Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 14.2(7)

© 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 4.2(7).

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

Note: The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Date	Description
March 17, 2023	Release 14.2(7w) became available.
November 20, 2022	Release 14.2(7v) became available. Added the resolved issues for this release.
July 22, 2022	Release 14.2(7u) became available. Added the resolved issues for this release.
May 16, 2022	In the Open Issues section, added bug CSCwa47686.
April 6, 2022	Release 14.2(7t) became available. Added the resolved issues for this release.
March 16, 2022	In the Open Issues section, added bugs CSCwa95241 and CSCwb14844.
February 14, 2022	Release 14.2(7s) became available. Added the resolved issues for this release.
December 16, 2021	In the Open 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 10, 2021	In the Open Issues section, added bug CSCwa18051.
October 7, 2021	Release 14.2(7q) became available. Added the open and resolved issues for this release.
September 24, 2021	In the Open Issues section, added bug CSCvz71778.
August 18, 2021	In the Compatibility Information section, for the N9K-C9316D-GX, N9K-C93600CD-GX, and N9K-C9364C-GX switches, added: 1G and 100MB speeds are not supported.
July 30, 2021	Release 14.2(7l) became available. Added the resolved issues for this release. Added the

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

Cisco Confidential

Date	Description
	same bugs as open in the 14.2(7f) release.
July 29, 2021	In the Modular Spine Switch Fabric Modules table, for N9K-C9504-FM, N9K-C9508-FM, and N9K-C9516-FM, changed the maximum to 6 and removed " Note: This fabric module is not supported in slot 21 nor 25."
July 6, 2021	In the Supported Hardware section, added the NXA-PAC-500W-PI and NXA-PAC-500W-PE PSUs.
June 24, 2021	In the Open Issues section, removed bug CSCvu07844. This was resolved in the 4.2(6) release.
June 21, 2021	In the Known Issues section, added bug CSCvu42069.
June 15, 2021	In the Open Issues section, added bug CSCvy43640.
April 30, 2021	In the Open Issues section, added bug CSCvy12057.
March 24, 2021	Added resolved issue CSCvx01777.
March 23, 2021	In the Open Issues section, added bug CSCvx70611. In the Resolved Issues section, added bug CSCvw51079.
March 16, 2021	Release 14.2(7f) 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-X9736C-FX	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet Cloud Scale line card	4	8	16

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

Cisco Confidential

Product ID	Description	Maximum Quantity		
		Cisco Nexus 9504	Cisco Nexus 9508	Cisco Nexus 9516
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
N9K-X9736PQ	Cisco Nexus 9500 36-port 40 Gigabit Ethernet line card	4	8	16

Table 3. Modular Spine Switch Fabric Modules

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
N9K-C9504-FM	Cisco Nexus 9504 classic fabric module	3	6
N9K-C9508-FM	Cisco Nexus 9508 classic fabric module	3	6
N9K-C9516-FM	Cisco Nexus 9516 classic fabric module	3	6

Table 4. 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 5. Fixed Spine Switches

Product ID	Description
N9K-C9316D-GX	Cisco Nexus 9300 platform switch with 16 10/40/100/400-Gigabit QSFP-DD ports (ports 1-16).

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

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-C9336PQ	Cisco Nexus 9336PQ switch, 36-port 40 Gigabit Ethernet QSFP
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 6. 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, 93128TX, 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, 93128TX, 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
N9K-PUV-1200W	1200W HVAC/HVDC dual-direction airflow power supply

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

Cisco Confidential

Table 7. 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-C9504-FAN	Fan tray for Cisco Nexus 9504 chassis
N9K-C9508-FAN	Fan tray for Cisco Nexus 9508 chassis
N9K-C9516-FAN	Fan tray for Cisco Nexus 9516 chassis
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 8. Fixed Leaf Switches

Product ID	Description
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 1 USB port.
N9K-C93600CD-GX	Cisco Nexus 93600CD-GX 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-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.
N9K-C9336C-FX2	Cisco Nexus C9336C-FX2 Top-of-rack (ToR) switch with 36 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: 1-Gigabit QSA is not supported on ports 1/1-6 and 1/33-36. The port profile feature supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.
N9K-C93108TC-FX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.

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

Cisco Confidential

Product ID	Description
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-C93180LC-EX	Cisco Nexus 9300 platform switch with 24 40-Gigabit front panel ports and 6 40/100-Gigabit QSFP28 spine-facing ports. The switch can be used as either a 24 40G port switch or a 12 100G port switch. If 100G is connected the Port1, Port 2 will be HW disabled.
N9K-C93180YC-EX	Cisco Nexus 9300 platform switch with 48 1/10/25-Gigabit front panel ports and 6-port 40/100 Gigabit QSFP28 spine-facing ports.
N9K-C93180YC-EX-24	Cisco Nexus 9300 platform switch with 24 1/10/25-Gigabit front panel ports and 6-port 40/100 Gigabit QSFP28 spine-facing ports.
N9K-C9372PX-E	Cisco Nexus 9372PX-E Top-of-rack (ToR) Layer 3 switch with 48 Port 1/10-Gigabit APIC-facing ports Ethernet SFP+ front panel ports and 6 40-Gbps Ethernet QSFP+ spine-facing ports Note: Only the downlink ports 1-16 and 33-48 are capable of supporting SFP1-10G-ZR SFP+.
N9K-C9372TX-E	Cisco Nexus 9372TX-E Top-of-rack (ToR) Layer 3 switch with 48 10GBASE-T (copper) front panel ports and 6 40-Gbps Ethernet QSFP+ spine-facing ports
N9K-C93120TX	Cisco Nexus 9300 platform switch with 96 1/10GBASE-T (copper) front panel ports and 6-port 40-Gigabit Ethernet QSFP spine-facing ports.
N9K-C93128TX	Cisco Nexus 9300 platform switch with 96 1/10GBASE-T (copper) front panel ports and 6 or

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

Product ID	Description
	8 40-Gigabit Ethernet QSFP spine-facing ports.
N9K-C9332PQ	Cisco Nexus 9332PQ Top-of-rack (ToR) Layer 3 switch with 26 APIC-facing ports and 6 fixed-Gigabit spine facing ports.
N9K-C9372PX	Cisco Nexus 9372PX Top-of-rack (ToR) Layer 3 switch with 48 Port 1/10-Gigabit APIC-facing ports Ethernet SFP+ front panel ports and 6 40-Gbps Ethernet QSFP+ spine-facing ports Note: Only the downlink ports 1-16 and 33-48 are capable of supporting SFP1-10G-ZR SFP+.
N9K-C9372TX	Cisco Nexus 9372TX Top-of-rack (ToR) Layer 3 switch with 48 1/10GBASE-T (copper) front panel ports and 6 40-Gbps Ethernet QSFP spine-facing ports
N9K-C9396PX	Cisco Nexus 9300 platform switch with 48 1/10-Gigabit SFP+ front panel ports and 6 or 12 40-Gigabit Ethernet QSFP spine-facing ports
N9K-C9396TX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) front panel ports and 6 or 12 40-Gigabit Ethernet QSFP spine-facing ports

Table 9. Expansion Modules

Product ID	Description
N9K-M12PQ	12-port or 8-port Gigabit Ethernet expansion module
N9K-M6PQ	6-port Gigabit Ethernet expansion module
N9K-M6PQ-E	6-port, 40 Gigabit Ethernet expansion module

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, 93128TX, 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, 93128TX, 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

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

Cisco Confidential

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

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

Cisco Confidential

Product ID	Description
	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

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

Cisco Confidential

Product ID	Description
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

Supported FEX Models

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

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

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

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

New Hardware Features

There are no new hardware features in this release.

New Software Features

For new software features, see the Cisco Application Policy Infrastructure Controller Release Notes, Release 4.2(7).

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 14.2(7) releases in which the bug exists. A bug might also exist in releases other than the 14.2(7) releases.

Bug ID	Description	Exists in
CSCwh78987	Breakout ports configured as port channel members are no longer part of the port channel post clean reload.	14.2(7v)

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

Bug ID	Description	Exists in
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).	14.2(7s) and later
CSCwf04501	There are actrl.Rule programming failures followed by a policy-mgr HAP reset.	14.2(7s) and later
CSCvz74911	A fault will be raised if the number of configured VLANS per port is greater than 256 for the MCP interface. A fault will be raised if total number of operational VLAN ports is greater than 2000 in the system at the MCP inst level.	14.2(7q) and later
CSCwa18051	An EMI fast retrain configuration requires a switch reload after applying the policy.	14.2(7q) and later
CSCwd83289	After reloading a leaf switch, HAL becomes unresponsive, other processes generate a core, and HAL continuously runs with ZMQ processing.	14.2(7q) 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	14.2(7q) and later
CSCvz95984	The value of " Use count" in the " show flow monitor" command output is incorrect.	14.2(7q) and 14.2(7s)
CSCwa18051	An EMI fast retrain configuration requires a switch reload after applying the policy.	14.2(7q)
CSCwa60256	EPLD fails to upgrade on the Cisco N9K-C93360YC-FX2 switch.	14.2(7q)
CSCwa92634	The power supply status flaps between failed and OK.	14.2(7q)
CSCvz85364	A switch reloads unexpectedly due to the " coop hap reset" error.	14.2(7l) through 14.2(7u)
CSCvz40389	Random link flap observed on the BCM-based physical leaf switches where there is EMI noise in the environment.	14.2(7l)
CSCvz71778	After the switch completes the upgrade, the switch will bring up its downlinks without any policy/configuration (VLANs). This can result in hosts hashing to these links and that traffic goes to a null route for a few minutes until the Cisco APIC handshake completes and the Cisco APIC is able to program the policy.	14.2(7l)
CSCvz78464	The first hop router of a Cisco ACI fabric with a rendezvous point inside the fabric will not generate a Protocol Independent Multicast (PIM) register packet and will not send the packet to all border leaf switches that are members of the anycast rendezvous point. As a result, the last hop router will fail to create an (S,G) entry.	14.2(7l)
CSCwb17229	The sysmgr process crashes unexpectedly, causing the line card to reload.	14.2(7f) through 14.2(7t)
CSCwb39899	A Cisco ACI leaf switch reloads with the following reset reason: Reset Reason for this card:	14.2(7f) through

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

Cisco Confidential

Bug ID	Description	Exists in
	Image Version : 14.2(7f) Reset Reason (LCM): Unknown (0) at time Tue Mar 22 13:01:28 2022 Reset Reason (SW): Reset triggered due to HA policy of Reset (16) at time Tue Mar 22 12:56:21 2022 Service (Additional Info): pim hap reset Reset Reason (HW): Reset triggered due to HA policy of Reset (16) at time Tue Mar 22 13:01:28 2022 Reset Cause (HW): 0x01 at time Tue Mar 22 13:01:28 2022 Reset internal (HW): 0x00 at time Tue Mar 22 13:01:28 2022	14.2(7t)
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.	14.2(7f) through 14.2(7s)
CSCvu58225	The NTP clock is not in synchronized between the supervisor and line card.	14.2(7f) through 14.2(7q)
CSCvw16997	The QSFP-100G-CWDM4-MSA-FEC interface flaps with a Cisco N9K-X9716D-GX line card after the line card reloads.	14.2(7f) through 14.2(7q)
CSCwa18165	Ether type 0x3737 is dropped by FX switches as ACL_DROP.	14.2(7f) through 14.2(7q)
CSCvw91752	Fault F0411 keeps on being raised, and it shows that PSUs on the chassis keep on failing and recovering. This issue is cosmetic and has no effect on the switch or traffic.	14.2(7f) through 14.2(7l)
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.	14.2(7f) and later
CSCvh11299	In COOP, the MAC IP address route has the wrong VNID, and endpoints are missing from the IP address DB of COOP.	14.2(7f) and later
CSCvp09949	Copy service traffic will fail to reach the TEP where the copy devices are connected. Traffic will not be seen on the spine switches.	14.2(7f) and later
CSCvs86972	Remote leaf switches and spine switches cannot be connected to from an external virtual machine.	14.2(7f) and later
CSCvt07021	A remote leaf switch or vPod is inactive after the deletion of the routable pool.	14.2(7f) and later
CSCvt16711	SSH cannot be used to connect from APIC to the leaf/spine switches using inband management and with the indband VRF table in enforced mode.	14.2(7f) and later
CSCvt73069	A Cisco ACI fabric is not fully fit after a Cisco APIC firmware upgrade.	14.2(7f) and later
CSCvt77359	SSH from an external virtual machine to the spine switches does not work due to the actrlMgmtRule rule not being created for the spine switch with " vzany cons for INB_VRF and L3out is prov" . SSH from an external virtual machine to a leaf switch is works.	14.2(7f) and later

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

Cisco Confidential

Bug ID	Description	Exists in
CSCvu08653	SSH from an external virtual machine does not work due to the actrlMgmtRule rule not being created.	14.2(7f) and later
CSCvu77935	Applications are slow when deployed in servers that are connected to a Tier-1 leaf switch.	14.2(7f) and later
CSCvv04106	Traffic classification is not correct in the sub-leaf switch (for the traffic coming from the mid leaf switch) when the Cisco ACI Multi-Pod COS-DSCP translation policy is enabled in the fabric.	14.2(7f) and later
CSCvy30381	After replacing the hardware for a leaf switch, the leaf switch front-panel ports are set to the admin-down state for 45 minutes.	14.2(7f) and later
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.	14.2(7f) and later
CSCwa12763	External route import for a VRF instance fails on a leaf switch after removing a shared services contract between two EPGs.	14.2(7f) 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.	14.2(7f) and later
CSCwa95241	An endpoint may fail to resolve an ARP request of another endpoint.	14.2(7f) 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.	14.2(7f) and later
CSCwd29346	An ACI switch's console may continuously output messages similar to: svc_ifc_eventmg (*****) Ran 7911 msecs in last 7924 msecs	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later

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

Bug ID	Description	Exists in
CSCwd75707	The parser drops IP-in-IP packets that are associated with IP traffic on ingress leaf switches.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later
CSCvx65787	PBR may not be applied at the provider leaf switch if an XR IP address or remote IP address endpoint gets programmed with sclass 1. This could happen as a result of a timing issue exposed by receiving a COOP bounce for an endpoint that is already in the bounced state.	14.2(7f) and 14.2(7l)
CSCvy69714	In a Cisco ACI Multi-Pod setup, a remote endpoint is available through an L2VPN EVPN BGP from a pod1 spine switch to a pod2 spine switch. However, on the pod2 spine switch, COOP is not updated with this endpoint from BGP.	14.2(7f) and 14.2(7l)
CSCvz22610	Routes are learned through an L3Out that are not exported to the other Cisco ACI Multi-Site fabric. However, a local fabric imported the routes from the other site.	14.2(7f) and 14.2(7l)
CSCvz64029	The following alert is generated in the Cisco APIC's GUI: The process lacp on Node <NodeId> stopped at <Timestamp> - more than <X> hours ago. Previously, it stopped at <Timestamp>, more than <X> hours ago, and <Y> other times before that. Clear this alert after lacp restarts successfully on this node.	14.2(7f) and 14.2(7l)
CSCvz65459	Multiple spine switches crash due to a kernel panic.	14.2(7f) and 14.2(7l)
CSCvq57414	HSRP/VRRP packets failed to flood locally in a service leaf switch, which causes a dual active state.	14.2(7f)
CSCvt15100	The logs are missing or have incomplete show techsupport output due to the /var/sysmgr/mem_logs partition being full. In an extreme case, the system may experience abnormal behaviour or crashes.	14.2(7f)
CSCvv63073	Any of the multicast modules (MRIB/IGMP/MCASTFWD/PIM/NGMVPN) generate a core, indicating a heartbeat failure.	14.2(7f)
CSCvv94004	An N9K-C9504-FM-E fabric module in an N9K-C9504 modular spine switch may fail.	14.2(7f)
CSCvx04772	The flow telemetry packets will not reach the collector and NI might not show the spine interfaces.	14.2(7f)
CSCvx41386	There is an inability to communicate with endpoints within the same bridge domain. When checking the endpoint MAC address on some leaf switches, there is a remote MAC endpoint, but the tunnel on which the endpoint is learned is not the tunnel to the leaf switches where the MAC address is connected locally.	14.2(7f)
CSCvx44791	In a vPC setup, after one of the leaf switches was upgraded or clean reloaded, the DSR VIP address was deleted from COOP and traffic to the VIP address was dropped by the spine switch proxy.	14.2(7f)
CSCvx45020	An endpoint is unable to reach its default gateway after a vMotion.	14.2(7f)

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

Cisco Confidential

Bug ID	Description	Exists in
CSCvx49448	When using OSPF HELLO timers set to 1 second, with a dead interval of 3 seconds, intermittently the OSPF adjacency will flap across multiple neighbors and VRF instances at the same time.	14.2(7f)
CSCvx61156	The LACP ports come up before all VLAN ports are provisioned.	14.2(7f)
CSCvx70611	ARP requests that should be flooded in encapsulation are instead flooded across encapsulations on the border leaf switches.	14.2(7f)
CSCvx70611	ARP requests that should be flooded in encapsulation are instead flooded across encapsulations on the border leaf switches.	14.2(7f)
CSCvx70611	ARP requests that should be flooded in encapsulation are instead flooded across encapsulations on the border leaf switches.	14.2(7f)
CSCvx76219	A tunnel connects to random IP addresses that do not exist in the ISIS table, which results in the following fault: F0475: Tunnel destination is not reachable	14.2(7f)
CSCvx82486	The mcastfwd module crashes and the switch reboots unexpectedly.	14.2(7f)
CSCvx84820	The following fault is generated: [Shard 32] failed to apply tree: SLA TCP port cannot be 0	14.2(7f)
CSCvx88519	The Cisco APIC raises fault " F2728 Latency feature is Inactive because PTP is not operational" . The affected switch displays a PTP Source IP Address of 0.0.0.0.	14.2(7f)
CSCvx93880	In a Cisco ACI Multi-Site setup, if the DHCP server and client on different sites, then the DHCP offer gets dropped on the modular spine switch due to a VLAN translate miss.	14.2(7f)
CSCvy01336	Traffic is sent to a device that does not match the expected device, and the symmetric hash is broken.	14.2(7f)
CSCvy07331	A leaked static route of the DHCP relay is deleted from a VRF instance.	14.2(7f)
CSCvy12057	After upgrading to the 14.2(6) or later release, if you boot from a SAN with a vPC configuration, then the Virtual Fiber Channel (VFC) interfaces associated with the member interfaces remain down until the port channel comes up. This results in errors on the end hosts when they are rebooted.	14.2(7f)
CSCvy12057	After upgrading to the 14.2(6) or later release, if you boot from a SAN with a vPC configuration, then the Virtual Fiber Channel (VFC) interfaces associated with the member interfaces remain down until the port channel comes up. This results in errors on the end hosts when they are rebooted.	14.2(7f)
CSCvy12057	SAN Boot with a vPC configuration will have issues after upgrading to the 4.2(6) or later release. The VFC interfaces associated with the member interfaces are down until the port channel comes up, and this results in errors on the end hosts when they are rebooted.	14.2(7f)
CSCvy13313	A Cisco ACI fabric switch reloads unexpectedly due to the NFM process initiating a HAP reset. This issue is caused by a heartbeat failure that is caused by the NI app. The issue is fixed in the 5.1 release of the app.	14.2(7f)

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

Cisco Confidential

Bug ID	Description	Exists in
CSCvy15585	MTS buffer in use build up consistently causes control-plane instability.	14.2(7f)
CSCvy17518	Cisco ACI leaf switch kernel panic due to the node process becoming out-of-memory.	14.2(7f)
CSCvy19681	Flow exports from the spine switch will not reach the collector. Spine switch interfaces might be missing in the packet flow path in NI.	14.2(7f)
CSCvy22243	A Cisco Nexus 2300 series FEX does not come online on Cisco Nexus 9000 switches or Cisco ACI leaf switches when a transceiver of type 40G-SR-BD QSFP is used.	14.2(7f)
CSCvy27363	After a stateful reload, some of the BFD sessions do not come up. The BFDC logs show that the port down event was pushed by BFD, but the port up event was not pushed, which causes sessions to stay down. The tcpdump command shows that there is no BFD control packet being sent.	14.2(7f)
CSCvy43728	After downgrading a Cisco ACI leaf or spine switch from a 4.2 release to a 3.2 release, you may notice that the switch becomes "active" in the fabric (acidiag fnvread), but the node is missing many policies allowing it to function properly, such as the BGP Route Reflector policies.	14.2(7f)
CSCvy46638	A Cisco 9336C-FX2 switch reboots with the following reason: reset-triggered-due-to-ha-policy-of-reset System version: 14.2(7f) Service: nfm hap reset	14.2(7f)
CSCvy48603	Telemetry packets are not exported to the collector.	14.2(7f)
CSCvy80235	There is intermittent flapping on a copper-based switch.	14.2(7f)
CSCvy81682	IGMP generates a core on the deletion of a tenant.	14.2(7f)
CSCvz07799	There might be a scenario when traffic between pods is blackholed when all Spines are reloaded/upgraded in a timeframe of less than 30 minutes.	14.2(7f)
CSCvz08565	IGMP generates a core due to memory corruption or a hearbeat failure. In addition, Mcastfwd generates a core due to a heartbeat failure.	14.2(7f)
CSCvz09521	PIM generates a core while collecting tech support or executing the 'show ip pim interface vrf all' command.	14.2(7f)
CSCvz74608	COOP session disconnection is reported for one of the leaf switches and fault F1360 is generated. The COOP packet can be dropped on the leaf switch and that can trigger TCP session reset and subsequent COOP session disconnection.	14.2(7f)
CSCwh72876	The EPM process crashed when there was no disk space was available at /var/sysmgr/tmp_logs/.	14.2(7f)

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.

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

Bug ID	Description	Fixed in
CSCvz81692	A vPC domain's interfaces are suspended/down when its peer node is decommissioned.	14.2(7v)
CSCvz85364	A switch reloads unexpectedly due to the " coop hap reset" error.	14.2(7v)
CSCwb17229	The sysmgr process crashes unexpectedly, causing the line card to reload.	14.2(7u)
CSCwb39899	A Cisco ACI leaf switch will reload with the following reset reason: Reset Reason for this card: Image Version : 14.2(7f) Reset Reason (LCM): Unknown (0) at time Tue Mar 22 13:01:28 2022 Reset Reason (SW): Reset triggered due to HA policy of Reset (16) at time Tue Mar 22 12:56:21 2022 Service (Additional Info): pim hap reset Reset Reason (HW): Reset triggered due to HA policy of Reset (16) at time Tue Mar 22 13:01:28 2022 Reset Cause (HW): 0x01 at time Tue Mar 22 13:01:28 2022 Reset internal (HW): 0x00 at time Tue Mar 22 13:01:28 2022	14.2(7u)
CSCvy28500	The vntag mgr process has a memory leak, which slows the process and results in the process crashing.	14.2(7t)
CSCvz95984	The value of " Use count" in the " show flow monitor" command output is incorrect.	14.2(7t)
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.	14.2(7t)
CSCvu58225	The NTP clock is not in synchronized between the supervisor and line card.	14.2(7s)
CSCvw16997	The QSFP-100G-CWDM4-MSA-FEC interface flaps with a Cisco N9K-X9716D-GX line card after the line card reloads.	14.2(7s)
CSCwa18051	An EMI fast retrain configuration requires a switch reload after applying the policy.	14.2(7s)
CSCwa18165	Ether type 0x3737 is dropped by FX switches as ACL_DROP.	14.2(7s)
CSCwa60256	EPLD fails to upgrade on the Cisco N9K-C93360YC-FX2 switch.	14.2(7s)
CSCwa92634	The power supply status flaps between failed and OK	14.2(7s)
CSCvw07625	The port security feature is configured, but does not take effect after the policy is applied under a vPC.	14.2(7q)
CSCvw91752	Fault F0411 keeps on being raised, and it shows that PSUs on the chassis keep on failing and recovering. This issue is cosmetic and has no effect on the switch or traffic.	14.2(7q)
CSCvx29134	There are events in Cisco ACI for fans being removed from and reinserted into the switches.	14.2(7q)
CSCvx45118	The ethpm process memory leaks on a leaf switch.	14.2(7q)

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

Cisco Confidential

Bug ID	Description	Fixed in
CSCvx65787	PBR may not be applied at the provider leaf switch if an XR IP address or remote IP address endpoint gets programmed with sclass 1. This could happen as a result of a timing issue exposed by receiving a COOP bounce for an endpoint that is already in the bounced state.	14.2(7q)
CSCvx76768	A static endpoint is configured on a vPC port channel (say Lea1-Leaf2 vPC). With the vPC leg on leaf2 down, if leaf1 is reloaded, the static endpoint is deleted in COOP as both of the legs are down. When leaf1 comes up, the switch restores the static endpoint, but does not update COOP post-restore, which causes the endpoint not to be added back in the COOP spine switch.	14.2(7q)
CSCvx83364	Known unicast packets received in a flood in encapsulation EPG get flooded in the bridge domain (in another leaf switch), which is not expected.	14.2(7q)
CSCvx86858	When a unicast ARP/GARP reply is received on a front panel port in an EPG with FIE enabled at the bridge domain/EPG, this can cause the reply to be flooded to other EPGs in the bridge domain.	14.2(7q)
CSCvy07418	A Cisco ACI-mode switch drops unicast proxy ARP requests received on FIE EPGs. This forces the client host to do retries and fall back to broadcast the proxy ARP request, which delays the ARP refresh.	14.2(7q)
CSCvy69104	If an endpoint exists as a dynamic endpoint and then gets configured as a Layer 4 to Layer 7 VIP address, EPM flushes the dynamic entry and waits for a new ARP to reprogram it as a static endpoint. The flush in EPM removes the entry in COOP, so endpoints communicating to the previous instance of the Layer 4 to Layer 7 VIP address could see a convergence issue until the next ARP is received.	14.2(7q)
CSCvy69714	In a Cisco ACI Multi-Pod setup, a remote endpoint is available through an L2VPN EVPN BGP from a pod1 spine switch to a pod2 spine switch. However, on the pod2 spine switch, COOP is not updated with this endpoint from BGP.	14.2(7q)
CSCvy82391	Flooded traffic such as ARP is dropped on transit FTAG leaf switches after the 4.2(6d) release. This issue appears to be a result of a code change due to issue CSCvx83364, which modified the code.	14.2(7q)
CSCvy90645	Multiple IPv6 subnets are configured under the bridge domain. When we change the IPv6 subnet from preferred (primary) to secondary or vice versa, the IPv6 subnet gets deleted from EPM and EPMC.	14.2(7q)
CSCvz12568	During live migration, some virtual machines receive a message regarding IPv6 duplicate address detection.	14.2(7q)
CSCvz21588	When Layer 2 data is received in a leaf switch, the tunnel between the ingress leaf switch and egress leaf switch is not created by default. A glean ACL takes care of punting the packet to the CPU that triggers tunnel creation. But, the original packet is forwarded with BDVNID, causing the packet to leak to other EPGs in the bridge domain on the other leaf switches.	14.2(7q)
CSCvz22610	Routes are learned through an L3Out that are not exported to the other Cisco ACI Multi-Site fabric. However, a local fabric imported the routes from the other site.	14.2(7q)
CSCvz27930	Endpoints may receive packets routed from a fabric with TTL0. TTL 1 packets arrive on a leaf switch and are meant to be routed. Some of those packets are redirected to the supervisor module by CoPP class for TTL exception (conformed bytes), and non-conformed packets are sent out of an egress downlink port with TTL0 instead of being dropped.	14.2(7q)

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

Cisco Confidential

Bug ID	Description	Fixed in
CSCvz40389	Random link flap observed on the BCM-based physical leaf switches where there is EMI noise in the environment.	14.2(7q)
CSCvz40909	An SSD-related fault (F3074 or F3073) is falsely raised for a switch SSD in the Cisco APIC GUI.	14.2(7q)
CSCvz47876	A Cisco ACI-mode switch drops unicast proxy ARP requests received on FIE EPGs. This forces the client host to do retries and fall back to broadcast the proxy ARP request, which delays the ARP refresh.	14.2(7q)
CSCvz55417	Fault F3525 might be observed for a PTP switch in a Cisco ACI fabric.	14.2(7q)
CSCvz61945	All switches in the fabric continuously reload with the reset-reason of reset-triggered-due-to-ha-policy-of-reset. Affected services include etlmc, ptplc, ipfib, sdkhal, epmc, and aclqos.	14.2(7q)
CSCvz64029	The following alert is generated in the Cisco APIC's GUI: The process lacp on Node <NodeId> stopped at <Timestamp> - more than <X> hours ago. Previously, it stopped at <Timestamp>, more than <X> hours ago, and <Y> other times before that. Clear this alert after lacp restarts successfully on this node.	14.2(7q)
CSCvz65459	Multiple spine switches crash due to a kernel panic.	14.2(7q)
CSCvz71778	After the switch completes the upgrade, the switch will bring up its downlinks without any policy/configuration (VLANs). This can result in hosts hashing to these links and that traffic goes to a null route for a few minutes until the Cisco APIC handshake completes and the Cisco APIC is able to program the policy.	14.2(7q)
CSCvz74608	COOP session disconnection is reported for one of the leaf switches and fault F1360 is generated. The COOP packet can be dropped on the leaf switch and that can trigger TCP session reset and subsequent COOP session disconnection.	14.2(7q)
CSCvz78464	The first hop router of a Cisco ACI fabric with a rendezvous point inside the fabric will not generate a Protocol Independent Multicast (PIM) register packet and will not send the packet to all border leaf switches that are members of the anycast rendezvous point. As a result, the last hop router will fail to create an (S,G) entry.	14.2(7q)
CSCvq57414	HSRP/VRRP packets failed to flood locally in a service leaf switch, which causes a dual active state.	14.2(7l)
CSCvt15100	The logs are missing or have incomplete show techsupport output due to the /var/sysmgr/mem_logs partition being full. In an extreme case, the system may experience abnormal behaviour or crashes.	14.2(7l)
CSCvv63073	Any of the multicast modules (MRIB/IGMP/MCASTFWD/PIM/NGMVPN) generate a core, indicating a heartbeat failure.	14.2(7l)
CSCvv94004	An N9K-C9504-FM-E fabric module in an N9K-C9504 modular spine switch may fail.	14.2(7l)
CSCvx04772	The flow telemetry packets will not reach the collector and NI might not show the spine interfaces.	14.2(7l)
CSCvx41386	There is an inability to communicate with endpoints within the same bridge domain. When checking the endpoint MAC address on some leaf switches, there is a remote MAC endpoint, but the tunnel on which the endpoint is learned is not the tunnel to the leaf switches where the MAC address is connected locally.	14.2(7l)

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

Bug ID	Description	Fixed in
CSCvx44791	In a vPC setup, after one of the leaf switches was upgraded or clean reloaded, the DSR VIP address was deleted from COOP and traffic to the VIP address was dropped by the spine switch proxy.	14.2(7l)
CSCvx45020	An endpoint is unable to reach its default gateway after a vMotion.	14.2(7l)
CSCvx49448	When using OSPF HELLO timers set to 1 second, with a dead interval of 3 seconds, intermittently the OSPF adjacency will flap across multiple neighbors and VRF instances at the same time.	14.2(7l)
CSCvx70611	ARP requests that should be flooded in encapsulation are instead flooded across encapsulations on the border leaf switches.	14.2(7l)
CSCvx70611	ARP requests that should be flooded in encapsulation are instead flooded across encapsulations on the border leaf switches.	14.2(7l)
CSCvx76219	A tunnel connects to random IP addresses that do not exist in the ISIS table, which results in the following fault: F0475: Tunnel destination is not reachable	14.2(7l)
CSCvx82486	The mcastfwd module crashes and the switch reboots unexpectedly.	14.2(7l)
CSCvx84820	The following fault is generated: [Shard 32] failed to apply tree: SLA TCP port cannot be 0	14.2(7l)
CSCvx88519	The Cisco APIC raises fault " F2728 Latency feature is Inactive because PTP is not operational" . The affected switch displays a PTP Source IP Address of 0.0.0.0.	14.2(7l)
CSCvx93880	In a Cisco ACI Multi-Site setup, if the DHCP server and client on different sites, then the DHCP offer gets dropped on the modular spine switch due to a VLAN translate miss.	14.2(7l)
CSCvy01336	Traffic is sent to a device that does not match the expected device, and the symmetric hash is broken.	14.2(7l)
CSCvy07331	A leaked static route of the DHCP relay is deleted from a VRF instance.	14.2(7l)
CSCvy12057	After upgrading to the 14.2(6) or later release, if you boot from a SAN with a vPC configuration, then the Virtual Fiber Channel (VFC) interfaces associated with the member interfaces remain down until the port channel comes up. This results in errors on the end hosts when they are rebooted.	14.2(7l)
CSCvy12057	SAN Boot with a vPC configuration will have issues after upgrading to the 4.2(6) or later release. The VFC interfaces associated with the member interfaces are down until the port channel comes up, and this results in errors on the end hosts when they are rebooted.	14.2(7l)
CSCvy13313	A Cisco ACI fabric switch reloads unexpectedly due to the NFM process initiating a HAP reset. This issue is caused by a heartbeat failure that is caused by the NI app. The issue is fixed in the 5.1 release of the app.	14.2(7l)
CSCvy15585	MTS buffer in use build up consistently causes control-plane instability.	14.2(7l)
CSCvy17518	Cisco ACI leaf switch kernel panic due to the node process becoming out-of-memory.	14.2(7l)

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

Cisco Confidential

Bug ID	Description	Fixed in
CSCvy19681	Flow exports from the spine switch will not reach the collector. Spine switch interfaces might be missing in the packet flow path in NI.	14.2(7l)
CSCvy22243	A Cisco Nexus 2300 series FEX does not come online on Cisco Nexus 9000 switches or Cisco ACI leaf switches when a transceiver of type 40G-SR-BD QSFP is used.	14.2(7l)
CSCvy27363	After a stateful reload, some of the BFD sessions do not come up. The BFDC logs show that the port down event was pushed by BFD, but the port up event was not pushed, which causes sessions to stay down. The tcpdump command shows that there is no BFD control packet being sent.	14.2(7l)
CSCvy43728	After downgrading a Cisco ACI leaf or spine switch from a 4.2 release to a 3.2 release, you may notice that the switch becomes "active" in the fabric (acidiag fnvread), but the node is missing many policies allowing it to function properly, such as the BGP Route Reflector policies.	14.2(7l)
CSCvy46638	A Cisco 9336C-FX2 switch reboots with the following reason: reset-triggered-due-to-ha-policy-of-reset System version: 14.2(7f) Service: nfm hap reset	14.2(7l)
CSCvy48603	Telemetry packets are not exported to the collector.	14.2(7l)
CSCvy80235	There is intermittent flapping on a copper-based switch.	14.2(7l)
CSCvy81682	IGMP generates a core on the deletion of a tenant.	14.2(7l)
CSCvz07799	There might be a scenario when traffic between pods is blackholed when all Spines are reloaded/upgraded in a timeframe of less than 30 minutes.	14.2(7l)
CSCvz08565	IGMP generates a core due to memory corruption or a hearbeat failure. In addition, Mcastfwd generates a core due to a heartbeat failure.	14.2(7l)
CSCvz09521	PIM generates a core while collecting tech support or executing the 'show ip pim interface vrf all' command.	14.2(7l)
CSCvu07510	There is high CPU usage due to the SNMP process.	14.2(7f)
CSCvv00971	Multi-destination traffic between sites stops working. The traffic is forwarded locally in a pod, but is not replicated properly to the remote sites. The spine switches get the traffic, but do not forward the traffic locally out to the ISN network switches.	14.2(7f)
CSCvv04140	When the "show interface <interface-name> transceiver" command is run for a Cisco ACI FEX interface, GLC-SX-MMD shows unknown. The port is up well.	14.2(7f)
CSCvv19842	With shared services inter-context traffic between remote leaf switches, there might be 2 to 3 minutes of traffic drop when upgrading the policy of the vPC pair leaf switch.	14.2(7f)
CSCvv21009	While using a Cisco N9K-C9364C-GX switch as the first or third hop leaf switch, a higher offset was observed during long duration PTP accuracy tests.	14.2(7f)
CSCvv89333	The BPDU filter doesn't work on a port channel, and STP BPDU is flooded to the encapsulation (EPG) instead of being dropped.	14.2(7f)

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

Cisco Confidential

Bug ID	Description	Fixed in
CSCvw03177	All Layer 3 multicast traffic originating within Cisco ACI and destined to a group located behind the L3Out is black holed. The traffic is sent from the server leaf switch toward the border leaf switch, but the traffic is not forwarded out of the L3Out interfaces. The source IP address that is connected to the Cisco ACI server leaf switch is not shown within the endpoint database of the border leaf switch. This only affects Layer 3 multicast traffic originating from within Cisco ACI and destined to an external group. The reverse process and pod-to-pod Layer 3 multicast works without any loss.	14.2(7f)
CSCvw15877	Regardless of the number of pods, when the TEP config changes from trail mode to path mode, the trail mode managed objects do not get deleted. In addition, some of the managed objects for ports get created for the path mode, and for some ports the managed object is not created.	14.2(7f)
CSCvw16566	Type 7 LSAs that are translated to Type 5 LSAs on a Cisco ACI border leaf switch don't have the forwarding address suppressed even though this is selected on the NSSA L3Out.	14.2(7f)
CSCvw19262	The port channel members are in a suspended state and the " show int e x/y" command shows that the interface operst is down. The LLFC/PFC operst for the port channel and member ports is up, as shown by the " show interface eth x/y flowcontrol" or " show interface eth x/y priority-flow-control" commands.	14.2(7f)
CSCvw19955	After disabling unicast routing in a bridge domain, IGMP snooping no longer works. The " show ip igmp snooping vlan <BD vlan>" command output shows " Multicast Routing enabled on VLAN" .	14.2(7f)
CSCvw20403	With scale zoning rule with stats enabled, periodically high CPU usage from the aclqos process is expected for stats collection.	14.2(7f)
CSCvw25118	The BFD protocol continues to send PDUs when the IPv6 neighbor becomes unavailable.	14.2(7f)
CSCvw33745	MLD V1 leave and MLD v1/v2 query packets cannot be tunneled when LLDP protocol tunneling is enabled.	14.2(7f)
CSCvw34334	Incorrect CPU utilization values under the " show process cpu sorted" command output.	14.2(7f)
CSCvw44520	The COOP process crashes after connectivity to the spine switch fails and recovers.	14.2(7f)
CSCvw49816	If the interface to the Tetration appliance (L3Out's external endpoint) on Leaf#1 goes down, the best route to Tetration in the spine switch's NFM does not change from Leaf#1 to Leaf#2, even though there is an update in spine switch's RIB that shows that the next-hop is changed to Leaf#2. This causes a failure to obtain flow information on the spine switches.	14.2(7f)
CSCvw51079	Some Bel Power 1100W DC PSUs have an issue in the firmware and need to be upgraded to a newer version.	14.2(7f)
CSCvw51774	IPv6 neighbor discovery doesn't work through an L3Out external bridge domain between border leaf switches when a subnet configured in the external subnets for an external EPG include a directly-connected IPv6 subnet (not ::/0).	14.2(7f)
CSCvw60119	When " show tenant <tenant name> vrf <vrf name> detail" is run from the Cisco APIC CLI, in a scale setup, the output is missing some node information.	14.2(7f)

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

Bug ID	Description	Fixed in
CSCvw62454	Not all exporters are programmed and the collector that is programmed gets only few flows, and all flows are not from the affected switch.	14.2(7f)
CSCvw66587	The inband connectivity is affected after a fabric node reboot.	14.2(7f)
CSCvw76305	In a Cisco ACI Multi-Site topology, when a bridge domain/EPG is not stretched, but a contract exits between the sites, silent hosts can't be reached.	14.2(7f)
CSCvw85874	Up to 30 seconds of routed multicast traffic loss is seen when remote learning is disabled on a border leaf switch under specific conditions.	14.2(7f)
CSCvw91341	After moving IP addresses to a new MAC address, the MAC address is considered as rogue and the rogue MAC address fault F3014 is raised. Fault F3083 can also be raised for IP addresses that are moved unexpectedly.	14.2(7f)
CSCvw92958	F3525 is reported in fabrics with high ARP and adjacency update activities. This bug is opened to add further optimization to the process to avoid the time stamp related updates and reduce the SSD writes.	14.2(7f)
CSCvw94285	A Cisco ACI leaf might intermittently become inactive due to ISIS adjacency being changed from UP to INIT, which is triggered by large amount 802.3x pause frames received from a front panel port that belongs to the same ASIC slice with all of the uplink ports.	14.2(7f)
CSCvx01777	On the Nexus 2000 Fabric extender model N2K-C2348TQ-10GE, some server facing ports may operate at 1G speed post auto-negotiation, even though the server and Fabric extender ports are configured to operate at 10G speed.	14.2(7f)
CSCvx04217	When Enabling " Enforce EPG VLAN Validation" feature in the Cisco ACI fabric under System Settings -> Fabric Wide Settings, the following validation failure error is seen: Error: 400 - Validation failed: Vlan ranges for an EPg cannot overlap Dn0=uni/tn-common/ap-Shared_ANP/epg-Rancher_EPG	14.2(7f)
CSCvx05716	A rare timing issue seen when an external router MAC address connected to an L3Out moves from being a local endpoint to another border leaf switch. EPM may core due to an invalid access of a freed data structure.	14.2(7f)
CSCvx10832	If a Cisco APIC is replaced with a new one on a different pair of leaf switches that do not have a suffix in the product ID (such as -EX or -FX), the Cisco APIC will not join the cluster. LLDP shows " not authorized" and ARP resolution fails on the Cisco APIC for 10.0.0.30. This has the same symptoms as CSCvq82478, but is not the same.	14.2(7f)
CSCvx16050	If PBR IP dataplane learning is not disabled on a service bridge domain, all IP packets coming back from the PBR node will cause unecessary notification to the software. There is no data plane impact, as learning is not happening. However, if there are a lot of endpoints crossing PBR devices, the notification may cause learning to be disabled on the ASIC for 60 seconds.	14.2(7f)
CSCvx28589	All vPCs on one node are in the LOCAL_UP_PEER_DOWN state, whereas in reality they are up on both sides.	14.2(7f)
CSCvx46437	When proxy ARP is being used, ARP is not resolved for endpoints in the same EPG when they are different pods. The ARP request from the endpoint is not flooded to the one spine switch that is used for cross pod flood traffic for that bridge domain. This happens because the traffic uses an ftag tree that has a transit leaf switch in the path to the correct spine switch. The ARP request is dropped on the transit leaf switch with ACL DROP instead of being flooded back to the spine switch.	14.2(7f)

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

Bug ID	Description	Fixed in
CSCvx47552	A spine switch reloads due to the crash of the aclqos and statsclient or eltmc process after connecting or disconnecting eth1/33 with SFP-10G-SR and after about 3 to 5. aclqos crashing occurs the most often.	14.2(7f)
CSCvx61624	A spine switch crashes due to " npv hap reset" without a core file when excecuting the " show feature" command on vsh.	14.2(7f)
CSCvx64940	Certain multicast states are missing and no joins are sent. PIM and NGMVPN are out of sync with regard to the stripe winner computation.	14.2(7f)

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 14.2(7) releases in which the bug exists. A bug might also exist in releases other than the 14.2(7) 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.	14.2(7f) 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.	14.2(7f) and later
CSCup65586	The **show interface** command shows the tunnel's Rx/Tx counters as 0.	14.2(7f) 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.	14.2(7f) and later
CSCup92534	Continuous " threshold exceeded" messages are generated from the fabric.	14.2(7f) and later
CSCuq39829	Switch rescue user (" admin") can log into fabric switches even when TACACS is selected as the default login realm.	14.2(7f) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	14.2(7f) 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.	14.2(7f) and later
CSCuq83910	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.	14.2(7f) and later
CSCuq92447	When modifying the L2Unknown Unicast parameter on a Bridge Domain (BD), interfaces on externally connected devices may bounce. Additionally, the endpoint cache for the BD is flushed and all endpoints will have to be re-learned.	14.2(7f) and later

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

Cisco Confidential

Bug ID	Description	Exists in
CSCuq93389	If an endpoint has multiple IPs, the endpoint will not be aged until all IPs go silent. If one of the IP addresses is reassigned to another server/host, the fabric detects it as an IP address move and forwarding will work as expected.	14.2(7f) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	14.2(7f) and later
CSCur81822	The access-port operational status is always "trunk".	14.2(7f) 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.	14.2(7f) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later
CSCuu11351	Traffic from the orphan port to the vPC pair is only updated on the destination node, so the traffic count shows as excess.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later

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

Bug ID	Description	Exists in
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	14.2(7f) and later
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to " 0" , and so traffic will still flow.	14.2(7f) and later
CSCuy96912	The debounce timer is not supported on 25G links.	14.2(7f) 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.	14.2(7f) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	14.2(7f) and later
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.	14.2(7f) and later
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	14.2(7f) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	14.2(7f) and later

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

Cisco Confidential

Bug ID	Description	Exists in
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.	14.2(7f) 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.	14.2(7f) and later
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later
CSCvp04772	During an upgrade on a dual-SUP system, the standby SUP may go into a failed state.	14.2(7f) and later
CSCvq56811	Output packets that are ERSPAN'd still have the PTP header. Wireshark might not be able to decode the packets, and instead shows frames with ethertype 0x8988.	14.2(7f) and later
CSCvq71034	There is a policy drop that occurs with L3Out transit cases.	14.2(7f) and later
CSCvr12912	A switch reloads due to a sysmgr heartbeat failure and sysmgr HAP reset.	14.2(7f) and later
CSCvr61096	In a port group that has ports of mixed speeds, the first port in the port group that has valid optics present and is not in the admin down state is processed. The ports that come up later are brought up if they are using the same speed; otherwise, they are put in the hw-disabled state. For example, if ports 14 and 15 are up and are using the 100G speed, then if ports 13 and 16 are using the 40G speed, these ports will be put in the hw-disabled state. After reloading or upgrading, you might not have the same interfaces in the port group in the UP state and in the hw-disabled state as you did before the reload or upgrade.	14.2(7f) and later
CSCvt53089	If a Cisco UCS fabric interconnect is deployed in the end host mode and is a peer to a Cisco ACI ToR switch, and CDP is enabled without LLDP, Blade switch MAC address move tracking is not feasible because CDP does not advertise the peer's MAC address. The blade switch MAC address entry for the fabric interconnect port MAC addresses is not seen in the output of the " show system internal epmc bladeswitch_mac all" command.	14.2(7f) and later

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

Bug ID	Description	Exists in
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.	14.2(7f) and later
CSCvv16647	A minor traffic outage is seen with a Cisco APIC downgrade.	14.2(7f) and later
CSCvw20049	A switch allows more storm traffic than the configured storm policer rate.	14.2(7f) and later
CSCvx52350	Traffic loss of may be seen after a trigger of removing and re-adding a port from a port channel while " no lacp suspend-individual" is present. The loss could be up to 15 minutes if traffic gets hashed onto the vPC leaf switch where EPM and EPMC are out of sync. The out of sync conditions maybe seen when traffic hashes onto the member port that is removed and re-added.	14.2(7f) 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.	14.2(7f) and later
N/A	IPN should preserve the CoS and DSCP values of a packet that enters IPN from the ACI spine switches. If there is a default policy on these nodes that change the CoS value based on the DSCP value or by any other mechanism, you must apply a policy to prevent the CoS value from being changed. At the minimum, the remarked CoS value should not be 4, 5, 6, or 7. If CoS is changed in the IPN, you must configure a DSCP-CoS translation policy in the APIC for the pod that translates queuing class information of the packet into the DSCP value in the outer header of the iVXLAN packet. You can also embed CoS by enabling CoS preservation. For more information, see the **Cisco APIC and QoS** KB article.	14.2(7f) 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)	14.2(7f) and later
N/A	The modular chassis Cisco ACI spine nodes, such as the Cisco Nexus 9508, support warm (stateless) standby where the state is not synched between the active and the standby supervisor modules. For an online insertion and removal (OIR) or reload of the active supervisor module, the standby supervisor module becomes active, but all modules in the switch are reset because the switchover is stateless. In the output of the **show system redundancy status** command, warm standby indicates stateless mode.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later

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

Cisco Confidential

Bug ID	Description	Exists in
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.	14.2(7f) and later
N/A	Multicast router functionality is not supported when IGMP queries are received with VxLAN encapsulation.	14.2(7f) and later
N/A	IGMP Querier election across multiple Endpoint Groups (EPGs) or Layer 2 outsides (External Bridged Network) in a given bridge domain is not supported. Only one EPG or Layer 2 outside for a given bridge domain should be extended to multiple multicast routers if any.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported in combination with static endpoint configurations.	14.2(7f) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported with Layer 2–only bridge domains. Such a configuration will not be blocked, but the configuration will not take effect as there is no Layer 3 learning in these bridge domains.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported with dynamic endpoint groups. Only static endpoint groups are supported.	14.2(7f) and later

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

Cisco Confidential

Bug ID	Description	Exists in
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.	14.2(7f) and later
N/A	Dynamic deployment of contracts based on instrImmedcy set to onDemand/lazy not supported; only immediate mode is supported.	14.2(7f) 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.	14.2(7f) 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.	14.2(7f) and later
N/A	Configurations for a virtual IP address can only be /32 or /128 prefix.	14.2(7f) 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.	14.2(7f) 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).	14.2(7f) and later
N/A	Learning through GARP will work only in ARP Flood Mode.	14.2(7f) 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(1)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 APIC (the controller cluster) to the ACI fabric, it is required to have a 10G interface on the ACI leaf. You cannot connect the APIC directly to the C9332PQ ACI leaf switch.

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

- 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.
- Cisco NX-OS uses an SSL library called CiscoSSL, which is a customized version of the OpenSSL library to support CVE fixes and FIPS compliance. Cisco maintains an extended support contract with OpenSSL. CVE fixes from OpenSSL upstream is regularly incorporated in the older versions of CiscoSSL library as well.
- The following tables provide compatibility information for specific hardware:

Table 12. Modular Spine Switch Compatibility Information

Product ID	Compatibility Information
N9K-C9336PQ	The Cisco N9K-C9336PQ switch is supported for multipod. The N9K-9336PQ switch is not supported for inter-site connectivity with Cisco ACI Multi-Site, but is supported for leaf switch-to-spine switch connectivity within a site. The N9K-9336PQ switch is not supported when multipod and Cisco ACI Multi-Site are deployed together.

Table 13. Modular Spine Switch Line Card Compatibility Information

Product ID	Compatibility Information
N9K-X9736C-FX	1-Gigabit QSA is not supported on ports 1/29-36. This line card supports the ability to add a fifth Fabric Module to the Cisco N9K-C9504 and N9K-C9508 switches. The fifth Fabric Module can only be inserted into slot 25.

Table 14. Modular Spine Switch Line Card Compatibility Information

Product ID	Compatibility Information
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-C93180LC-EX	This switch has the following limitations: • The top and bottom ports must use the same speed. If there is a speed mismatch, the top port takes precedence and bottom port will be error disabled. Both ports both must be used in either the 40

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

Product ID	Compatibility Information
	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 15. Fixed Spine Switches Compatibility Information

Product ID	Compatibility Information
N9K-C9364C	You can deploy multipod or Cisco ACI Multi-Site separately (but not together) on the Cisco N9K-9364C switch starting in the 3.1 release. You can deploy multipod and Cisco ACI Multi-Site together on the Cisco N9K-9364C switch starting in the 3.2 release. A 930W-DC PSU (NXA-PDC-930W-PE or NXA-PDC-930W-PI) is supported in redundancy mode if 3.5W QSFP+ modules or passive QSFP cables are used and the system is used in 40C ambient temperature or less; for other optics or a higher ambient temperature, a 930W-DC PSU is supported only with 2 PSUs in non-redundancy mode. 1-Gigabit QSA is not supported on ports 1/49-64. This switch supports the following PSUs: • NXA-PAC-1200W-PE • NXA-PAC-1200W-PI • N9K-PUV-1200W • NXA-PDC-930W-PE • NXA-PDC-930W-PI
N9K-C9316D-GX	1G and 100MB speeds are not supported.

Table 16. Fixed Leaf Switches Compatibility Information

Product ID	Compatibility Information
N9K-C93180YC-EX	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
N9K-C9364C-GX	This switch has the following limitations: • 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 port 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 40G or 100G. Similarly, ports 60-64 can operate only on 10G or 40G or 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

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

Cisco Confidential

Product ID	Compatibility Information
	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. • 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-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 Top-of-Rack (ToR) switch ○ When using port optical insertion and removal (OIR) ○ When performing bulk flaps of ports on the ToR switch
N9K-C93600CD-GX	This switch has the following limitations: • 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 40G, or 100G; you cannot mix the speeds. • Mixed speeds of 10G and 40G, or 10G and 100G, or 40G 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**

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

Cisco Confidential

Product ID	Compatibility Information
	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. • 1G and 100MB speeds are not supported.
N9K-C9332PQ	To connect the Cisco APIC to the Cisco ACI fabric, you must have a 10G interface on the ACI leaf switch. You cannot connect the APIC directly to the N9332PQ ACI leaf switch.

- The following table provides MACsec and CloudSec compatibility information for specific hardware:

Table 17. MACsec and CloudSec Support

Product ID	Hardware Type	MACsec Support	CloudSec Support
N9K-C93108TC-FX	Switch	Yes	No
N9K-C93180YC-FX	Switch	Yes	No
N9K-c93216TC-FX2	Switch	Yes	No
N9K-C93240YC-FX2	Switch	Yes	No
N9K-C9332C	Switch	Yes	Yes, only on the last 8 ports
N9K-C93360YC-FX2	Switch	Yes	No
N9K-C9336C-FX2	Switch	Yes	No
N9K-C9348GC-FXP	Switch	Yes, only with 10G+	No
N9K-C9364C	Switch	Yes	Yes, only on the last 16 ports
N9K-X9736C-FX	Line Card	Yes	Yes, only on the last 8 ports

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

- The following additional MACsec and CloudSec compatibility restrictions apply:
    - MACsec is not supported with 1G speed on Cisco ACI leaf switch.
    - MACsec is supported only on the leaf switch ports where an L3Out is enabled. For example, MACsec between a Cisco ACI leaf switch and any computer host is not supported. Only switch-to-switch mode is supported.
    - When using copper ports, the copper cables must be connected directly the peer device (standalone N9k) in 10G mode.
    - A 10G copper SFP module on the peer is not supported.
    - 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.

  Note: See the Cisco Application Policy Infrastructure Controller Release Notes, Release 4.2(7) 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

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

Cisco Confidential

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.
- 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 14.2(7) 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 use 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.

- 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.

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.

© 2021-2025 Cisco Systems, Inc. All rights reserved.