Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 14.2(5)

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(5)*.

For more information about this product, see Related Content.

Date	Description
July 19, 2022	In the Open Issues section, added bugs CSCwb17229 and CSCwb39899.
May 16, 2022	In the Open Issues section, added bug CSCwa47686.
March 16, 2022	In the Open Issues section, added bugs CSCwa95241 and CSCwb14844.
February 10, 2022	In the Open Issues section, added bug CSCwa18165.
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 15, 2021	In the Open Issues section, added bug CSCvw91752, CSCvx29134, CSCvy69104, CSCvz61945, and CSCvz64029.
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.
August 10, 2021	Added open issues CSCvy30381, CSCvq57414, CSCvx76219, CSCvy43728, CSCvx84820, CSCvx70611, CSCvx41386, CSCvy13313, CSCvx49448, CSCvy01336, and CSCvy80235.
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.
July 6, 2021	In the Supported Hardware section, added the NXA-PAC-500W-PI and NXA-PAC-500W-PE PSUs.
June 21, 2021	Added known issue CSCvu42069.
June 15, 2021	Added open issue CSCvy43640.
May 17, 2021	Added open issues CSCvq57414, CSCvx01777, CSCvx41386, CSCvx70611, and CSCvy01336.

Date	Description
March 23, 2021	Added open issue CSCvx70611.
March 10, 2021	Added open issue CSCvv04106.
February 22, 2021	Added open issue CSCvx18314.
February 15, 2021	Added open issue CSCvt80543.
January 22, 2021	Added open issues CSCvw98902 and CSCvt73069.
January 19, 2021	In the Known Behaviors section, changed the following sentence: The Cisco Nexus 9508 ACI-mode switch supports warm (stateless) standby where the state is not synched between the active and the standby supervisor modules. To: 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.
October 19, 2020	Release 14.2(5n) became available. Added the open and resolved issues for this release.
September 15, 2020	Release 14.2(5l) became available. Added the open and resolved issues for this release.
August 22, 2020	Release 14.2(5k) became available.

Contents

Supported Hardware

The following sections list the 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
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

Product ID	Description	Minimum	Maximum
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).
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

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

Table 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

Product ID	Description
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.
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.

Product ID	Description
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 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

Product ID	Description
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
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

Product ID	Description
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.

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

Table 11 Fixed Leaf Switch Fans

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

Product ID	Description
NXA-FAN-160CFM-PE	Blue port side exhaust fan
NXA-FAN-160CFM-PI	Burgundy port side intake fan
NXA-FAN-30CFM-B	Burgundy port side intake fan
NXA-FAN-30CFM-F	Blue port side exhaust fan
NXA-FAN-35CFM-PE	Blue port side exhaust fan
NXA-FAN-35CFM-PI	Burgundy port side intake fan
NXA-FAN-65CFM-PE	Blue port side exhaust fan
NXA-SFAN-65CFM-PE	Blue port side exhaust fan
NXA-FAN-65CFM-PI	Burgundy port side intake fan
NXA-SFAN-65CFM-PI	Burgundy port side intake fan

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(5)*.

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(5) releases in which the bug exists. A bug might also exist in releases other than the 14.2(5) releases.

Bug ID	Description	Exists In
CSCvw20119	Traffic blackholed for the entire subnet is advertised from the L3Out. The nexthop in the HAL software and the actual hardware entry in the TCAM do not match. This issue happens only if the route is present in TCAM. Routes in TRIE are not exposed to this problem.	14.2(5n) and later
CSCvv79140	PTP packets are dropped on Cisco ACI switches when the total size of the packets is larger than 128 bytes.	14.2(5l) and later
CSCvy80235	There is intermittent flapping on a copper-based switch.	14.2(5l) and later
CSCvs71658	Multiple crashes are observed on EX platforms for the following processes: SDKHAL, ELTMC, ACLQOS, IPFIB, and EPMC.	14.2(5l) and 14.2(5k)
CSCvu82773	Leaf switches reload due to an IGMP snooping process crash.	14.2(5l) and 14.2(5k)
CSCvu98988	COOP inter-pod communication stops working after reloading a switch because an endpoint entry goes missing from the COOP database. Also, one or more VNID registrations remain incomplete in BGP in the spine switch.	14.2(5l) and 14.2(5k)
CSCvv73935	A line card or fabric module reloads unexpectedly, causing an outage.	14.2(5l) and 14.2(5k)
CSCvv93409	Traffic is inadvertently contract dropped due to a problem with ACLQOS zoning-rules. Per the below zoning-rules, a rule matching the source and destination pcTag/class appears to be set and enabled. However, traffic is confirmed to be missing this rule and instead matches another rule, such as an implicit deny. In this case, the permit rule with priority 11 should be hit first before the deny rule with priority 12. The filter for the permit rule is "default," so all traffic with this source and destination pcTag/class should match.	14.2(5l) and 14.2(5k)
CSCvv99166	The mvdxn process in a fabric module leaks 4kb memory every 15 minutes.	14.2(5l) and 14.2(5k)

Bug ID	Description	Exists In
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(5k) 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(5k) 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(5k) 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(5k) and later
CSCvq19279	BFD sessions keep flapping between a -GX leaf and spine switches. The command "show system internal bfd event-history session" shows multiple instances of the Echo function failing: "Session state changed: 3(Up) -> 1(Down), New diag: 2(Echo Function Failed)"	14.2(5k) and later
CSCvq57414	HSRP/VRRP packets failed to flood locally in a service leaf switch, which causes a dual active state.	14.2(5k) and later
CSCvr16588	In a GOLF setup on a spine switch, when the bridge domain subnets and endpoints are associated to a newer VRF table and the older VRF table is deleted, after changing the VRF table (detaching the old VRF table and attaching a new VRF table) it takes long time (approximately 30 minutes) for host routes of endpoints to be advertised to CSR GOLF.	14.2(5k) and later
CSCvr71280	If the line card virtual shell (vsh_lc) crashes, the vsh_lc command must be rerun to get the CLI output. There is no other functional impact.	14.2(5k) and later
CSCvs05012	When the source is in one tenant's L3Out and the destination is in a different tenant's L3Out, the IPv6 Neighbor Discovery Protocol is not triggered for a destination IPv6 address. When both the source and destination are in the same tenant's L3Outs (even if in different VRF tables), this issue is not seen.	14.2(5k) and later
CSCvs05377	The LLDP neighbor in a modular spine switch is not formed after an upgrade.	14.2(5k) and later
CSCvs06516	A multicast route does not get programmed in the hardware. This causes the multicast traffic to be dropped.	14.2(5k) and later

Bug ID	Description	Exists In
CSCvs27994	Ping to a bridge domain's IPv6 link local address may fail.	14.2(5k) and later
CSCvs40360	Filing an enhancement to add SNMPv3 support for using AES-256 encryption.	14.2(5k) and later
CSCvs77436	Error message "No handlers could be found for logger "root"" appears when doing a moquery for certain objects.	14.2(5k) and later
CSCvs77484	A spine switch fabric module or line card is reloaded unexpectedly due to a kernel panic. The stack trace includes the following statement: Kernel panic - not syncing: Out of memory: system-wide panic_on_oom is enabled.	14.2(5k) and later
CSCvs86972	Remote leaf switches and spine switches cannot be connected to from an external virtual machine.	14.2(5k) and later
CSCvt14717	Random flaps on individual breakout ports.	14.2(5k) 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(5k) and later
CSCvt31976	There is no DOM output for DOM-supported breakout capable optics.	14.2(5k) and later
CSCvt38188	After moving static bindings from EPG A to EPG B, static bindings are not deployed on the leaf switch.	14.2(5k) and later
CSCvt50510	Modular spine switch C4 exhibits bad PTP corrections. This in turn is propagated to the underlying tier-1 and tier-2 leaf switches.	14.2(5k) and later
CSCvt63819	A border leaf switch reflects a certain number of EIGRP routes received from the external router. This issue is seen after a batch of routes are withdrawn downstream of the external router.	14.2(5k) and later

Bug ID	Description	Exists In
CSCvt71580	When the remote leaf switch POD redundancy feature is enabled, the RL tries to create a new coop connection with another POD's spine. This fabricSecurityTokenMo is periodically sent to all the nodes in the fabric every hour from the Cisco APIC and the fabricSecurityTokenMo's MD5 key is also sent, which is used for coop connection. When there is a connectivity issue between the Cisco APIC cluster, sometimes the minority Cisco APIC sends the fabricSecurityTokenMo, which causes inconsistent md5 keys in some of the nodes.	14.2(5k) and later
CSCvt73069	A Cisco ACI fabric is not fully fit after a Cisco APIC firmware upgrade.	14.2(5k) 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(5k) and later
CSCvt80543	The Cisco N9K-C9316D-GX spine switches encounter a SDKHAL process crash if the route hardware scale limits are exceeded.	14.2(5k) and later
CSCvu02527	On initiating a policy-based upgrade, if the new image has new version of the EPLD firmware, it gets upgraded by the installer. In this process irrespective of whether active SUP had an EPLD upgrade or not, if the standby SUP undergoes an EPLD firmware upgrade, then the following messages appear in the active SUP console: Module 28 EPLD upgrade is successful. EPLD upgrade procedure cannot be interrupted, signo 11 <<<<< EPLD upgrade procedure cannot be interrupted, signo 11 <<<<< CAUTION !! Shell command interrupted during EPLD upgrade Please run the command again to finish the upgrade Sending Exit command to upgrade process Upgrade process will exit in a few moments. The spine switch will not auto-reload after printing the above message.	14.2(5k) and later
CSCvu07510	There is high CPU usage due to the SNMP process.	14.2(5k) and later
CSCvu07844	When a Cisco N9K-C93180LC-EX, N9K-93180YC-EX, or N9K-C93108TC-EX leaf switch receives control, data, or BUM traffic from the front panel ports with the storm policer configured for BUM traffic, the storm policer will not get enforced. As such, the switch will let all such traffic through the system.	14.2(5k) and later
CSCvu08653	SSH from an external virtual machine does not work due to the actrlMgmtRule rule not being created.	14.2(5k) and later

Bug ID	Description	Exists In
CSCvu16473	While enabling remote leaf direct forwarding for an existing remote leaf vPC pair, connectivity loss is seen when communicating from the pod that the remote leaf is not a part of to the bridge domain pervasive gateway on the remote leaf pair. Specifically, the issue is seen after the first vPC member is decommissioned and brought back into the fabric and after the second leaf switch is decommissioned and reloaded, and before the second leaf switch is brought back into the fabric. This only affects reachability to the pervasive gateway on the remote leaf from sources in the pod of which the remote leaf is not a part.	14.2(5k) and later
CSCvu16987	A Cisco ACI leaf switch reboots due to an ICMPv6 HAP reset.	14.2(5k) and later
CSCvu22736	There is an event in which the syslog message is masked and does not provide details about the issue. The main syslog message is not seen, but rate-throttled syslog messages are seen.	14.2(5k) and later
CSCvu47561	The remote site VRF SVI IP address is missing in COOP.	14.2(5k) and later
CSCvu48811	When a Cisco ACI switch is configured in a "maintenance mode" (mmode), a banner is displayed to the user indicating the operating mode of the switch.	14.2(5k) and later
CSCvu77758	The import BGP route target that should allow leaf switches to import L3Out routes from another site is missing. As a result, the leaf switches with the missing route targets will not be able to learn routes from the remote site L3Out.	14.2(5k) and later
CSCvu84587	VTEP endpoints are learned and set to bounce on some leaf switches. A single VTEP IP address could be seen as local on one vPC pair, but as an IP XR with bounce on another leaf switch pair.	14.2(5k) and later
CSCvu97674	For all Avago 100/25G links, link up delay may be observed.	14.2(5k) and later
CSCvv01585	Pinging the inband-mgmt of a switch that is running in the ACI mode sometimes fails. This happens between leaf switches and also between leaf switches and spine switches.	14.2(5k) and later
CSCvv03974	There is an IGMP snooping memory leak after a fabric port flap trigger.	14.2(5k) 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(5k) and later

Bug ID	Description	Exists In
CSCvv05220	There is a minimal memory leak whenever a VRF table is deleted. The leak is experienced only for the VRF tables that have static routes, and is independent of the traffic.	14.2(5k) and later
CSCvv09950	Under "show platform internal counters port internal detail" of a Cisco N9K-C9504-FM-E switch, a lot of values seem to be stuck at 4294967295. This issue is cosmetic. There is no impact to traffic.	14.2(5k) and later
CSCvv10995	There is a very small memory leak the switches. The leak might not be noticed and might not lead to memory exhaustion.	14.2(5k) and later
CSCvv12422	Interleak route map with the deny action is not supported. The GUI and REST API display an error if there is a deny action in route map for interleak policies.	14.2(5k) and later
CSCvv15984	vPC switches crash due to a vPC HAP reset and a vPC HAP reset core is generated.	14.2(5k) and later
CSCvv17496	A FEX link takes a long time (5+ minutes) to come up.	14.2(5k) and later
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(5k) and later
CSCvv20902	The BGP route is supposed to update the VRF label, but COOP still has the label of the previous VRF.	14.2(5k) and later
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(5k) and later
CSCvv22833	The link between a Cisco ACI leaf switch and some 3rd party hardware appears to be up, but no traffic is received on the ingress Cisco ACI side. The link will also show as Paired (P) and be up on the Cisco ACI side.	14.2(5k) and later
CSCvv26814	30 to 60 seconds of traffic is lost for intra-POD or inter-POD remote leaf switch traffic after rebooting a switch.	14.2(5k) and later

Bug ID	Description	Exists In
CSCvv27817	DHCP unicast renewal ACKs are NOT forwarded across the fabric to clients. This traffic is sourced from port 67 destined to port 68. The regular Discover, Offer, Request, Acknowledge (DORA) process and unicast ACKs function correctly. This traffic is sourced from port 67 destined to port 67. The DHCP renewals are incorrectly being punted to the CPU as ISTACK_SUP_CODE_DHCP_SNOOP on the ingress leaf switch.	14.2(5k) and later
CSCvv33100	The IPS port is not down when an RX cable is removed on a Cisco ACI leaf switch 1G port. An ACI switch with 1G fiber would signal a peer IOS device, such as a Catalyst 6000 series switch, with flow control auto/desired to turn on the flow control.	14.2(5k) and later
CSCvv33487	Fault F1545 is raised for atomic counter drops.	14.2(5k) and later
CSCvv37006	There are PIM/PIM6 memory leaks after importing a configuration and flapping the ports.	14.2(5k) and later
CSCvv39277	After an upgrade, for one of the VRF tables, the BGP route map is missing on the spine switch, which results in bridge domain prefixes not being advertised.	14.2(5k) and later
CSCvv43137	The Cisco ACI fabric raises the critical event E4208052 "Operation errors internal error in BGP instance"."Inconsistent flags" messages are logged in "show bgp internal event-history logs".	14.2(5k) and later
CSCvv45095	After deleting then re-adding the fabric AS, the COOP process on a spine switch that is participating in multipod might generate a core file.	14.2(5k) and later
CSCvv46475	In a scaled setup, spine switches may experience delayed endpoint sync after a reload.	14.2(5k) and later
CSCvv48256	A fault for rogue endpoint detection not being generated. Rogue endpoint detection works correctly as expected.	14.2(5k) and later
CSCvv48587	Legitimate inter-pod glean traffic is dropped by the loop avoidance ACL on spine switches with or without Cisco ACI Multi-Pod QoS policies.	14.2(5k) and later
CSCvv52793	VFC flaps occur during exceptionally high DHCP traffic.	14.2(5k) and later

Bug ID	Description	Exists In
CSCvv61025	Forwarded control traffic, such as DHCP, ARP, IGMP, HSRP, and ND, does not go through storm policers.	14.2(5k) and later
CSCvv61715	A system's 100G/25G link can take up to 3 minutes to link up.	14.2(5k) and later
CSCvv75224	IPv6 BGP route with recursive next-hop is programmed in the software, but not programmed in the hardware. Traffic destined to this route is blackholed.	14.2(5k) and later
CSCvv78885	A stale route map entry is causes unexpected route leaking.	14.2(5k) and later
CSCvv79054	When OSPF is configured with "default leak policy" and "Leak default route only," the "deny-all" route map is expected to be in the OSPF process, but another route map was pushed to the OSPF process.	14.2(5k) and later
CSCvv85355	ARP responses egress the leaf switch with VLAN 0.	14.2(5k) and later
CSCvv89037	The kernel hits a HAP reset or generates a core.	14.2(5k) and later
CSCvv92044	Frequent link flaps are seen on a spine switch that is using the QSFP-100G-SM-SR transceiver.	14.2(5k) and later
CSCvv92317	Egress buffer drops are seen on Cisco N9K-C9336C-FX2 or N9K-C93240YC-FX2 switches without congestion.	14.2(5k) and later
CSCvv95800	A spine switch reloads unexpectedly due to the service on the linecard having a hap-reset.	14.2(5k) and later
CSCvv97363	On Cisco ACI fabric switches, syslog messages are generated from the incorrect VRF/Interface/source IP address.	14.2(5k) and later
CSCvw03621	Fault F1821 is raised because the /var/sysmgr directory became full on the module. The affected module might reset unexpectedly.	14.2(5k) and later

Bug ID	Description	Exists In
CSCvw06833	A subnet pushed by Openstack does not get configured onto the leaf node. In Cisco APIC, fault F1372 shows for the leaf switch: "Failed to deploy policy to service 5 on node with id X of fabric non-retrievable(3:21363)" A core is created on the Policy Element because of this fault.	14.2(5k) and later
CSCvw07282	On a modular spine switch, an unconnected port's switching state is disabled, which means it is out of service. The issue is that after reloading a line card, all of the ports on that line card change to switching state enabled, even if the port is not connected to anything. This issue is mostly cosmetic; there is no real impact if an unconnected port has switching state enabled.	14.2(5k) and later
CSCvw16121	An IGMPv3 leave causes multicast route OIL to be deleted even when there is an existing receiver subscribed to the group. Multicast traffic interrupted until the existing receivers send a report in response to a general query.	14.2(5k) and later
CSCvw19911	After disabling unicast routing from a bridge domain, the static pervasive route is still present on the leaf switch.	14.2(5k) and later
CSCvw27406	There is a discrepancy in the "show ip pim group-range vrf all" in iBash and VSH mode.	14.2(5k) and later
CSCvw28828	The e1/47 and e1/48 of a Cisco N9K-C9348GC-FXP switch will delay down the other port when reloading the switch.	14.2(5k) and later
CSCvw41780	After deleting an EPG, endpoints learned in a separate EPG are no longer advertised out as host routes.	14.2(5k) and later
CSCvw43179	Leaf switches crash with the following reset reason: Reason: reset-triggered-due-to-ha-policy-of-reset Service: mld hap reset	14.2(5k) and later
CSCvw49683	The next hop is not reachable if ARP has already been resolved. Continuous ARP requests and replies are seen for the directly connected IP address/next hop. The ARP resolution is not programmed in the hardware Forwarding Information Base (FIB) table.	14.2(5k) and later
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(5k) and later
CSCvw98902	BGP does not come up if the bridge domain enforcement flag is set under the VRF if both peers do not accept connection to port 179.	14.2(5k) and later
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(5k) and later

Bug ID	Description	Exists In
CSCvx18314	DTEP between leaf switches fails to get established due to a rare timing condition. This results in XRs not being programmed between these leaf switches. If the XR was for a multicast source, then Layer 3 multicast from that source would get affected.	14.2(5k) and later
CSCvx29134	**There are events in Cisco ACI for fans being removed from and reinserted into the switches.**	14.2(5k) and later
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(5k) and later
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(5k) and later
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(5k) and later
CSCvx70611	ARP requests that should be flooded in encapsulation are instead flooded across encapsulations on the border leaf switches.	14.2(5k) and later
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(5k) and later
CSCvx84820	The following fault is generated:[Shard 32] failed to apply tree: SLA TCP port cannot be 0	14.2(5k) and later
CSCvy01336	Traffic is sent to a device that does not match the expected device, and the symmetric hash is broken.	14.2(5k) and later
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.show system reset-reason...Reason: reset-triggered-due-to-ha-policy-of-resetService:nfm hap reset...	14.2(5k) 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(5k) and later

Bug ID	Description	Exists In
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(5k) and later
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(5k) and later
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(5k) and later
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(5k) and later
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(5k) 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(5k) and later
CSCwa18165	Ether type 0x3737 is dropped by FX switches as ACL_DROP.	14.2(5k) 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(5k) and later
CSCwa95241	An endpoint may fail to resolve an ARP request of another endpoint.	14.2(5k) and later

Bug ID	Description	Exists In
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(5k) and later
CSCwb14844	There is a long delay when connecting a Cisco N9K-C9336C-FX2 switch to a Mellanox NIC/40G using QSFP-40G-SR4. The status of the port is down, not connected.	14.2(5k) and later
CSCwb17229	The sysmgr process crashes unexpectedly, causing the line card to reload.	14.2(5k) and later
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(5k) 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(5k) 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(5k) 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(5k) and later

Bug ID	Description	Exists In
CSCvs87079	The CL91-RS-FEC mode is rejected for 25G ports on certain platforms and F1186 is raised.	14.2(5k)
CSCvv37217	A /32 host-based route is not advertised out of an L3Out when there is endpoint instability in the network for the specific host.	14.2(5k)

Resolved Issues

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

Bug ID	Description	Fixed in
CSCvs71658	Multiple crashes are observed on EX platforms for the following processes: SDKHAL, ELTMC, ACLQOS, IPFIB, and EPMC.	14.2(5n)
CSCvu82773	Leaf switches reload due to an IGMP snooping process crash.	14.2(5n)
CSCvu98988	COOP inter-pod communication stops working after reloading a switch because an endpoint entry goes missing from the COOP database. Also, one or more VNID registrations remain incomplete in BGP in the spine switch.	14.2(5n)
CSCvv73935	A line card or fabric module reloads unexpectedly, causing an outage.	14.2(5n)
CSCvv81732	When a leaf switch's published endpoints nexthop tunnel address is 0.0.0.0, a spine switch's COOP hits an assertion failure, which results in a crash.	14.2(5n)
CSCvv93409	Traffic is inadvertently contract dropped due to a problem with ACLQOS zoning-rules. Per the below zoning-rules, a rule matching the source and destination pcTag/class appears to be set and enabled. However, traffic is confirmed to be missing this rule and instead matches another rule, such as an implicit deny. In this case, the permit rule with priority 11 should be hit first before the deny rule with priority 12. The filter for the permit rule is "default," so all traffic with this source and destination pcTag/class should match.	14.2(5n)
CSCvv99166	The mvdxn process in a fabric module leaks 4kb memory every 15 minutes.	14.2(5n)
CSCvs87079	The CL91-RS-FEC mode is rejected for 25G ports on certain platforms and F1186 is raised.	14.2(5l)
CSCvv37217	A /32 host-based route is not advertised out of an L3Out when there is endpoint instability in the network for the specific host.	14.2(5l)
CSCvr91674	Whenever a switch hits a burst of PCIe, DRAM, or MCE errors, sometimes the device_test process crashes, which can cause the switch to reload.	14.2(5k) and

Bug ID	Description	Fixed in
		later
CSCvs500 80	When fast-link failover hardware support is enabled with more than 32 fabric ports, the sdkhal process aborts.	14.2(5k) and later
CSCvt1673 2	The SDK HAL process generates a core file with the decode pointing to the phy flash update.	14.2(5k) and later
CSCvv198 05	The process could run out of memory if the port flaps a few hundred times from the peer end.	14.2(5k) and later
CSCvi5889 5	CRC errors increment randomly on 10G interfaces of Cisco N9K-C93180YC-EX switches to which there are devices using Mellanox ConnectX3 adapters connected at the other end. In addition, links do not come up or take a long time to come up.	14.2(5k)
CSCvo397 15	When downgrading a Cisco ACI fabric, the OSPF neighbors go down after downgrading the Cisco APICs from a 3.2 or later release to a pre-3.2 release. After the upgrade, the switches are still running a 13.2 or later release.	14.2(5k)
CSCvq049 73	During the upgrade process, if the Bidirectional Forwarding Detection (BFD) process goes down with the reason "Control Detection Time Expired," applications dependent on BFD, such as BGP, also go down. Immediately after that, the leaf switch brings the BGP up without BFD being up. On an L3Out, this may create traffic blackholing, as BGP might remain up on the leaf switch peer until it times out while the leaf switch reloads for the upgrade.	14.2(5k)
CSCvr0040 6	A HAL crash can be trigger on border leaf switches if a multicast group IP is added as a policy prefix. This configuration is not supported.	14.2(5k)
CSCvr0242 1	N2348TQ tiburon fex randomly reboots. Crash in the 'tiburon' and/or 'ethpc' service may be observed in syslogs immediately prior to reload event.	14.2(5k)
CSCvr1469 9	The Cisco APIC raises fault F1925 for the switch's in-band IP address:IpCktEP provisioning failure for ip <IP_ADDRESS> on node with id <NODE_ID> with name <NODE_NAME> from pod <POD_ID> of fabric <FABRIC_NAME>.	14.2(5k)
CSCvr2259 9	Remote endpoints are learned behind the wrong leaf switch, which causes traffic to get forwarded and dropped at wrong the leaf switch.	14.2(5k)
CSCvr3704 0	The Cisco Nexus N9K-C9336C-FX2's 100G ports stay link-up even though the neighbor port is down. Even after unplugging the cable, the ports stay link-up.	14.2(5k)
CSCvs050 24	vPC interfaces enter the error-disabled state.	14.2(5k)
CSCvs176 88	Interfaces on a leaf switch show no input rate.	14.2(5k)

Bug ID	Description	Fixed in
CSCvs18150	After a certain set of steps, it is observed that the deny-external-tag route-map used for transit routing loop prevention gets set back to the default tag 4294967295. Since routes arriving in Cisco ACI with this tag are denied from being installed in the routing table, if the VRF table that has the route-tag policy is providing transit for another VRF table in Cisco ACI (for instance and inside and outside vrf with a fw connecting them) and the non-transit VRF table has the default route-tag policy, routes from the non-transit VRF table would not be installed in the transit VRF table.

This bug is also particularly impactful in scenarios where transit routing is being used and OSPF or EIGRP is used on a vPC border leaf switch pair. vPC border leaf switches peer with each other, so if member A gets a transit route from BGP, redistributes into OSPF, and then advertises to member B (since they are peers)...without a loop prevention mechanism, member B would install the route through OSPF since it has a better admin distance and would then advertise back into BGP. This VRF tag is set on redistribution of BGP > OSPF and then as a table map in OSPF that blocks routes with the tag from getting installed in the routing table. When hitting this bug, the route-map used for redistributing into OSPF still sets the tag to the correct value. However, the table map no longer matches the correct tag. Rather, it matches the default tag. As a result, member A (could be B) would install the route through OSPF pointing to B. It would then redistribute it back into BGP with the med set to 1. The rest of the fabric (including member B) would install the BGP route pointing to member A since its med is better than the original route's med. | 14.2(5k) |
| CSCvs31340 | In ACI 4.1 releases, FEX port-channel member interfaces (NIF) can no longer be configured as SPAN source interfaces.

The following Fault is raised and the SPAN session remains operationally down:

F1199: Span source interface sys/phys-[eth1/x] on node xxx in failed state reason Configuration not supported on this TOR. | 14.2(5k) |
| CSCvs33445 | The following conditions must be true:

1) An L3Out External EPG has both a 0/0 prefix and multiple specific prefixes (e.g. One 0/0 prefix and two /24 prefixes).

2) All specific prefixes are deleted in the same POST or all specific prefixes are created in the same POST.

This issue is most commonly seen during a configuration rollback, where all specific prefixes are either deleted or recreated. | 14.2(5k) |
CSCvs50966	The sclass of an endpoint (MAC/IP address) is programmed with 0, thereby causing traffic to drop.	14.2(5k)
CSCvs56976	The event manager service crashes on a fabric node when using testapi to trigger the eventmgr (running a switch node) for sending a log message to a remote syslog server.	14.2(5k)
CSCvs60566	A vPC pair of leaf switches go into the split brain mode, causing traffic duplication.	14.2(5k)

Bug ID	Description	Fixed in
CSCvs97879	According to the documentation, power supply slot 1 is on the left and slot 2 is on the right. However, N9K-C9332C recognizes the left side slot as slot 2 and vice versa.	14.2(5k)
CSCvs97901	A node resets with the following crash reason: Last reset at 0 usecs after Fri Jan 31 05:52:52 2020 CST Reason: reset-triggered-due-to-ha-policy-of-reset System version: 13.2(4d) Service: vntag_mgr hap reset	14.2(5k)
CSCvt09775	The policy element crashes due to database space exhaustion with B22 FEX devices in the fabric.	14.2(5k)
CSCvt25313	Inter-pod traffic drops in spine switches that have -EX, -FX, or later suffixes in the product ID. Packets destined to hosts in remote a pod reach the spine switch and are not forwarded despite having the routes present with the correct destination tunnels.	14.2(5k)
CSCvt25383	The pervasive static route is missing on the spine node.	14.2(5k)
CSCvt25585	16 GB leaf switches that do not have -EX, -FX, or a later suffix in the product ID have a kernel panic without generating any core files. The dmesg output shows it was due to an out-of-memory issue.	14.2(5k)
CSCvt40739	The tunnel destination IP address is not leaned in EPM as a VTEP endpoint	14.2(5k)
CSCvt44961	After an upgrade, some ports do not come up.	14.2(5k)
CSCvt49916	A spine switch reloads with the following reason: reset-triggered-due-to-ha-policy-of-reset	14.2(5k)
CSCvt49963	ARP requests between EPGs in the same bridge domain are being redirected by the service graph without using an Allow Any contract.	14.2(5k)
CSCvt57119	A Cisco ACI leaf switch sends traffic that is untagged for a particular VLAN even though it is configured as trunk (tagged).	14.2(5k)
CSCvt60814	A port-facing interpod network gets the "Uncalibrated" status, and PTP packets are not seen as received on the spine switch.	14.2(5k)
CSCvt64042	The policy element crashes once during a misconfiguration.	14.2(5k)
CSCvt64733	An ARP request from the endpoint behind the remote leaf switch is received on the leaf switch and is flooded to the spine switch as expected (ARP flooding enabled on bridge domain). This can be an issue in cases where the endpoint is behind something such as a Fabric	14.2(5k)

Bug ID	Description	Fixed in
	Interconnects, in which it may be expected behavior to delete the MAC address if the endpoint receives the same MAC address back from the upstream leaf switches.	
CSCvt6724 9	After upgrading to release 4.2(3q), Fibre Channel interfaces go down with faults F3248 and F3307.	14.2(5k)
CSCvt6745 8	Graceful restart does not show up in "show ip ospf" when enabled.	14.2(5k)
CSCvt6935 6	The sysmgr process generates core files from time to time on spine switches.	14.2(5k)
CSCvt7213 7	Traffic on the interface of leaf switch model N9K-C9348GC-FXP is being received at the rate of 125k pps and is not being dropped.	14.2(5k)
CSCvt8266 9	In a scaled setup, some public bridge domain subnets are not advertised out of the L3Outs to the neighbors.	14.2(5k)
CSCvt8549 9	When using host-based routing (/32 host route for an endpoint advertised to an L3Out) and the border leaf node loses all of its adjacency, the border leaf node keeps the /32 host route and still advertises the route to the routing protocol for about 15 minutes. This can cause a blackhole of traffic from external routes to the /32 host routes.	14.2(5k)
CSCvt9403 9	A leaf switch crashes and reloads due to "nfm hap reset".	14.2(5k)
CSCvt9441 2	In a situation in which a virtual machine cannot reach the default gateway, Cisco ACI declares a VMNIC as down, but VMware vCenter displays that VMNIC as connected.	14.2(5k)
CSCvt9836 4	When dot1 preserve is set and dscp/cos map translation is disable, single chip spine (non-modular) will mistreat traffic received with outer Cos 6 from IPN/ISN. It will be classifies as traceroute and dropped.	14.2(5k)
CSCvu016 39	There are faults for failed contract rules and prefixes on switches prior to the -EX switches. Furthermore, traffic that is destined to an L3Out gets dropped because the compute leaf switches do not have the external prefix programmed in ns shim GST-TCAM. You might also see that leaf switches prior to the -EX switches do not have all contracts programmed correctly in the hardware.	14.2(5k)
CSCvu080 65	If inter-VRF DHCP relay is used, it may be observed that DHCP breaks after performing any activity that causes the client VRF to get removed and re-deployed on the client leaf nodes.	14.2(5k)
CSCvu145 45	An admin down or blacklisted port shows as being delivered power with the command "show power inline".	14.2(5k)
CSCvu157 12	If a spine switch's PTEP is configured as the multipod L3Out router ID and the router ID is later changed, the spine switch's PTEP loopback gets deleted and the MP BGP session goes down.	14.2(5k)
CSCvu157 51	The following event can be seen on the spine node: [E4204936][transition][warning][sys] %URIB-4-SYSLOG_SL_MSG_WARNING: URIB-5-	14.2(5k)

Bug ID	Description	Fixed in
	RPATH_DELETE: message repeated 1 times in last 220162 sec	
CSCvu21147	QSFP-40G-SR-BD SFP shows DOM from "show interface ethernet x/x transceiver details" even though it is unsupported hardware.	14.2(5k)
CSCvu23722	Removing GBIC from leaf node uplink ports/spine node ports generates the "Transceiver has been removed" syslog event on unused ports on the spine node and some leaf nodes. The reason is that when GBIC is removed, it triggers these events internally due to using the wrong ethpmFcot managed object logic. The only affect the internal managed object class update behavior. You can safely ignore the events.	14.2(5k)
CSCvu26947	If a rogue file grows too large, it can cause out of memory condition on a spine switch line card or fabric module without proactively alerting the user to the memory leak, and the line card or fabric module will reload.	14.2(5k)
CSCvu27617	Fault F0098 gets raised and the statistics in the Capacity dashboards do not show for the affected leaf nodes.	14.2(5k)
CSCvu27791	Paths to L1/L2 devices do not get programmed although they are tracked as up. This happens in an active-standby deployment.	14.2(5k)
CSCvu28036	There is traffic loss as a result of a missing endpoint in COOP on the spine node.	14.2(5k)
CSCvu30087	After reloading, clean reloading, or upgrading, a remote leaf node, the remote leaf node L3Outs no longer advertise host routes.	14.2(5k)
CSCvu32603	The AVAGO QSFP-100G-SR4 transceiver displays an abnormal Rx dbm value in a down port.	14.2(5k)
CSCvu40050	The spine node KIC database is missing the v4 default route from RIB. This causes in-band return traffic to drop on the way back to the border leaf nodes.	14.2(5k)
CSCvu43096	This is an enhancement to fix the show command outputs on ACI leaf switches related to Auto-Recovery and Delay-Restore feature for vPC. In ACI, Auto-Recovery is not available but instead Peer Dead Interval is implemented via MO vpcInstPol and its attribute deadIntvl. Delay Restore is implemented in ACI but not configurable. It is applied with a default value that may depend on each ACI release. However, `show vpc` command output on leaf switches shows Auto-Recovery as enabled (also with a wrong default interval value instead of the one for vpcInstPol.deadIntvl). The same command does not provide any information regarding Delay Restore.	14.2(5k)
CSCvu51926	ARP packets are dropped in the fabric that should be implicitly permitted using zoning-rule programming. The packet drop reason is SECURITY_GROUP_DENY and the specific rule is implicit deny for the VRF table.	14.2(5k)

Bug ID	Description	Fixed in
CSCvu535 58	When walking through SNMP targets, SNMP generates a core file on a spine switch.	14.2(5k)
CSCvu610 24	Zoning-rules are not programmed in the hardware after reloading a switch.	14.2(5k)
CSCvu642 19	There are 2 issues caused by this bug: 1. Dataplane learning is disabled, meaning learning is only occurs only from ARP and neighbor discovery. This issue occurs on all platforms. 2. Hitbit does not get updated during learn disable, causing endpoints to age prematurely. This issue occurs with N9K-C9336C-FX2, N9K-C93240YC-FX2, N9K-C93240YC-FX2Z, and later switches.	14.2(5k)
CSCvu724 16	Triggered by a physical layer issue, such as fiber or a bad transceiver, a link flap may happen every now and then. However, it is uncommon to have continuous flaps when the node is left unattended over an extended period, such as having 688,000 flaps over a year. Each time after the fabric link flaps, one dbgRemotePort managed object is added to the policyElement database. After a long time flapping like this, unexpected memory allocation and access can be triggered for the Nexus OS process, such as policy_mgr or ethpm. This defect is to enhance the object-store to reduce the impact for such scenarios.	14.2(5k)
CSCvu824 81	There is a memory leak on the leaf switches due to the vlan_mgr process.	14.2(5k)
CSCvu905 28	A switch reboots unexpectedly due to "ip hap reset".	14.2(5k)
CSCvv008 66	The FCOE host may experience latency and disconnections after a fabric monitoring policy change for the switch nodes. The monitor policy itself does not include an FCOE setting. However, the policy change may result in impact to FCOE traffic on the switch level.	14.2(5k)
CSCvv088 50	A fabric node device crashes with an SDKHAL core.	14.2(5k)
CSCvv098 10	For a local receiver that is added by IGMP, there is no extranet linkage for the S,G in the source VRF table. The source VRF table's S,G outgoing interface (OIF) will not contain the extranet OIF list.	14.2(5k)
CSCvv112 92	ACLLOG-5-ACLLOG_PKTLOG syslog event messages are not generated after an upgrade.	14.2(5k)
CSCvv137 22	A leaf switch crashes due to a routing loop in the IPFIB process.	14.2(5k)
CSCvv205 45	DHCP generates a core file due to a memory leak issue. If you have a remote leaf switch and DHCP configured, you can run into this issue.	14.2(5k)

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(5) releases in which the known behavior exists. A bug might also exist in releases other than the 14.2(5) 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(5k) 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(5k) and later
CSCup65586	The show interface command shows the tunnel's Rx/Tx counters as 0.	14.2(5k) 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(5k) and later
CSCup92534	Continuous "threshold exceeded" messages are generated from the fabric.	14.2(5k) and later
CSCuq39829	Switch rescue user ("admin") can log into fabric switches even when TACACS is selected as the default login realm.	14.2(5k) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	14.2(5k) 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(5k) 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(5k) 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(5k) and later

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(5k) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	14.2(5k) and later
CSCur81822	The access-port operational status is always "trunk".	14.2(5k) 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(5k) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	14.2(5k) 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(5k) 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(5k) 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(5k) 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(5k) and later

Bug ID	Description	Exists In
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(5k) 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(5k) 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(5k) 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(5k) 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(5k) and later
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.	14.2(5k) 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(5k) 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(5k) 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(5k) 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(5k) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	14.2(5k) and later
CSCuy22288	The iping **command's** replies get dropped by the QOS ingress policer.	14.2(5k) and later

Bug ID	Description	Exists In
CSCuy25780	An overlapping or duplicate prefix/subnet could cause the valid prefixes not to be installed because of batching behavior on a switch. This can happen during an upgrade to the 1.2(2) release.	14.2(5k) 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(5k) 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(5k) 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(5k) and later
CSCuy96912	The debounce timer is not supported on 25G links.	14.2(5k) 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(5k) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	14.2(5k) 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(5k) and later
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	14.2(5k) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	14.2(5k) and later
CSCva98767	The front port of the QSA and GLC-T 1G module has a 10 to 15-second delay as it comes up from the insertion process.	14.2(5k) 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(5k) and later

Bug ID	Description	Exists In
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	14.2(5k) 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(5k) 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(5k) 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(5k) 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(5k) and later
CSCvp04772	During an upgrade on a dual-SUP system, the standby SUP may go into a failed state.	14.2(5k) 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(5k) and later
CSCvq71034	There is a policy drop that occurs with L3Out transit cases.	14.2(5k) and later
CSCvr12912	A switch reloads due to a sysmgr heartbeat failure and sysmgr HAP reset.	14.2(5k) 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(5k) and later

Bug ID	Description	Exists In
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(5k) and later
CSCvu42069	The event log shows VTEP tunnel down and up events. The down time and up time are the same, and there is no fault message.	14.2(5k) 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(5k) 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(5k) 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(5k) 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(5k) 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(5k) 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(5k) and later

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(5k) and later
N/A	Multicast router functionality is not supported when IGMP queries are received with VxLAN encapsulation.	14.2(5k) 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(5k) 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(5k) 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(5k) 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(5k) 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(5k) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported in combination with static endpoint configurations.	14.2(5k) 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(5k) 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(5k) and later

Bug ID	Description	Exists In
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(5k) 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(5k) and later
N/A	No fault will be raised if the IP/MAC Ckt endpoint prefix configured is outside of the bridge domain subnet range. This is because a user can configure bridge domain subnet and IP/MAC Ckt endpoint in any order and so this is not error condition. If the final configuration is such that a configured IP/MAC Ckt endpoint prefix is outside all bridge domain subnets, the configuration has no impact and is not an error condition.	14.2(5k) and later
N/A	Dynamic deployment of contracts based on instrImmedcy set to onDemand/lazy not supported; only immediate mode is supported.	14.2(5k) 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(5k) 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(5k) and later
N/A	Configurations for a virtual IP address can only be /32 or /128 prefix.	14.2(5k) 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(5k) 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(5k) and later
N/A	Learning through GARP will work only in ARP Flood Mode.	14.2(5k) and later

Compatibility Information

- For the supported optics per device, see the [Cisco Optics-to-Device Compatibility Matrix](#).

- Link level flow control is not supported on ACI-mode switches.

- 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.

- 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.

- 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

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. ■ When a Cisco N9K-C9348GC-FXP switch has only one PSU inserted and connected, the PSU status for the empty PSU slot will be displayed as "shut" instead of "absent" due to a hardware limitation. ■ 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 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	Description

Product ID	Description
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	Description
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

Product ID	Description
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 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.

Product ID	Description

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
	
	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

Product ID	Description
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 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

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

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(5)* for policy information.

— UDP DestPort 161: SNMP. These cannot be blocked through contracts. Creating an SNMP ClientGroup with a list of Client-IP Addresses restricts SNMP access to only those configured Client-IP Addresses. If no Client-IP address is configured, SNMP packets are allowed from anywhere.

— TCP SrcPort 179: BGP

— TCP DstPort 179: BGP

— OSPF

— UDP DstPort 67: BOOTP/DHCP

— UDP DstPort 68: BOOTP/DHCP

— IGMP

— PIM

— UDP SrcPort 53: DNS replies

— TCP SrcPort 25: SMTP replies

— TCP DstPort 443: HTTPS

— UDP SrcPort 123: NTP

— UDP DstPort 123: NTP

■ Leaf switches and spine switches typically have memory utilization of approximately 70% to 75%, even in a new deployment where no configuration has been pushed. This amount of memory utilization is due to the Cisco ACI-specific processes, which take up more memory compared to a standalone Nexus deployment. The memory utilization is not a problem unless it exceeds 90%. You can open a Cisco TAC case to troubleshoot proactively when memory utilization is more than 85%.

■ The Cisco APIC GUI incorrectly reports more memory used than is actually used. To calculate the appropriate amount of memory used, run the "show system internal kernel meminfo | egrep "MemT|MemA"" command on the desired switch. Divide MemAvailable by MemTotal, multiply that number by 100, then subtract that number from 100.

— Example: 10680000 / 24499856 = 0.436 x 100 = 43.6% Free, 100% - 43.6% = 56.4% Used

■ 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(5) 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.

- 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 Cisco APIC 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: 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.