



Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 15.2(3)

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 5.2\(3\)](#).

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

Date	Description
April 27, 2023	Added the N9K-M12PQ, N9K-M6PQ, and N9K-M6PQ-E expansion modules to the No Longer Supported section.
October 10, 2022	In the Open Issues section, added bug CSCvz81692.
September 6, 2022	In the Open Issues section, added bug CSCwc23246.
July 28, 2022	In the Open Issues section, added bug CSCvz31425.
May 16, 2022	In the Open Issues section, added bug CSCwa47686.
April 12, 2022	In the Open Issues section, added bug CSCwa64076.
April 4, 2022	In the Open Issues section, added bug CSCwa92634.
March 16, 2022	In the Open Issues section, added bugs CSCwa67425, CSCwa95241, CSCwb12482, and CSCwb14844.
February 11, 2022	In the Open Issues section, added bug CSCvy10791.
February 8, 2022	In the Open Issues section, added bug CSCwa86674.
February 2, 2022	In the Open Issues section, added bug CSCwa72232.
December 18, 2021	Release 15.2(3g) became available; there are no changes to this document for this release. See the Cisco Application Policy Infrastructure Controller Release Notes, Release 5.2(3) for the changes in this release.
December 16, 2021	In the Resolved Issues section, added bug CSCvz95984.
December 7, 2021	In the Compatibility Information section, for the N9K-C9364C-GX switch, added: All 4 fans must be operational, otherwise the switch will power down due to a fan policy trigger.
November 24, 2021	Release 15.2(3f) became available. In the Resolved Issues section, added bug CSCwa12371.
November 17, 2021	In the Open Issues section, added bug CSCwa12371.
November 4, 2021	In the Resolved Issues section, added bug CSCvy08825.

Date	Description
October 29, 2021	Added the Cisco N9K-C9332D-GX2B switch to the fixed spine switch tables.
October 18, 2021	Release 15.2(3e) 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

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

Product ID	Description	Minimum	Maximum
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

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-C9332D-GX2B	Cisco Nexus 9300 platform switch with 32p 400/100-Gbps QSFP-DD ports and 2p 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).

Product ID	Description
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
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

Product ID	Description
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-C9332D-GX2B	Cisco Nexus 9300 platform switch with 32p 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 64.

Product ID	Description
N9K-C9336C-FX2	<p>Cisco Nexus C9336C-FX2 Top-of-rack (ToR) switch with 36 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports.</p> <p>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.</p>
N9K-C93108TC-FX	<p>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.</p> <p>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.</p>
N9K-C93108TC-FX-24	<p>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.</p> <p>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.</p>
N9K-C93180YC-FX	<p>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.</p> <p>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.</p>
N9K-C93180YC-FX-24	<p>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.</p> <p>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.</p>
N9K-C9348GC-FXP	<p>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.</p>
N9K-C93108TC-EX	<p>Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) front panel ports and 6 40/100-Gigabit QSFP28 spine facing ports.</p>
N9K-C93108TC-EX-24	<p>Cisco Nexus 9300 platform switch with 24 1/10GBASE-T (copper) front panel ports and 6 40/100-Gigabit QSFP28 spine facing ports.</p>
N9K-C93180LC-EX	<p>Cisco Nexus 9300 platform switch with 24 40-Gigabit front panel ports and 6 40/100-Gigabit QSFP28 spine-facing ports.</p> <p>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.</p>
N9K-C93180YC-EX	<p>Cisco Nexus 9300 platform switch with 48 1/10/25-Gigabit front panel ports and 6-port 40/100 Gigabit QSFP28 spine-facing ports.</p>
N9K-C93180YC-EX-24	<p>Cisco Nexus 9300 platform switch with 24 1/10/25-Gigabit front panel ports and 6-port 40/100 Gigabit QSFP28 spine-facing ports.</p>
N9K-C93120TX	<p>Cisco Nexus 9300 platform switch with 96 1/10GBASE-T (copper) front panel ports and 6-</p>

Product ID	Description
	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
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

Product ID	Description
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
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 ToR 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
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

Product Type	Product ID
	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 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/fex_tables.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

- The Cisco Nexus 9332D-GX2B switch (N9K-C9332D-GX2B) has 32 400/100-Gbps QSFP-DD ports and 2 1/10 SFP+ ports.
- The 400Gb to 100Gb breakout feature is now supported on the access ports of the following switches:
 - N9K-C9332D-GX2B with the QDD-4X100G-FR-S or QDD-400G-DR4-S transceiver
- The 400Gb to 100Gb breakout feature is now supported on the fabric ports of the following switches:
 - N9K-C9332D-GX2B with the QDD-4X100G-FR-S transceiver

- The Platform Insights Engine (PIE) was implemented, which is an on-switch real time root cause analysis application. To keep the service outages in a switch to a minimum, PIE gives insights about the current issues on the switch and visibility about the potential issues. For more information, see the [Cisco APIC and Platform Insights Engine](#) document.
- MACsec is now supported on the following switches:
 - N9K-C9332D-GX2B
- For PTP, the N9K-C93108TC-FX3P switch now supports the 1G speed.

New Software Features

For new software features, see the [Cisco Application Policy Infrastructure Controller Release Notes, Release 5.2\(3\)](#).

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 15.2(3) releases in which the bug exists. A bug might also exist in releases other than the 15.2(3) releases.

Bug ID	Description	Exists in
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.	15.2(3g) and later
CSCwac04832	A PBR health-group goes down due to IPSLA issues only on one side of the vPC leaf switch pair.	15.2(3g) and later
CSCwf53105	"vsh" process generates multiple core files on switches after starting OnDemand Techsupport collection for leaf switches.	15.2(3g) and later
CSCwf53105	"vsh" process generates multiple core files on switches after starting OnDemand Techsupport collection for leaf switches.	15.2(3g) 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.	15.2(3g) and later
CSCwa67425	In a remote leaf switch vPC setup, when the primary vPC reloads, the remaining peer suspends its vPC legs as if it lost its fabric links, even though it has not.	15.2(3f) and later
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.	15.2(3e) and later
CSCvw89840	Traffic originating from a vPC TEP is dropped for Layer 2 multicast and unknown unicast traffic when pod redundancy is triggered.	15.2(3e) and later

Bug ID	Description	Exists in
CSCvx31008	This issue is only observed with Cisco ACI Multi-Pod Infra OSPF B2B links. When B2B OSPF link goes down, the faults for the multiPodDirect instance will not be raised. Also, the Oper for the multiPodDirect instance will not be updated in the DME DB.	15.2(3e) and later
CSCvy10791	<p>When an EPG is configured to deploy VLAN X on port X, a VXLAN ID called FD_VNID (or fabric encap) is assigned to VLAN X. This VXLAN ID is derived from the VLAN pool to which VLAN X belongs. If there are multiple VLAN pools that can provide a VXLAN ID for the same VLAN ID, the current logic picks the VLAN pool that is hit first during the internal lookup.</p> <p>This may cause inconsistent VXLAN ID assignment for VLANs on each leaf switch after an upgrade or stateless reload (i.e. clean reload).</p> <p>In the past, bug CSCvj84175 implemented an option to prevent the VLAN pools with overlapping ranges from being associated to the same EPG. However, the knob does not resolve the issue if the configuration was already done before the option was enabled.</p> <p>This enhancement is to implement more robust logic to prevent randomness in VLAN pool selection even if there are overlapping configurations.</p>	15.2(3e) and later
CSCvy31805	The PBR destination group for bypass action is not properly programmed with PBR service graph for service devices behind I3out 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.	15.2(3e) and later
CSCvy94715	PTP is not supported on breakout ports.	15.2(3e) and later
CSCvz31425	A Nexus Insights app sees some extra flows sent from the spine switch when a subnet is added in a VRF instance for flow telemetry and then the subnet is deleted.	15.2(3e) and later
CSCvz49846	In summary, a remote leaf switch endpoint is sending multicast traffic, and after a spine switch fabric card reload, traffic is hashed on this fabric card, which causes the endpoint to be marked as rogue and penalize that endpoint for 30 minutes.	15.2(3e) and later
CSCvz59605	Day-1 issue, and is present in standalone as well. PIM interface policy with MD5 authentication applied only on Cisco ACI side Layer 3 SVI interface. The PIM neighborship with an external leaf switch still exists on the border leaf switches with the expiry timer as "N/A." Non-direct border leaf switches takes more time than direct border leaf switches to update "N/A" as the expired state.	15.2(3e) and later
CSCvz81692	A vPC domain's interfaces are suspended/down when its peer node is decommissioned.	15.2(3e) and later
CSCvz84323	The OSPF route is not installed in the routing table even though the OSPF database learned the route. This happens for type-5/type-7 LSAs that have a forwarding address in them.	15.2(3e) and later
CSCvz90912	When applying SMUs to Cisco ACI Nexus 9000 modular switch, there is a data syncing issue causing SMU inconsistency between two supervisors. This leads to an extra reload of the switch to reach the SMU consistent state.	15.2(3e) and later
CSCvz95266	When applying SMUs to Cisco ACI Nexus 9000 modular switch, depending on the supervisor boot up sequence, a duplicate configuration may be pushed to the active supervisor, causing an incorrect update of the SMU status. A correctly activated SMU's status will be incorrectly shown as "failed", "already exists".	15.2(3e) and later
CSCvz97089	There is traffic duplication when using PIM over the SVI L3Out feature during the reload of a border leaf switch.	15.2(3e) and later

Bug ID	Description	Exists in
CSCvz98691	With the PIM over SVI L3Out feature, a PIM neighbor is shown to have a direct connection toward an external box, even though the STP has this link down and the connection is through another border leaf switch. This affects traffic to and from this border leaf switch.	15.2(3e) and later
CSCwa12763	External route import for a VRF instance fails on a leaf switch after removing a shared services contract between two EPGs.	15.2(3e) and later
CSCwa47686	For a Cisco ACI fabric with more than 128 leaf switches in a given pod, such as 210 leaf switches in a single pod deployment, after enabling PTP globally, only 128 leaf switches are able to enable PTP. The remaining 82 leaf switches fail to enable PTP due to the error F2728 latency-enable-failed.	15.2(3e) and later
CSCwa64076	Some activated SMUs do not show with the "show version" command or in the GUI.	15.2(3e) and later
CSCwa72232	400G fiber optics fail to come up with the Cisco N9K-C93600CD-GX switch in spine mode when upgrading from the 15.2(2e) release to the 5.2(3e) release.	15.2(3e) and later
CSCwa86674	A switch reloads due to a kernel panic that is triggered by oom-killer. The dmesg log gets filled with messages similar to the following example: [1056.987558] Shared skb 0xffff880828f7f9c0 users -1519686805 dma vaddr 0xffff880828f94018 dma paddr 0x0000000828f94018 len 1802201963 vrf 1802201963 protocol 27499	15.2(3e) and later
CSCwa92634	The power supply status flaps between failed and OK.	15.2(3e) and later
CSCwa95241	An endpoint may fail to resolve an ARP request of another endpoint.	15.2(3e) and later
CSCwb08081	A route profile that matches on community list and sets the local pref and community is not working post upgrade to 5.2.x release. route-map imp-l3out-L3OUT_WAN-peer-2359297, permit, sequence 4201 Match clauses: community (community-list filter): peer16389-2359297-exc-ext-in-L3OUT_WAN_COMMUNITY-rgcom Set clauses: local-preference 200 community xxxxx:101 xxxxx:500 xxxxx:601 xxxxy:4 additive The match clause works as expected, but the set clause is ignored.	15.2(3e) and later
CSCwb12482	A 40G link takes an extended amount of time for link bring up to occur.	15.2(3e) and later
CSCwb14844	There is a long delay when connecting a Cisco N9K-C9336C-FX2 switch to a Mellanox NIC/40G using QSFP-40G-SR4. The status of the port is down, not connected.	15.2(3e) and later

Bug ID	Description	Exists in
CSCwb46927	<ul style="list-style-type: none"> 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. 	15.2(3e) and later
CSCwc23246	A leaf switch will unexpectedly reboot.	15.2(3e) and later
CSCwd29346	An ACI switch's console may continuously output messages similar to: svc_ifc_eventmg (*****) Ran 7911 msec in last 7924 msec	15.2(3e) and later
CSCwd36295	The BFD process crashes in Cisco ACI switches and the BFD process is listed in the output of the "show cores" command.	15.2(3e) and later
CSCwd44102	Fiber interfaces (QSA) show up as "Fcot Copper" in the USD port information. When a 10G Fiber optics with Copper QSA is inserted, the fcot gets updated as Copper only instead of Fiber.	15.2(3e) and later
CSCwd49996	A leaf switch may the bring front panel ports up before programming all policies, which causes traffic to get dropped. The issue is specific to a stateless reload or upgrade. The root cause is related to pushing the configuration from the Cisco APIC to the leaf switch, resolving all objects and programming hardware tables. In some corner cases, the bootstrap may be falsely claimed as completed and ports go up while the actual hardware tables are not fully ready. Another symptom may be a significant amount of time needed for bootstrap (possibly over an hour).	15.2(3e) and later
CSCwd75707	The parser drops IP-in-IP packets that are associated with IP traffic on ingress leaf switches.	15.2(3e) and later
CSCwd83289	After reloading a leaf switch, HAL becomes unresponsive, other processes generate a core, and HAL continuously runs with ZMQ processing.	15.2(3e) and later
CSCwf04145	100Mb links that use GLC-T/GLC-TE transceivers do not come up.	15.2(3e) and later
CSCwf04501	There are actrl.Rule programming failures followed by a policy-mgr HAP reset.	15.2(3e) and later
CSCwf57396	The 30 second input rate and 30 second output rate show values beyond 30 seconds for an interface that is disabled	15.2(3e) and later
CSCwh07391	Traffic coming from ISN or IPN may get misclassified as iTraceroute or will not preserve cos correctly. On any FM that was reloaded, dot1p preserve may have not been set correctly post reload.	15.2(3e) and later
CSCwh71704	When one of the vPC peers reloads and comes up, the non-reloaded peer is seen to be suspending the vPC interfaces.	15.2(3e) and later
CSCwh72876	The EPM process crashed when there was no disk space was available at /var/sysmgr/tmp_logs/.	15.2(3e) and later
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.	15.2(3e) and later

Bug ID	Description	Exists in
CSCwh78987	Breakout ports configured as port channel members are no longer part of the port channel post clean reload.	15.2(3e) and later
CSCwi89535	A leaf switch crashes due to the following reason: Reason: Reset triggered due to HA policy of ResetService: sdkhal hap resetVersion : xxx	15.2(3e) and later
CSCwa12371	Cisco ACI nodes are unable to sync with an IPv6 NTP server over an in-band management network. The packet sent counter gets incremented. However, NTP packets cannot be seen going out from the leaf node using tcpdump or ELAM. Disabling and re-enabling the NTP server or reloading the leaf switch does not fix the issue.	15.2(3e)

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
CSCwa12371	Cisco ACI nodes are unable to sync with an IPv6 NTP server over an in-band management network. The packet sent counter gets incremented. However, NTP packets cannot be seen going out from the leaf node using tcpdump or ELAM. Disabling and re-enabling the NTP server or reloading the leaf switch does not fix the issue.	15.2(3f)
CSCvy77935	Applications are slow when deployed in servers that are connected to a Tier-1 leaf switch.	15.2(3e)
CSCvy08825	DOM values are incorrectly displayed when the "show interface ethernet 1/x transceiver details" command is executed in ibash or when viewed in Web GUI. But, when the "show platform internal usd xcvr dom raw port" command is executed under vsh_lc mode, the correct values on each lane can be seen. After the command is executed in the vsh_lc mode, switching back to ibash and executing the same command returns the correct values.	15.2(3e)
CSCvy19256	Traffic is dropped for PBR redirected traffic from a remote leaf switch to the service device in the main pod, after the bridge domain used for service device hosting is completely undeployed from all of the remote leaf switches.	15.2(3e)
CSCvy24531	If PBR is configured with an L3Out, the fault "tracked-as-down" may be incorrectly raised on the destGrp object in the switch. There is no impact to traffic.	15.2(3e)
CSCvy31884	In the ipmgr process, for the primary subnet, an overlapping address is ignored and is not pushed to EPM. But, this is not enforced for secondary subnets.	15.2(3e)
CSCvy32902	The static source-specific multicast exact configuration of the quality level for EEC2/eEEC has a lower preference than the internal OSC. The same is not true if the user considers the quality level coming on wire. There, the EEC2/eEEC are treated equally, and so there is no problem. On Tx, the switch always considers and transmits the quality level as locked. We do not recommend that you configure the exact Tx in the field.	15.2(3e)

Bug ID	Description	Fixed in
CSCvy39316	This condition occurs when the RX quality level is hardcoded to eEEC. In this case, eEEC will be treated lower than ST3. The issue will not be seen if a user considers the incoming quality level on wire, which is the recommended configuration. On wire, it is treated similar to ST3.	15.2(3e)
CSCvy43640	A leaf node crashes when PFC or LLFC is enabled on a stretched fabric or a Multi-tier fabric. PFC and LLFC is mainly used for FCoE and RoCE. For a stretched fabric, when a transit leaf node that has connectivity to spine nodes in both locations receives the traffic that matches the QoS class with No-Drop-Cos and PFC enabled, the transit leaf node crashes. For a Multi-tier fabric, when a tier-2 leaf node receives the traffic that matches the QoS class with No-Drop-Cos and PFC enabled, the tier-2 leaf node crashes.	15.2(3e)
CSCvz95984	The value of " Use count" in the " show flow monitor" command output is incorrect.	15.2(3e)

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 15.2(3) releases in which the bug exists. A bug might also exist in releases other than the 15.2(3) 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.	15.2(3e) 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.	15.2(3e) and later
CSCup65586	The show interface command shows the tunnel's Rx/Tx counters as 0.	15.2(3e) and later
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.	15.2(3e) and later
CSCup92534	Continuous " threshold exceeded" messages are generated from the fabric.	15.2(3e) and later
CSCuq39829	Switch rescue user (" admin") can log into fabric switches even when TACACS is selected as the default login realm.	15.2(3e) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	15.2(3e) and later
CSCuq77095	When the command show ip ospf vrf <vrf_name> is run from bash on the border leaf, the checksum field in the output always shows a zero value.	15.2(3e) and later

Bug ID	Description	Exists in
CSCug83910	When an IP address moves from one MAC behind one ToR to another MAC behind another ToR, 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 ToR 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.	15.2(3e) and later
CSCug92447	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.	15.2(3e) and later
CSCug93389	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.	15.2(3e) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	15.2(3e) and later
CSCur81822	The access-port operational status is always "trunk" .	15.2(3e) 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.	15.2(3e) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) and later
CSCus61748	If the TOR 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 TOR 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 TOR 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.	15.2(3e) and later
CSCut59020	If Backbone and NSSA areas are on the same leaf, and default route leak is enabled, Type-5 LSAs cannot be redistributed to the Backbone area.	15.2(3e) 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.	15.2(3e) and later

Bug ID	Description	Exists in
CSCuu11351	Traffic from the orphan port to the vPC pair is only updated on the destination node, so the traffic count shows as excess.	15.2(3e) 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.	15.2(3e) and later
CSCuv57302	Atomic counters on the border leaf do not increment for traffic from an endpoint group going to the Layer 3 out interface.	15.2(3e) and later
CSCuv57315	Atomic counters on the border leaf do not increment for traffic from the Layer 3 out interface to an internal remote endpoint group.	15.2(3e) and later
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	15.2(3e) and later
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	15.2(3e) and later
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.	15.2(3e) 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.	15.2(3e) and later
CSCuy56975	You must configure different router MACs for SVI on each border leaf 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.	15.2(3e) and later
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to "0", and so traffic will still flow.	15.2(3e) and later
CSCuy96912	The debounce timer is not supported on 25G links.	15.2(3e) 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.	15.2(3e) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	15.2(3e) and later

Bug ID	Description	Exists in
CSCuz13810	When subnet check is enabled, a ToR 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 ToRs.	15.2(3e) and later
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	15.2(3e) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) and later
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	15.2(3e) and later
CSCvb49451	In the case of endpoints in two different TOR pairs across a spine switch that are trying to communicate, an endpoint does not get relearned after being deleted on the local TOR pair. However, the endpoint still has its entries on the remote TOR pair.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) and later
CSCvn94400	There is a traffic blackhole that lasts anywhere from a few seconds to a few mins after a border leaf switch is restored.	15.2(3e) and later
CSCvp04772	During an upgrade on a dual-SUP system, the standby SUP may go into a failed state.	15.2(3e) and later
CSCvq71034	There is a policy drop that occurs with L3Out transit cases.	15.2(3e) and later
CSCvr12912	A switch reloads due to a sysmgr heartbeat failure and sysmgr HAP reset.	15.2(3e) and later
CSCvr61096	<p>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.</p> <p>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.</p>	15.2(3e) and later

Bug ID	Description	Exists in
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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) and later
CSCvx75128	Redirected traffic coming back from a service device behind an L3Out is dropped on the second leg.	15.2(3e) and later
CSCvy06135	The leaf switch techsupport with a specified time range fails when the space "/mnt/ifs/log" gets filled up by more than 80%.	15.2(3e) 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.	15.2(3e) and later
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.	15.2(3e) 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 iVLAN packet. You can also embed CoS by enabling CoS preservation. For more information, see the Cisco APIC and QoS KB article.	15.2(3e) 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)	15.2(3e) and later

Bug ID	Description	Exists in
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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) and later
N/A	Multicast router functionality is not supported when IGMP queries are received with VxLAN encapsulation.	15.2(3e) and later
N/A	IGMP Querier election across multiple Endpoint Groups (EPGs) or Layer 2 outsiders (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.	15.2(3e) 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.	15.2(3e) and later
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, 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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported in combination with static endpoint configurations.	15.2(3e) and later

Bug ID	Description	Exists in
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.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported with dynamic endpoint groups. Only static endpoint groups are supported.	15.2(3e) 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.	15.2(3e) and later
N/A	Dynamic deployment of contracts based on instrImmedcy set to onDemand/lazy not supported; only immediate mode is supported.	15.2(3e) 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.	15.2(3e) 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.	15.2(3e) and later
N/A	Configurations for a virtual IP address can only be /32 or /128 prefix.	15.2(3e) and later
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.	15.2(3e) 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).	15.2(3e) and later
N/A	Learning through GARP will work only in ARP Flood Mode.	15.2(3e) 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.
- 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.
- 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."

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-

Product ID	Compatibility Information
	<p>Site together on the Cisco N9K-9364C switch starting in the 3.2 release.</p> <p>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.</p> <p>1-Gigabit QSA is not supported on ports 1/49-64.</p> <p>This switch supports the following PSUs:</p> <ul style="list-style-type: none"> • NXA-PAC-1200W-PE • NXA-PAC-1200W-PI • N9K-PUV-1200W • NXA-PDC-930W-PE • NXA-PDC-930W-PI
N9K-C9332D-GX2B	<p>The following information applies to this switch:</p> <ul style="list-style-type: none"> • Ports 33 and 34 do not support the following things: <ul style="list-style-type: none"> ◦ 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	<p>This switch has the following limitations:</p> <ul style="list-style-type: none"> • The switch will power down in 2 minutes after the first fan failure. The switch can be powered up only after replacing the failed fan. • 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. • All 4 fans must be operational, otherwise the switch will power down due to a fan policy trigger.
N9K-C93600CD-GX	<p>This switch has the following limitations:</p>

Product ID	Compatibility Information
	<ul style="list-style-type: none"> • 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 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-C9332D-GX2B	<p>The following information applies to this switch:</p> <ul style="list-style-type: none"> • Ports 33 and 34 do not support the following things: <ul style="list-style-type: none"> ◦ 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	<p>Auto-negotiation and forward error correction are not supported when you use this switch is as a leaf switch.</p>
N9K-C93180YC-FX3	<p>The following information applies to this switch:</p> <ul style="list-style-type: none"> • The following ports are not supported: <ul style="list-style-type: none"> ◦ 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

Product ID	Compatibility Information
	<p>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.</p>
N9K-C9336C-FX2	<p>The following information applies to this switch:</p> <ul style="list-style-type: none"> 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: <pre>ifav124-leaf5# cat /sys/kernel/cisco_board_info/hw_change_bits</pre> 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, 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: <ul style="list-style-type: none"> After reloading the Top-of-Rack (ToR) switch When using port optical insertion and removal (OIR) When performing bulk flaps of ports on the ToR switch
N9K-C93240YC-FX2	<p>The following information applies when this switch is configured with port-side intake airflow:</p> <ul style="list-style-type: none"> 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. <p>The following information applies when this switch is configured with port-side exhaust airflow:</p> <ul style="list-style-type: none"> 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. <p>The following information applies regardless of the airflow direction:</p> <ul style="list-style-type: none"> 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	<p>The following information applies to this switch:</p> <ul style="list-style-type: none"> 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,

Product ID	Compatibility Information
	<p>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.</p> <ul style="list-style-type: none"> • 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	<p>This switch supports the following PSUs:</p> <ul style="list-style-type: none"> • NXA-PAC-350W-PI • NXA-PAC-350W-PE • NXA-PAC-1100W-PI • NXA-PAC-1100W-PE <p>The following information applies to this switch:</p> <ul style="list-style-type: none"> • 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	<p>This switch does not support the 10G GLC-T optic.</p>
N9K-C93180YC-FX	<p>The following information applies to this switch:</p> <ul style="list-style-type: none"> • 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	<p>This switch does not support the 10G GLC-T optic.</p>
N9K-C93180YC-EX	<p>The following information applies to this switch:</p> <ul style="list-style-type: none"> • The following FEC modes are not supported on N9K-C93180YC-EX ports 1 through 48 when running in 25G speed: <ul style="list-style-type: none"> ◦ 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

Product ID	Compatibility Information
	<p>copper cables connected.</p> <ul style="list-style-type: none"> 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-C93180LC-EX	<p>This switch has the following limitations:</p> <ul style="list-style-type: none"> 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. See the [Cisco Application Policy Infrastructure Controller Release Notes, Release 5.2\(2\)](#) for policy information.
 - 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
- 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 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.

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.

© 2023-2024 Cisco Systems, Inc. All rights reserved.