Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 16.1(5)

© 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 6.1(5).

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

Date	Description
May 4, 2026	Defect CSCwt09384 was added to the list of Open issues.
March 23, 2026	Defect CSCwt38698 was added to the list of Open issues.
February 27, 2026	Defect CSCws82819 was added to the list of Open issues.
December 15, 2025	Release 16.1(5e) became available.

Supported hardware

Table 1. Modular Spine Switches

Product ID	Description
N9K-C9408	Cisco Nexus 9408 modular chassis switch with up to 128 200/100-Gigabit (256 100-Gigabit by 200G-to-2x100G breakout) ports using N9K-X9400-16W or 64 400/200/100-Gigabit (256 100-Gigabit by 400G-to-4x100G breakout) ports using N9K-X9400-8D.
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 9408	Cisco Nexus 9504	Cisco Nexus 9508	Cisco Nexus 9516
N9K-X9400-8D	Cisco Nexus 9400 8-port 400 Gigabit QSFP-DD linecard expansion module	8	N/A	N/A	N/A
N9K-X9400-16W	Cisco Nexus 9400 16-port 200 Gigabit linecard expansion module	8	N/A	N/A	N/A
N9K-X9716D-GX	Cisco Nexus 9500 16-port 400 Gigabit Ethernet QSFP line card	N/A	4	8	N/A
N9K-X9736C-FX	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet Cloud Scale line card	N/A	4	8	16
N9K-X9736Q-FX	Cisco Nexus 9500 36-port 40 Gigabit Ethernet Cloud Scale line card	N/A	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.	N/A	4	8	16
N9K-X9736C-FX3	Cisco Nexus 9500 36-port 100-Gigabit Ethernet QSFP28 line card	N/A	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

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

Cisco Confidential

Product ID	Description
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-C9400-SUP-A	Cisco Nexus 9400 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-H2R	Cisco Nexus 9300 platform switch, which is a deep-buffer, 32-port 400G fixed switch.
N9k-C9364C-H1	Cisco Nexus 9300 platform switch, which is a 64-port 100G fixed switch.
N9K-C9364D-GX2A	Cisco Nexus 9300 platform switch with 64 400/100-Gigabit QSFP-DD ports and 2 1/10 SFP+ ports.
N9K-C9348D-GX2A	Cisco Nexus 9300 platform switch with 48 400/100-Gigabit QSFP-DD ports and 2 1/10 SFP+ ports.
N9K-C9332D-GX2B	Cisco Nexus 9300 platform switch with 32p 400/100-Gigabit 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).
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 10/100/1000BASE-T 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.

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

Cisco Confidential

Table 7. Fixed Spine Switch Power Supply Units

Product ID	Description
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
NXA-PAC-1500W-PE	1500W AC power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance
NXA-PAC-1500W-PI	1500W AC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance
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-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-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

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

Cisco Confidential

Product ID	Description
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
N9K-C9400-FAN-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-35CFM-PE	Blue port side exhaust fan
NXA-FAN-35CFM-PI	Burgundy port side intake fan

Table 9. Modular Leaf Switches

Product ID	Description
N9K-C9408	Cisco Nexus 9408 modular chassis switch with up to 128 200/100-Gigabit (256 100-Gigabit by 200G-to-2x100G breakout) ports using N9K-X9400-16W or 64 400/200/100-Gigabit (256 100-Gigabit by 400G-to-4x100G breakout) ports using N9K-X9400-8D or N9K-X9400-22L (10/25/50G) in leaf-only mode.

Table 10. Modular Leaf Switch Line Cards

Product ID	Description
N9K-X9400-8D	Cisco Nexus 9400 8-port 400 Gigabit QSFP-DD linecard expansion module
N9K-X9400-16W	Cisco Nexus 9400 16-port 200 Gigabit linecard expansion module
N9K-X9400-22L	Cisco Nexus 9400 22-port 10/25/50 Gigabit linecard expansion module

Table 11. Modular Leaf Switch Supervisor and System Controller Modules

Product ID	Description
N9K-C9400-SUP-A	Cisco Nexus 9400 Series supervisor module

Table 12. Fixed Leaf Switches

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

Product ID	Description
N9K-C9332D-H2R	Cisco Nexus 9300 platform switch, which is a deep-buffer, 32-port 400G fixed switch.
N9k-C9364C-H1	Cisco Nexus 9300 platform switch, which is a 64-port 100G fixed switch.
N9K-C93400LD-H1	Cisco Nexus 9300 platform switch, which is a 48-port 50G, 4-port 400G fixed switch.
N9K-C9364D-GX2A	Cisco Nexus 9300 platform switch with 64 400/100-Gigabit QSFP-DD ports and 2 1/10 SFP+ ports.
N9K-C9348D-GX2A	Cisco Nexus 9300 platform switch with 48 400/100-Gigabit QSFP-DD ports and 2 1/10 SFP+ ports.
N9K-C9332D-GX2B	Cisco Nexus 9300 platform switch with 32p 400/100-Gigabit 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 10/100/1000BASE-T 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-C93180YC-FX3H	Cisco Nexus 9300 platform switch with 24 100M/1/10/25-Gigabit Ethernet SFP28 ports, 6 40/100-Gigabit QSFP28 ports, one management port (10/100/1000BASE-T), one console port (RS-232), and one USB port.
N9K-C93108TC-FX3H	Cisco Nexus 9300 platform switch with 24 100M/1/10-GBASE-T (copper) 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 and one SFP+), one console port (RS-232), and one USB port.
N9K-C93108TC-FX3	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 and one SFP+), one console port (RS-232), and one USB port.
N9K-C9348GC-FX3	Cisco Nexus 9300 platform switch with 48 100M/1-GBASE-T (copper) ports, 4 1/10/25-Gigabit SFP28 ports, two 40/100G QSFP28 ports, two management ports (one 10/100/1000BASE-T and one SFP+), 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

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

Product ID	Description
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 9336C-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 9336C-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 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-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.

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

Cisco Confidential

Product ID	Description
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.

Table 13. 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-1500W-PE	1500W AC power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance
NXA-PAC-1500W-PI	1500W AC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance
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

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

Product ID	Description
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
NXA-PDC-715W-PI	715W DC power supply, port side intake pluggable Note: This power supply is supported only by the Cisco Nexus 93108TC-FX3P ACI-mode switch.
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 and 9348GC-FX3 ACI-mode switches.
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 and 9348GC-FX3 ACI-mode switches.
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

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

Cisco Confidential

Product ID	Description
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 14. 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
N9K-C9400-FAN-PI	Burgundy port side intake 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-SFAN-30CFM-PE	Blue port side exhaust fan
NXA-SFAN-30CFM-PI	Burgundy port side intake fan
NXA-FAN-35CFM-PE	Blue port side exhaust fan
NXA-SFAN-35CFM-PE	Blue port side exhaust fan
NXA-FAN-35CFM-PI	Burgundy port side intake fan
NXA-SFAN-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

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

Cisco Confidential

No longer supported hardware

The following hardware is not supported:

Product Type	Product ID
Spine switches	• N9K-C9336PQ
Modular spine switch line cards	• N9K-X9736PQ
Modular spine switch fabric modules	• N9K-C9504-FM • N9K-C9508-FM • N9K-C9516-FM
Leaf Switches	• N9K-C93180LC-EX • N9K-C93120TX • N9K-C93128TX • N9K-C9332PQ • N9K-C9372PX • N9K-C9372PX-E • N9K-C9372TX • N9K-C9372TX-E • N9K-C9396PX • N9K-C9396TX
Expansion Modules	• N9K-M12PQ • N9K-M6PQ • N9K-M6PQ-E
Fabric Extenders	• N2K-C2332TQ-10GT • N2K-C2348TQ-10GE • N2K-C2232PP-10GE • N2K-C2232TM-E-10GE • N2K-C2348TQ-10G-E
Supervisors	• N9K-SUP-A • N9K-SUP-B

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.

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

Cisco Confidential

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

Release overview

Cisco ACI version 16.1(5) is a maintenance release that prioritizes system stability through critical bug fixes and code hardening. This release follows Cisco's established maintenance release cadence and is recommended for customers seeking enhanced reliability and stability in their ACI deployments.

Vital support enhancements

Essential improvements have been incorporated to address specific customer deployment requirements.

Hardware support additions:

- QDD-2X100-LR4-S optics support on N9K-C9332D-GX2B platform (downlink only).
- N9K-C93180YC-FX3 switch allows all downlink ports to be converted as uplink.

  For APIC-related enhancements, see the Cisco Application Policy Infrastructure Controller Release Notes, Release 6.1(5).

Changes in behavior

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

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
CSCwq44924	Failure for switch to forward DHCP offer in second stage relay due to stale entry in kic_database.	16.1(5e)
CSCwq46899	Switch will crash due to PIM HAP reset. Core will be generated.	16.1(5e)
CSCwq58400	Id to dst-ip mapping will have different dest ip address in below 2 commands. 1. moquery -c tunnelIf | egrep " id |dest" 2. Spine(vsh)# show interface | grep " Tunnel[0-300]|source|Tunnel[0-9]"	16.1(5e)
CSCwr43590	Memory corruption in core due to buffer overflow.	16.1(5e)

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

Cisco Confidential

Bug ID	Description	Fixed in
CSCwr44994	Leaf switch reboots with Fan policy trigger.	16.1(5e)
CSCwr58680	Unexpected reloads right after adding a couple of SNMP client and trap destinations to ACI fabric.	16.1(5e)
CSCwr73418	Endpoints are not found from ACI spine's COOP database, which results in proxy traffic drop. Now when the EP requests come from the leaf to the spine via COOP, it is forwarded to the other spine and then back to same spine as both the spines think that the other one is the shard owner for that EP.	16.1(5e)
CSCwr83306	Traffic loss occurs due to a stale remote IP (XR) endpoint entry in the Endpoint Table. The remote IP (XR) tunnel entry in the EP table still points to the endpoint's original location, rather than its new (current) location. As a result, traffic destined for the remote endpoint is lost instead of being redirected to the intended destination.	16.1(5e)
CSCwr91125	Defect required to track investigation and provide a fix for control-plane memory leak on ACI leaf nodes running Kwai 5.2(6x). The issue can result in unexpected leaf reloads due to OOM kernel panic.	16.1(5e)
CSCws28239	ICMPv6 Packet Too Big" packet has MTU value inside (Let's say call it "Too Big MTU" here). In case leaf has received "ICMPv6 Packet Too Big" from an external router, the leaf keeps "Too Big MTU" into kernel memory until it is reloaded. After that, the leaf cannot handle MTU properly in the kernel level. Under some condition with "Too Big MTU", the leaf could drop BGP UPDATE message which is sent by the leaf itself and which is bigger than "Too Big MTU". This leads to BGP neighbor flap.	16.1(5e)
CSCwp66868	In a ACI Border Gateway topology, traffic to a remote NXOS EP gets dropped in the ACI leaf due to missing contract/policy rule, when that remote EP is learned on the leaf with the BD pcTag.	16.1(5e)
CSCwk35253	After undeploy and redeploy of templates (FP, FR, TP, L3out, App), one of the port channels is failing to be configured.	16.1(5e)
CSCwq59239	In a multi-pod setup with inband management enabled, after upgrade of the first set of spines, inter-pod traffic gets dropped for some of the EPs local to POD1. The drops occur on the spines of POD3, as the corresponding remote EPs are missing from COOP.	16.1(5e)
CSCwq02977	Remote site services endpoints flapping from time to time raising faults. Can affect PBR traffic if threshold is configured with down action as drop.	16.1(5e)

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 16.1(5) releases in which the bug exists. A bug might also exist in releases other than the 16.1(5) releases.

Bug ID	Description	Exists in
CSCwt38698	SSD Fault F3074 or F3073 is raised on APIC for switch with Micron drive, but the P/E (Attribute 173), RRE (Attribute 1), and GBB (Attribute 5) values do not exceed predefined thresholds.	16.1(5e)

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

Bug ID	Description	Exists in
CSCws82819	N9K-C9408 will experience boot loop with dt_helper process crash.	16.1(5e)
CSCwq38840	Interface in 1G speed intermittent failures observed where interfaces stop receiving (RX) traffic on the copper port. Model: N9K-C93108TC-FX3P.	16.1(5e)
CSCws40839	In case of VxLAN Gateway between ACI - ACI fabrics, stretched VRF L3 traffic is getting permitted between endpoints even though no contract exists between them (This scenario happens only when BD is not stretched). Even in case of VxLAN gateway between NX - ACI fabrics, similar problem of traffic flowing from NX site to ACI site can be observed when no contracts between endpoints and BD is not stretched.	16.1(5e)
CSCwr42598	There is single interface between IPN and Remote Leaf that is acting as Uplink as well as tenant L3Out, this is an Edge RL feature. There is vpc on back-to-back link between RL1 and RL2 and there is orphan port between RL and host/EPG. The orphan port between RL1 and host/EPG is shut down. Then the native IP packets coming from IPN are reaching RL1, then bounced from RL1 to RL2 but are not forwarded from RL2 to host/EPG during the bounce timer.	16.1(5e)
CSCwj43167	When traffic is sent for non existent EP in different VRF, tglean process CPU utilisation might reach 100%. CPU utilisation will come down as soon as traffic stops.	16.1(5e)
CSCwt09384	SDK HAL core, leaf crash during Ts collection on GX2, H2, H1 model leaf.	16.1(5e)

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 16.1(5) releases in which the bug exists. A bug might also exist in releases other than the 16.1(5) releases.

Bug ID	Description	Exists in
CSCwp17600	EVPN type 4 route RD will still be generated with old router id.	16.1(5e)
CSCvq85886	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.	16.1(5e)
CSCwn52554	N9K-C93320TC-FX3P switch when you connect Copper based QSFP cable, then link may not come up or have link stability issues.	16.1(5e)
CSCvw89840	Traffic originating from a vPC TEP is dropped for Layer 2 multicast and unknown unicast traffic when pod redundancy is triggered.	16.1(5e)
CSCwd89607	When endpoint rogue detection or endpoint loop control is enabled with first hop security, the fabric might flag incorrect endpoint moves. This might lead to loss of traffic or the disabling of bridge domain learning.	16.1(5e)
CSCwf74167	An endpoint does not receive a DHCP response when First-Hop Security (FHS) is enabled.	16.1(5e)
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.	16.1(5e)
CSCup65586	The **show interface** command shows the tunnel's Rx/Tx counters as 0.	16.1(5e)

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

Cisco Confidential

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.	16.1(5e)
CSCup92534	Continuous " threshold exceeded" messages are generated from the fabric.	16.1(5e)
CSCuq39829	Switch rescue user (" admin") can log into fabric switches even when TACACS is selected as the default login realm.	16.1(5e)
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
CSCur81822	The access-port operational status is always " trunk" .	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
CSCuv57302	Atomic counters on the border leaf switch do not increment for traffic from an endpoint group going to the Layer 3 out interface.	16.1(5e)
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.	16.1(5e)
CSCuv57316	TEP counters from the border leaf switch to remote leaf switch nodes do not increment.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
CSCuy06749	Traffic is dropped between two isolated EPGs.	16.1(5e)
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	16.1(5e)

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

Cisco Confidential

Bug ID	Description	Exists in
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to "0", and so traffic will still flow.	16.1(5e)
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.	16.1(5e)
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	16.1(5e)
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.	16.1(5e)
CSCvn94400	There is a traffic blackhole that lasts anywhere from a few seconds to a few mins after a border leaf switch is restored.	16.1(5e)
CSCvp04772	During an upgrade on a dual-SUP system, the standby SUP may go into a failed state.	16.1(5e)
CSCvq71034	There is a policy drop that occurs with L3Out transit cases.	16.1(5e)
CSCvr12912	A switch reloads due to a sysmgr heartbeat failure and sysmgr HAP reset.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
CSCvy06135	The leaf switch techsupport with a specified time range fails when the space "/mnt/ifc/log" gets filled up by more than 80%.	16.1(5e)
CSCvy71586	400G port is automatically broken out into 4 breakout ports. After performing online insertion and removal (OIR) of a 400G transceiver, one of the breakout ports has the "SFP not inserted" or "SFP missing" state.	16.1(5e)
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.	16.1(5e)

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

Cisco Confidential

Bug ID	Description	Exists in
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.	16.1(5e)
CSCwd64518	A virtual machine has connectivity loss when the destination virtual machine is migrated using vMotion. This issue happens only if microsegmentation is enabled on the EPG.	16.1(5e)
CSCwd95467	With N9K-X9400-16W LEM, a pair of odd and even number ports such as port 1/1 and 1/2 must work as the same link type: downlink or fabric link because of CSCwd95467. This consideration is not applicable to N9K-X9400-8D.	16.1(5e)
CSCwe08179	A peer vPC leg goes down after swapping a 16 port LEM with an 8 port LEM. The following error shows in the " show vpc" output: " Peer does not have corresponding vPC" . The leg on the peer switch immediately comes up, but traffic is still disrupted.	16.1(5e)
CSCwe33967	After deleting or adding a VRF instance, the BGP peer session picks up the default timer values instead of the configured values. This is evidenced by the holdIntvl and kaIntvl values in the bgpPeerEntry managed object in the policy engine. The issue happens intermittently.	16.1(5e)
CSCwe41508	As a result of new features, certain PIDs running ACI release 6.1(1) software in 32-bit architecture will see increase in memory consumption and their process virtual address space. This particular issue is seen with a trigger of 500 bridge domain (BD) deletions and addition in a scale configuration of 64k fvrspath scale, 1980 BDs along with 123k policycam entries. In release 6.1(1) with a 32-bit image, process memory could run close to the limit of 4GB." In this scenario, EPM is running at 3.9GB. During the vlan creation as part of the above trigger, EPM attempts to retrieve sclass corresponding to the vlan through DME and DME access is failing. Memory map failures are seen through the instance of EPM. The DME failure may be due to mmap failures.	16.1(5e)
CSCwe97510	When AN On-Enforce is enabled on QDD-4ZQ100G-COPPER breakouts on switches with -GX or -GX2 in the product ID, the links do not come up.	16.1(5e)
CSCwf88389	After an SVI member port flap, ECMP hashing no longer uses the flapped SVI's path and instead uses other SVI paths.	16.1(5e)
CSCwf90351	With the rogue endpoint feature, a MAC address gets flagged as rogue. A leaf switch ignores any further moves of the rogue endpoint for 15 minutes, which can cause an outage. Traffic coming from a FEX vPC carries the Physical Tunnel Endpoint (PTEP) as the source IP address of the outer header (SIPo) instead of the FEX vPC Tunnel Endpoint (TEP).	16.1(5e)
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.	16.1(5e)

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

Bug ID	Description	Exists in
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.	16.1(5e)
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)	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
N/A	Multicast router functionality is not supported when IGMP queries are received with VxLAN encapsulation.	16.1(5e)
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.	16.1(5e)
N/A	The rate of the number of IGMP reports sent to a leaf switch should be limited to 1000 reports per second.	16.1(5e)

© 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.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
N/A	An IP/MAC Ckt endpoint configuration is not supported in combination with static endpoint configurations.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
N/A	An IP/MAC Ckt endpoint configuration is not supported with dynamic endpoint groups. Only static endpoint groups are supported.	16.1(5e)
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.	16.1(5e)
N/A	Dynamic deployment of contracts based on instrImmedcy set to onDemand/lazy not supported; only immediate mode is supported.	16.1(5e)
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.	16.1(5e)
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.	16.1(5e)
N/A	Configurations for a virtual IP address can only be /32 or /128 prefix.	16.1(5e)

© 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.	16.1(5e)
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).	16.1(5e)
N/A	Learning through GARP will work only in ARP Flood Mode.	16.1(5e)

Compatibility information

- For the supported optics per device, see the Cisco Optics-to-Device Compatibility Matrix.
- SFP-10/25G-CSR does not support the Base-R FEC configuration mode in ACI/NXOS. The default FEC mode for SFP-10/25G-CSR is FC-FEC.
- SFP-50G-SR-S, SFP-50G-LR-S, SFP-50G-CUxM are supported on N9K-X9400-22L on N9K-C9408 and N9K-C93400LD-H1.
- QSFP-200G-SL4 and QSFP-100G-SL4 do not support breakout.
- Auto-negotiation is not supported on 200G and 400G optics.
- 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.
- 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.

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

Cisco Confidential

- SFP-10G-T-X cannot be used for APIC connectivity. The APIC GUI provides configuration options for the SFP-10G-T-X transceiver because it is supported for connectivity between leaf switches and endpoints; however, this transceiver is not supported for APIC-to-leaf switch connections.
- 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."
- 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.
- For N9K-C9332D-H2R, one SPAN session is statically reserved due to hardware limitations. Configure only 1 -15 bidirectional fabric SPAN sessions.
- Copper cables are not supported on N9K-X9736C-FX3 (spine linecard).

Table 15. 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	No	No
N9K-X9736C-FX	5	5	5	5	5	5
N9K-X9736Q-FX	5	5	5	5	5	5
N9K-X9732C-EX	4	4	4	4	4	4
N9K-9736C-FX3	5	5	5	5	5	5

Table 16. Modular Spine Switch Line Card Compatibility Information

Product ID	Compatibility Information
N9K-X9400-16W	A port enters the hw-disabled state if you connect a 10G, 40G, or 100G cable next to a port that is connected as 200G-SR. This is because ports next to each other, such as port 1-2, 3-4, or 15-16, share the same MAC address, and only the 200G port comes up between the two.
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 17. Fixed Spine Switches Compatibility Information

Product ID	Compatibility Information
N9K-C9408	This switch has the following limitations: • You cannot use the 2x100G breakout speed in this release. This consideration is applicable to both N9K-X9400-16W and N9K-X9400-8D LEMs. • PTP and SyncE are not supported. • With the N9K-X9400-16W LEM, the pair of port in the same row must be used as either 40/100G or 10G with QSA. For example, if 1/1 is used as 10G with QSA, 1/2 can be used as 10G with QSA but not as 40/100G. If a 40/100G optic is inserted to 1/2, the port becomes hw-disabled. This consideration is not applicable to N9K-X9400-8D LEM.

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

Cisco Confidential

Product ID	Compatibility Information
	• The SFP management port on the supervisor module does not work. • Each chassis supports up to 32 high power optics in total: QDD-400G-ZR-S and QDD-400G-ZRP-S
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: • 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 • This switch supports the following PSUs: ◦ NXA-PAC-1500W-PE ◦ NXA-PAC-1500W-PI ◦ NXA-PDC-1100W-PI • If you use port-side exhaust for this switch, do not overcrowd the switch in the rack. Make sure that you have 2 rack units of empty space above and below the switch.
N9K-C9332D-H2R	Can work as a spine or leaf switch; default is spine switch. This switch has the following limitation: High power optics, 400G-ZR-S/400G-ZRP-S, are supported only on odd ports.

Table 18. Fixed Leaf Switches Compatibility Information

Product ID	Compatibility Information
N9K-C9364C-H1	Can work as a spine or leaf switch; default is leaf switch. This switch has the following limitations:

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

Cisco Confidential

Product ID	Compatibility Information
	• Macsec feature is supported only on the last 16 ports - from 1/49-64. • 10G QSA + 40G+100G can be mixed in a quad; no speed limitation in a quad. • Access Breakout 4X10 and 4X25 are supported only on the first port of every quad (1,5,9,13....). • With breakout on the first port, other 3 ports in a quad are hardware-disabled. • Breakout is not supported on converted downlinks of the last quad on port 1/61 as it will disable fabric port 63-64.
N9K-C93400LD-H1	This switch has the following limitations: • Leaf switch only (not a spine switch). • Does not work with 4+1 redundancy with PE fans.
N9K-C9408	This switch has the following limitations: • You cannot use the 2x100G breakout speed in this release. This consideration is applicable to both N9K-X9400-16W and N9K-X9400-8D LEMs. • Only ports 1 to 6 support port profiles for both the 8D and 16C line-card Ethernet modules (LEMs). • If a port profile is already configured on a LEM and you replace that LEM with a different LEM type, the switch sets the status of the new LEM to " lem-type-mismatch" and Cisco APIC raises the following fault: " Module opertational state changed to LEM type mismatch, please make sure no other lem-type port profile is configured in this slot." To use the new LEM type, you must reload the chassis. • PTP and SyncE are not supported. • With the N9K-X9400-16W LEM, the pair of port in the same row must be used as either 40/100G or 10G with QSA. For example, if 1/1 is used as 10G with QSA, 1/2 can be used as 10G with QSA but not as 40/100G. If a 40/100G optic is inserted to 1/2, the port becomes hw-disabled. This consideration is not applicable to N9K-X9400-8D LEM. • You cannot configure breakout on even ports of the N9K-X9400-16W LEM. • After you configure breakout on an odd port of the N9K-X9400-16W LEM, the next even port will be HW-disabled. • You see the LEM type mismatch status if you swap LEMs when a port profile configuration exists. • The SFP management port on the supervisor module does not work. • Some ARP to gateway packets get dropped with the reason of " ACL_DROP." • Each chassis supports up to 32 high power optics in total: QDD-400G-ZR-S and QDD-400G-ZRP-S • FEXes are not supported.
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 in a quad (1-4,5-8...21-24) However, mixed mode of 40G and 100G in a quad is supported. Based on the port bring up sequence, the port that comes up is retained and the port having speed mismatch 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. • 4X10 and 4X25 breakout is supported on the odd ports 1,3,5,7, etc. • 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:

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

Cisco Confidential

Product ID	Compatibility Information
	• 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 that comes up is retained and the port having speed mismatch 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. • For ports 1 to 24, breakout 4X10 and 4X25 are supported on odd ports. • Ports 25-26 and 27-28 form respective port pairs, and each pair can operate with 4x10, 10G, or 4x25G speed. • 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-X9400-22L	This leaf expansion module has the following limitations: • Link bring up between 50G-SR-S on N9K-X9400-22L and 25G-SR-S to APIC will not come UP. • N9K-X9400-22L has port restrictions wherein a quad of ports will support only 10G/25G or 50G speed. • Remove the N9K-X9400-22L for EWS while downgrading from 6.1.2 release.
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 • This switch supports the following PSUs: ◦ NXA-PAC-1500W-PE ◦ NXA-PAC-1500W-PI ◦ NXA-PDC-1100W-PI

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

Cisco Confidential

Product ID	Compatibility Information
	• If you use port-side exhaust for this switch, do not overcrowd the switch in the rack. Make sure that you have 2 rack units of empty space above and below the switch.
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. • When using this switch as a FEX, QoS stats (as shown by the " show queuing interface ethernet" CLI command) are not supported on the parent Cisco ACI leaf switch. • Auto-negotiation is not supported with 10G speed on ports 1 through 48 and ports 49 through 52 with QSA.
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, 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.

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

Cisco Confidential

Product ID	Compatibility Information
	• 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 • 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

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

Product ID	Compatibility Information
	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.

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 6.1(5) 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

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

    - 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.
- 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.

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

Cisco Confidential

- 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.

© 2026 Cisco Systems, Inc. All rights reserved.

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

Cisco Confidential