YC-EX ports 1 through 48 when running in 25G speed: ◦ cl91-rs-fec ◦ cons16-rs-fec ◦ ieee-rs-fec • Auto-negotiation is not supported if you convert port 51 or 52 to a downlink and you have 40/100G copper cables connected. • Ports 1, 4, 5, 8, 9, 12, 13, 16, 37, 40, 41, 44, 45, and 48 are capable of supporting the 10G GLC-T and SFP-10G-T-X optics. After you configure these ports to use 10G GLC-T or SFP-10G-T-X, these ports will be the only ports on the switch that can support 10G GLC-T or SFP-10G-T-X. Without being configured for 10G GLC-T nor SFP-10G-T-X, these ports behave as normal switch ports. • If you configure port 12 for 10G GLC-T, then ports 10, 11, and 14 must either be left empty or must deploy only DACs. • Ports 49 through 54 can be configured to use 10G GLC-T or can be normal ports, regardless of the configuration of the other ports. • When using the SFP-10G-T-X optic on a port, the you must either leave the physically adjacent ports empty or only deploy direct attach cables (DACs) to those ports. • If you insert a non-DAC optic in a port that is physically adjacent to a port that is capable of supporting a 10G GLC-T optic, and later you insert a GLC-T optic into the GLC-T-capable port, the GLC-T optic will be hw-disabled. To bring up the GLC-T port, you must shut down the non-DAC port, then run the **shut** and **no shut** commands on the GLC-T port.

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	Compatibility Information
N9K-C93180LC-EX	This switch has the following limitations: • The top and bottom ports must use the same speed. If there is a speed mismatch, the top port takes precedence and bottom port will be error disabled. Both ports both must be used in either the 40 Gbps or 10 Gbps mode. • Ports 26 and 28 are hardware disabled. • This release supports 40 and 100 Gbps for the front panel ports. The uplink ports can be used at the 100 Gbps speed. • Port profiles and breakout ports are not supported on the same port.

Table 16. CloudSec Support

Product ID	Hardware Type	CloudSec Support
N9K-C9332C	Switch	Yes, only on the last 8 ports
N9K-C9364C	Switch	Yes, only on the last 16 ports
N9K-X9736C-FX	Line Card	Yes, only on the last 8 ports

- The following additional CloudSec compatibility restrictions apply:
  - CloudSec only works with spine switches in Cisco ACI and only works between sites managed by Cisco ACI Multi-Site.
  - For CloudSec to work properly, all of the spine switch links that participate in Cisco ACI Multi-Site must have MACsec/CloudSec support.

Usage Guidelines

- The current list of protocols that are allowed (and cannot be blocked through contracts) include the following. Some of the protocols have SrcPort/DstPort distinction.
  - UDP DestPort 161: SNMP. These cannot be blocked through contracts. Creating an SNMP ClientGroup with a list of Client-IP Addresses restricts SNMP access to only those configured Client-IP Addresses. If no Client-IP address is configured, SNMP packets are allowed from anywhere.
  - TCP SrcPort 179: BGP
  - TCP DstPort 179: BGP
  - OSPF
  - UDP DstPort 67: BOOTP/DHCP
  - UDP DstPort 68: BOOTP/DHCP
  - IGMP
  - PIM
  - UDP SrcPort 53: DNS replies
  - TCP SrcPort 25: SMTP replies

© 2025 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

    - TCP DstPort 443: HTTPS
    - UDP SrcPort 123: NTP
    - UDP DstPort 123: NTP
- Leaf switches and spine switches typically have memory utilization of approximately 70% to 75%, even in a new deployment where no configuration has been pushed. This amount of memory utilization is due to the Cisco ACI-specific processes, which take up more memory compared to a standalone Nexus deployment. The memory utilization is not a problem unless it exceeds 90%. You can open a Cisco TAC case to troubleshoot proactively when memory utilization is more than 85%.
- Leaf and spine switches from two different fabrics cannot be connected regardless of whether the links are administratively kept down.
- If you replace a switch where a Cisco APIC is connected, make sure that the Cisco APIC has two connections: one active/backup to the replaced switch and another to a different switch. Otherwise, the Cisco APIC will not join the cluster after you replace the switch.
- Only one instance of OSPF (or any multi-instance process using the managed object hierarchy for configurations) can have the write access to operate the database. Due to this, the operational database is limited to the default OSPF process alone and the multipodInternal instance does not store any operational data. To debug an OSPF instance ospf-multipodInternal, use the command in VSH prompt. Do not use ibash because some ibash commands depend on Operational data stored in the database.
- When you enable or disable Federal Information Processing Standards (FIPS) on a Cisco ACI fabric, you must reload each of the switches in the fabric for the change to take effect. The configured scale profile setting is lost when you issue the first reload after changing the FIPS configuration. The switch remains operational, but it uses the default port scale profile. This issue does not happen on subsequent reloads if the FIPS configuration has not changed.
    - FIPS is supported on Cisco NX-OS release 15.2(2) or later. If you must downgrade the firmware from a release that supports FIPS to a release that does not support FIPS, you must first disable FIPS on the Cisco ACI fabric and reload all of the switches in the fabric.
- You cannot use the breakout feature on a port that has a port profile configured on a Cisco N9K-C93180LC-EX switch. With a port profile on an access port, the port is converted to an uplink, and breakout is not supported on an uplink. With a port profile on a fabric port, the port is converted to a downlink. Breakout is currently supported only on ports 1 through 24.
- On Cisco 93180LC-EX Switches, ports 25 and 27 are the native uplink ports. Using a port profile, if you convert ports 25 and 27 to downlink ports, ports 29, 30, 31, and 32 are still available as four native uplink ports. Because of the threshold on the number of ports (which is maximum of 12 ports) that can be converted, you can convert 8 more downlink ports to uplink ports. For example, ports 1, 3, 5, 7, 9, 13, 15, 17 are converted to uplink ports and ports 29, 30, 31 and 32 are the 4 native uplink ports, which is the maximum uplink port limit on Cisco 93180LC-EX switches.
    - When the switch is in this state and if the port profile configuration is deleted on ports 25 and 27, ports 25 and 27 are converted back to uplink ports, but there are already 12 uplink ports on the switch in the example. To accommodate ports 25 and 27 as uplink ports, 2 random ports from the port range 1, 3, 5, 7, 9, 13, 15, 17 are denied the uplink conversion; the chosen ports cannot be controlled by the user. Therefore, it is mandatory to clear all the faults before reloading the leaf node to avoid any unexpected behavior regarding the port type. If a node is reloaded without clearing the

© 2025 Cisco and/or its affiliates. All rights reserved.

port profile faults, especially when there is a fault related to limit-exceed, the ports might be in an unexpected mode.

- When using a 25G Mellanox cable that is connected to a Mellanox NIC, you can set the ACI leaf switch port to run at a speed of 25G or 10G.
- You cannot enable auto-negotiation on the spine switch or leaf switch side with 40G or 100G CR4 optics. For 40G copper transceivers, you must disable auto-negotiation and set the speed to 40G. For 100G copper transceivers, you must disable auto-negotiation on the remote end and set the speed to 100G.
- You cannot enable auto-negotiation on an active QSFP to SFP/SFP+ Adapter (QSA) module. You can enable auto-negotiation only on a passive QSA module. The following example CLI command shows an active QSA module:

  module-1# **show platform internal usd port info | grep -A 10 "Eth1/42"**

  Port 107.0 (Eth1/42) : Admin UP Link DOWN Cfg_Fec Disabled Fec Disabled Fcot Copper retimer 0x116c0100

  AN_cfg Yes AN_operSt No In_debounce 0, Debounce-Time 100000 usecs SM sm qsa: Yes

  The following example CLI command shows a passive QSA module:

  module-1# **show platform internal usd port info | grep -A 10 "Eth1/43"**

  Port 109.0 (Eth1/43) : Admin UP Link UP Cfg_Fec Disabled Fec Disabled Fcot Copper retimer 0x116c0100

  AN_cfg Yes AN_operSt No In_debounce 0, Debounce-Time 100000 usecs SM sm qsa: Passive
- You can enable auto-negotiation for 10G, 25G, 40G, or 100G on downlink ports on a Cisco ACI leaf switch. However, you cannot enable auto-negotiation on spine ports and uplink ports on a Cisco ACI leaf switch. Therefore, if the Inter-Pod Network (IPN) is connected to the spine ports using copper cables, you should disable auto-negotiation on the peer node that is the IPN port. Similarly, if a remote leaf switch is connected to the IPN using copper cables on the uplink port, you should disable auto-negotiation on the peer node that is the IPN port.
- A 25G link that is using the IEEE-RS-FEC mode can communicate with a link that is using the CL16-RS-FEC mode. There will not be a FEC mismatch and the link will not be impacted.
- When the provider edge router is an IOS XR device, the router does not support route re-origination from one EVPN stitching site to another EVPN stitching site.

Related Content

See the Cisco Application Policy Infrastructure Controller (APIC) page for the documentation.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, send your comments to apic-docfeedback@cisco.com. We appreciate your feedback.

© 2025 Cisco and/or its affiliates. All rights reserved.

Legal Information

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: http://www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2025 Cisco Systems, Inc. All rights reserved.