Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 14.2(2)

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, bugs, and limitations for the Cisco NX-OS software. Use this document in combination with the *Cisco Application Policy Infrastructure Controller Release Notes, Release 4.2(2)*.

Additional product documentation is listed in the "Related Documentation" section.

Release notes are sometimes updated with new information about restrictions and issues. See the following website for the most recent version of the *Cisco Nexus 9000 ACI-Mode Switches Release Notes*:

https://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-release-notes-list.html

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

Table 1 shows the online change history for this document.

Table 1. Online History Change

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.
August 18, 2021	In the Compatibility Information section, for the N9K-C9316D-GX and N9K-C93600CD-GX switches, added: 1G and 100MB speeds are not supported.
August 10, 2021	In the Open Issues section, added bugs CSCvy30381, CSCvq57414, and CSCvx76219.
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 24, 2021	In the Open Issues section, added bug CSCvu07844.
June 15, 2021	In the Open Issues section, added bug CSCvy43640.

Date	Description
May 17, 2021	In the Open Issues section, added bug CSCvq57414.
February 22, 2021	In the Open Issues section, added bug CSCvx18314.
January 22, 2021	In the Open Issues section, added bug 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.
March 13, 2020	14.2(2e): In the Resolved Issues section, added bug CSCvr98827. Added known behaviors CSCvq56811.
December 5, 2019	14.2(2e): In the Open Issues section, added bug CSCvr76947.
November 14, 2019	14.2(2g): Release 14.2(2g) became available. Added the resolved issues for this release.
November 13, 2019	In the Compatibility Information section, for the N9K-C93600CD-GX switch, added the following bullets: ■ 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.
October 30, 2019	14.2(2f): Release 14.2(2f) became available. Added the open and resolved issues for this release.
October 17, 2019	14.2(2e): In the Resolved Issues section, added bugs CSCvq42672, CSCvq97092, CSCvq99155, and CSCvr52067.
October 14, 2019	14.2(2e): In the Open Issues section, added bug CSCvr51479.
October 10, 2019	14.2(2e): Release 14.2(2e) became available.

Cisco Confidential

Contents

This document includes the following sections:

- Contents
- Supported Hardware
- Supported FEX Models
- New and Changed Information
- Installation Notes
- Compatibility Information
- Usage Guidelines
- Issues
- Related Documentation

Cisco Confidential

Supported Hardware

The following sections list the supported hardware.

Table 2 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 3 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 4 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

Cisco Confidential

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 5 Modular Spine Switch Supervisor and System Controller Modules

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

Table 6 Fixed Spine Switches

Product ID	Description
N9K-C9316D-GX	Cisco Nexus N9K-C9316D-GX is a 1 RU spine switch with 16 10/40/100/400-Gigabit QSFP-DD ports (ports 1-16). Note: This switch does not support 100G-CR4 cables.
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 9364C switch is a 2-rack unit (RU), fixed-port switch designed for spine-leaf-APIC deployment in data centers. This switch supports 64 40/100-Gigabit QSFP28 ports and two 1/10-Gigabit SFP+ ports. The last 16 of the QSFP28 ports are colored green to indicate that they support wire-rate MACsec encryption.

Table 7 Fixed Spine Switch 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

Cisco Confidential

NXA-FAN-160CFM-PI	Burgundy port side intake fan
NXA-FAN-35CFM-PE	Blue port side exhaust fan
NXA-FAN-35CFM-PI	Burgundy port side intake fan

Table 8 Fixed Leaf Switches

Product ID	Description
N9K-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 9336C-FX2 Top-of-rack (ToR) switch with 36 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: 1-Gigabit QSA is not supported on ports 1/1-6 and 1/33-36. The port profile feature supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.
N9K-C93108TC-FX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.
N9K-C93108TC-FX-24	Cisco Nexus 9300 platform switch with 24 1/10GBASE-T (copper) front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.
N9K-C93180YC-FX	Cisco Nexus 9300 platform switch with 48 1/10/25-Gigabit Ethernet SFP28 front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. The SFP28 ports support 1-, 10-, and 25-Gigabit Ethernet connections and 8-, 16-, and 32-Gigabit Fibre Channel connections. Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.

Cisco Confidential

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
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

Cisco Confidential

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
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
NXA-PAC-1100W-PE2	1100W AC power supply, port side exhaust pluggable
NXA-PAC-1100W-PI2	1100W AC power supply, port side intake pluggable
N9K-PAC-650W	650W AC Power supply, port side intake pluggable
N9K-PAC-650W-B	650W AC Power supply, port side exhaust pluggable
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-PHV-1100W-PE	1100W HVAC/HVDC power supply, port-side exhaust
NXA-PHV-1100W-PI	1100W HVAC/HVDC power supply, port-side intake

Cisco Confidential

NXA-PAC-1200W-PE	1200W AC Power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance Note: This power supply is supported only by the Cisco Nexus 93120TX, 93128TX, and 9336PQ ACI-mode switches.
NXA-PAC-1200W-PI	1200W AC Power supply, port side intake pluggable, with higher fan speeds for NEBS compliance Note: This power supply is supported only by the Cisco Nexus 93120TX, 93128TX, and 9336PQ ACI-mode switches.
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-500W-PE	500W AC Power supply, port side exhaust pluggable
NXA-PAC-500W-PI	500W AC Power supply, port side intake pluggable
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.
N9K-PUV-1200W	1200W HVAC/HVDC dual-direction airflow power supply Note: This power supply is supported only by the Cisco Nexus 93120TX, 93128TX, and 9336PQ ACI-mode switches
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

Cisco Confidential

NXA-FAN-160CFM2-PE	Blue port side exhaust fan
NXA-FAN-160CFM2-PI	Burgundy port side intake fan
NXA-FAN-160CFM-PE	Blue port side exhaust fan
NXA-FAN-160CFM-PI	Burgundy port side intake fan
NXA-FAN-30CFM-B	Burgundy port side intake fan
NXA-FAN-30CFM-F	Blue port side exhaust fan
NXA-FAN-35CFM-PE	Blue port side exhaust fan
NXA-FAN-35CFM-PI	Burgundy port side intake fan
NXA-FAN-65CFM-PE	Blue port side exhaust fan
NXA-SFAN-65CFM-PE	Blue port side exhaust fan
NXA-FAN-65CFM-PI	Burgundy port side intake fan
NXA-SFAN-65CFM-PI	Burgundy port side intake fan

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 and Changed Information

This section lists the new and changed features in this release.

- New Hardware Features

Cisco Confidential

- New Software Features

New Hardware Features

The following hardware features are now available:

- The Cisco Nexus 93600CD-GX switch (N9K-C93600CD-GX) is a 1-rack unit (RU), fixed-port switch designed for spine-leaf-APIC deployment in data centers. It has 28 10/40/100-Gigabit Ethernet QSFP28 ports (ports 1-28), 8 10/40/100/400-Gigabit QSFP-DD ports (ports 29-36), two management ports (one 10/100/1000BASE-T port and one 1-Gbps SFP+ port), one console port (RS-232), and 1 USB port.
- The Cisco Nexus switch N9K-C9316D-GX is a 1-rack unit (RU), fixed-port switch designed for spine-leaf-APIC deployment in data centers. It has 16 10/40/100/400-Gigabit QSFP-DD ports (ports 1-16), two management ports (one 10/100/1000BASE-T port and one 1-Gbps SFP+ port), one console port (RS-232), and 1 USB port.

New Software Features

For new software features, see the *Cisco Application Policy Infrastructure Controller Release Notes, Release 4.2(2)*.

Changes in Behavior

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

Installation Notes

The following procedure installs a Gigabit Ethernet module (GEM) in a top-of-rack switch:

1. Clear the switch's current configuration by using the setup-clean-config command.
2. Power off the switch by disconnecting the power.
3. Replace the current GEM card with the new GEM card.
4. Power on the switch.

For other installation instructions, see the *Cisco Application Centric Infrastructure Fabric Hardware Installation Guide*.

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:

Cisco Confidential

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

Table 12 Modular Spine Switch Compatibility Information

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

Table 13 Modular Spine Switch Line Card Compatibility Information

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

Table 14 Modular Spine Switch Line Card Compatibility Information

Product ID	Compatibility Information

Cisco Confidential

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

Cisco Confidential

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

Cisco Confidential

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. ■ 4X100 breakout is not supported. ■ The Hardware Abstraction Layer (HAL) will spike and the console can hang if 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. ■ This switch does not support 100G-CR4 cables. ■ 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

Cisco Confidential

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.

Cisco Confidential

Note: See the *Cisco Application Policy Infrastructure Controller Release Notes, Release 4.2(2)* for policy information.

- UDP DestPort 161: SNMP. These cannot be blocked through contracts. Creating an SNMP ClientGroup with a list of Client-IP Addresses restricts SNMP access to only those configured Client-IP Addresses. If no Client-IP address is configured, SNMP packets are allowed from anywhere.
- TCP SrcPort 179: BGP
- TCP DstPort 179: BGP
- OSPF
- UDP DstPort 67: BOOTP/DHCP
- UDP DstPort 68: BOOTP/DHCP
- IGMP
- PIM
- UDP SrcPort 53: DNS replies
- TCP SrcPort 25: SMTP replies
- TCP DstPort 443: HTTPS
- UDP SrcPort 123: NTP
- UDP DstPort 123: NTP

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

Cisco Confidential

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(2) or later. If you must downgrade the firmware from a release that supports FIPS to a release that does not support FIPS, you must first disable FIPS on the Cisco ACI fabric and reload all of the switches in the fabric.

- You cannot use the breakout feature on a port that has a port profile configured on a Cisco N9K-C93180LC-EX switch. With a port profile on an access port, the port is converted to an uplink, and breakout is not supported on an uplink. With a port profile on a fabric port, the port is converted to a downlink. Breakout is currently supported only on ports 1 through 24.
- On Cisco 93180LC-EX Switches, ports 25 and 27 are the native uplink ports. Using a port profile, if you convert ports 25 and 27 to downlink ports, ports 29, 30, 31, and 32 are still available as four native uplink ports. Because of the threshold on the number of ports (which is maximum of 12 ports) that can be converted, you can convert 8 more downlink ports to uplink ports. For example, ports 1, 3, 5, 7, 9, 13, 15, 17 are converted to uplink ports and ports 29, 30, 31 and 32 are the 4 native uplink ports, which is the maximum uplink port limit on Cisco 93180LC-EX switches.

  When the switch is in this state and if the port profile configuration is deleted on ports 25 and 27, ports 25 and 27 are converted back to uplink ports, but there are already 12 uplink ports on the switch in the example. To accommodate ports 25 and 27 as uplink ports, 2 random ports from the port range 1, 3, 5, 7, 9, 13, 15, 17 are denied the uplink conversion; the chosen ports cannot be controlled by the user. Therefore, it is mandatory to clear all the faults before reloading the leaf node to avoid any unexpected behavior regarding the port type. If a node is reloaded without clearing the port profile faults, especially when there is a fault related to limit-exceed, the ports might be in an unexpected mode.
- When using a 25G Mellanox cable that is connected to a Mellanox NIC, you can set the ACI leaf switch port to run at a speed of 25G or 10G.
- You cannot 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.

Issues

This section contains lists of open and resolved issues and known behaviors.

- Open Issues
- Resolved Issues
- Known Behaviors

Open Issues

This section lists the open bugs. 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(2) releases in which the bug exists. A bug might also exist in releases other than the 14.2(2) releases.

Table 18 Open Issues in This Release

Bug ID	Description	Exists In
CSCvq57414	HSRP/VRRP packets failed to flood locally in a service leaf switch, which causes a dual active state.	14.2(2g) 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(2g) and later
CSCvs21844	A data packet in a Layer 2-only VLAN dedicated for an FTD cluster communication link (CCL) has an invalid L4 checksum (expected). An N9K-C9364C switch seems to be stomping these packets when getting them from the IPN on interfaces 65-66. This results in an incorrect l2 frame checksum, and a data packet from FTD-A sent to FTD-B through Cisco ACI cannot be processed.	14.2(2f) and later
CSCvs40360	Filing an enhancement to add SNMPv3 support for using AES-256 encryption.	14.2(2f) and later
CSCvs79018	Running log_trace_bl_print_tool will result in the decoded binary logs referencing the system localtime. This will likely differ from the timezone of the collected techsupports.	14.2(2f) and later
CSCvt00457	Egress FCOE traffic is incorrectly classified in the default class (Level 3) instead of the no drop class (Level 1). This issue causes priority flow control to not transmit pause frames to throttle the traffic during heavy congestion. Level1 Rx Drop Pkts/Bytes will then result.	14.2(2f) and later
CSCvt07621	The fabric module on a modular spine switch crashes with sdkhal-hap-reset and generates an sdkhal core file.	14.2(2f) and later
CSCvt14717	Random flaps on individual breakout ports.	14.2(2f) and later
CSCvt44961	After an upgrade, some ports do not come up.	14.2(2f) and later
CSCvt67458	Graceful restart does not show up in "show ip ospf" when enabled.	14.2(2f) and later
CSCvu21147	QSFP-40G-SR-BD SFP shows DOM from "show interface ethernet x/x transceiver details" even though it is unsupported hardware.	14.2(2f) and later
CSCvv08850	A fabric node device crashes with an SDKHAL core.	14.2(2f) and later
CSCvv19805	The process could run out of memory if the port is flapped million times from the peer end.	14.2(2f) 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(2f) and later
CSCvv92317	Egress buffer drops are seen on Cisco N9K-C9336C-FX2 or N9K-C93240YC-FX2 switches without congestion.	14.2(2f) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvr63689	The ACI N93360YC-FX2 leaf switch becomes inactive.	14.2(2e) through 14.2(2f)
CSCvs00854	A switch's DC PSU went to the "shut" or "failed" state, and the PSU LED was amber color.	14.2(2e) through 14.2(2f)
CSCvq85886	When an ARP request is generated from one endpoint to another endpoint in an isolated EPG, an ARP glean request is generated for the first endpoint.	14.2(2e) 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(2e) 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(2e) and later
CSCvn16192	CRC errors increment on a leaf switch front panel port, fabric ports, and spine switch ports in a fabric with switches whose model names end with -EX, -FX, or later.	14.2(2e) and later
CSCvo39715	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(2e) and later
CSCvo86795	SAN port channel bringup will be unsuccessful when a new vendor switch is connected and the Organizationally Unique Identifier (OUI) of the switch is not present in the OUI list.	14.2(2e) and later
CSCvp00694	If a tenant is undeployed from the Multi-Site Orchestrator, the limited VRF instances on the spine switch created for GOLF can be stuck in deletion based on timing. This issue occurs in a topology/ConfigA single pod setup, with a VRF instance and bridge domain that is stretched across sites and GOLF host route is enabled.	14.2(2e) 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(2e) and later
CSCvp50075	A leaf switch experiences an unexpected reload due to a HAP reset.	14.2(2e) and later
CSCvp72312	A contract that is provided by an EPG using a bridge domain with subnet X and that is consumed by an L3Out EPG causes a leak of subnet X from VRF B to VRF A. The existing non-pervasive static route in VRF A is replaced by a pervasive route in pointing to spine switch V4 proxy. After the contract leaking subnet A is removed, the pervasive static route persists.	14.2(2e) and later
CSCvq19269	When a vPC leg goes down and comes back up, a long traffic drop (1-2 minutes) may occur.	14.2(2e) 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(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvq20215	Some ECMP paths may be flap between "multipath" and "non-multipath." For example, if the configured EBGP MAX ECMP number is 10 and there are 16 BGP ECMP paths for a prefix in the BGP routing table, then 5 paths change between multipath and non-multipath whenever the BGP bestpath calculation is run.	14.2(2e) and later
CSCvq33855	The VMAC of an ACI SVI is used for an endpoint refresh instead of the PMAC. For example, it can cause an endpoint refresh issue, if the endpoint is reachable through OTV (when HSRP filtration is used and the HSRP MAC is set as a virtual MAC address).	14.2(2e) and later
CSCvq57414	HSRP/VRRP packets failed to flood locally in a service leaf switch, which causes a dual active state.	14.2(2e) and later
CSCvq68352	A last hop router does not generate the S,G tree.	14.2(2e) and later
CSCvq87843	There is an RPF lookup failure for MRIB.	14.2(2e) and later
CSCvq98750	In Cisco ACI when using MAC pinning with a vPC, prior to reloading when you run the 'show vpc brief' command on the CLI, the command shows that the vPC is passing consistency checks. However, after reloading the leaf switch, the vPC then properly displays the consistency check as 'Not Applicable'.	14.2(2e) and later
CSCvr02421	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(2e) and later
CSCvr05132	Traffic drops occur if there is redirect traffic towards a remote leaf switch and a Layer 2 service device is deployed behind a remote leaf switch in RL Direct mode.	14.2(2e) and later
CSCvr09108	An interface does not come up when a new link is connected. However, from the DOM data, the signals are present.	14.2(2e) and later
CSCvr10127	1) Deploy the breakout configuration. 2) Deploy a port channel or vPC configuration on these broken-out ports. 3) Remove the breakout configuration. The port channel or vPC configuration is still present in the APIC. 4) Deploy the breakout configuration. This action could cause the Hal process to spin and console stuck. This issue occurs when the vPC or port channel configuration is present even before the breakout is applied.	14.2(2e) and later
CSCvr11460	In a source outside and receiver inside scenario, PIM on the NBL is computing the stripe winner, when there is local receiver. As a result of this, PIM sends out joins over the fabric. However, the route is not currently propagated to the other border leaf switches, which is the expected behavior.	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvr12912	A switch reloads due to a sysmgr heartbeat failure and sysmgr HAP reset.	14.2(2e) 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(2e) and later
CSCvr17706	There is a kernel-panic out of memory crash. The following logs appear in the kernel traces: Kernel panic - not syncing: Out of memory: system-wide panic_on_oom is enabled	14.2(2e) and later
CSCvr27624	A switch reloads due to a sysmgr heartbeat failure and sysmgr HAP reset. This bug is a duplicate of CSCvr12912.	14.2(2e) and later
CSCvr27776	The sysmgr core got cleaned up as part of last core restore logic upon system reboot.	14.2(2e) and later
CSCvr31315	The SPAN manager process crashes when a SPAN session is deleted.	14.2(2e) and later
CSCvr31458	When a physical interface or transceiver is defective or has misconfigured speed, the fault displayed in the GUI has additional whitespaces.	14.2(2e) and later
CSCvr46681	Leaf switches crash and generate a core file after invoking the ACI snapshot rollback.	14.2(2e) and later
CSCvr46867	A Cisco ACI modular spine switch (N9504 chassis) with redundant supervisor modules (N9K-SUP-A) had an unexpected series of switchovers during a 6 minute period.	14.2(2e) and later
CSCvr47042	After removing a transceiver or cable from the interface, the port LED remains green. A port is physically down, but the "show interface" command says that the port is still up.	14.2(2e) and later
CSCvr49904	Traffic with a UDP destination port of 8472 is dropped on ingress by the ACI fabric.	14.2(2e) and later
CSCvr50031	The iBash "show interface ethernet <portnum>" command does not show CRC and stomped CRC errors.	14.2(2e) and later
CSCvr51479	MAC counters are reset when using the port shutdown command on the N9K-C9316D-GX and N9K-C93600CD-GX platforms. As a result, the MAC counters are not preserved.	14.2(2e) and later
CSCvr52832	There are no aggregated stats from L3If to L3ExtOut.	14.2(2e) and later
CSCvr60813	If fabric bidirectional forwarding detection (BFD) is enabled, the tenant BFD flaps upon enabling or disabling priority flow control (PFC) globally on any class when the peer node is an N9K-C9316D-GX or N9K-C93600CD-GX switch.	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvr76947	After upgrading leaf switches and after the switches come online on the target firmware version, reloading the chassis causes a failure to boot and a crash to the Loader> prompt with nothing left in the bootflash from which to boot.	14.2(2e) and later
CSCvr79911	An LLDP/CDP MAC address entry gets stuck in the blade switch table on a leaf switch in a vPC. The entry can get stuck if the MAC address flaps and hits the move detection interval, which stops all learning for the address. Use the following command to verify if a switch has a stale MAC address entry: module-1# show system internal epmc bladeswitch_mac all	14.2(2e) and later
CSCvr80292	A remote leaf switch is stuck in the "inactive" state after being registered into the fabric.	14.2(2e) and later
CSCvr83337	A Cisco ACI leaf switch unexpectedly reloads and generates a core file.	14.2(2e) and later
CSCvr84189	10Gbase-ER optics will not link up.	14.2(2e) and later
CSCvr87506	A leaf switch reboots 2 times, and on the 2nd reboot the state is not preserved. This causes the front panel ports to come online before being correctly configured and connected into the fabric, which can have traffic impact.	14.2(2e) and later
CSCvr87847	PIM border leaf switches do not see each other as PIM neighbors, resulting in potentially both being the stripe-winner and sending upstream joins. This can result in duplicate packets.	14.2(2e) and later
CSCvr88009	The Netflow (nfm) process crashes during configuration changes.	14.2(2e) and later
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(2e) and later
CSCvr95697	On a border leaf switch, some of the routes that are removed from the routing table are found to be not removed from BGP VPNv4 prefixes.	14.2(2e) and later
CSCvr98827	Some of the control plane packets are incorrectly classified as the user class and are reported as dropped in single chip spine switches. The statistics are incorrect because the packets are not actually dropped.	14.2(2e) and later
CSCvs02955	When running "show system internal epm endpoint all summary" on an FX leaf, the command output is cut short.	14.2(2e) and later
CSCvs04956	There is a memory leak with svc_ifc_streame.	14.2(2e) and later
CSCvs06119	Multiple switches crash and generate a core file at same time when applying an NTP policy.	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvs08304	The spine outerdstip, which indicates that the egress TEP is connecting to the Tetration network, is not updated when an egress L3Out in the mgmt:inb VRF fails over to a redundant L3Out on another leaf switch.	14.2(2e) and later
CSCvs10395	Leaf switch downlinks all go down at one time due to FabricTrack.	14.2(2e) and later
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(2e) and later
CSCvs29384	Fault F3525 (high SSD usage) is observed.	14.2(2e) and later
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(2e) and later
CSCvs34065	The "get_bkout_cfg failed" error displays when the following vsh_lc cli command is executed: vsh_lc -c "show system internal port-client event-history all"	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvs40299	The policy_mgr process on an ACI leaf switch has a memory leak and results in an unexpected reload. The problem can happen over a long period of time, such as a year. Depending on when individual switches were last rebooted, multiple devices could experience the reload at around the same time.	14.2(2e) and later
CSCvs41818	Port 1/2 on N9k-C9364C flaps continuously and does not come up.	14.2(2e) and later
CSCvs45414	A N9K-X9736PQ linecard in an ACI mode Nexus 9500 spine switch unexpectedly reloads. The following output is seen in the command "show system reset-reason module 1": `show system reset-reason module 1` ********** module reset reason (1) ********** 0) At 2019-12-01T00:00:00.00 Reason: line-card-not-responding Service:Line card not responding => [Failures < MAX] : powercycle Version:	14.2(2e) and later
CSCvs49377	After a virtual machine is vMotioned, traffic begins to drop the source from that endpoint. When running "show logging ip access-list internal packet-log deny" on the leaf switch, you can see policy drops for the endpoint.	14.2(2e) and later
CSCvs56556	A local AS configuration is not applied to the eBGP neighbor on a Cisco ACI border leaf switch, which results in the switch sending the fabric ASN (configured in the BGP Route Reflector policy) in the OPEN messages, which makes the neighbor reject the session because of the "bad remote-as" reason.	14.2(2e) and later
CSCvs56978	Connectivity between a server EPG and external L3Out EPG can be broken for some subnets that are configured with an external subnet for an external EPG.	14.2(2e) and later
CSCvs57186	After a link to a Cisco ACI leaf switch flaps, ARP continuously refreshes, and unicast traffic to a neighboring device is non-functional. In a packet capture, the leaf switch continuously sends ARP requests for the neighboring device, even though that device is sending ARP responses. When running "show ip arp vrf tenant:vrf", the age of the ARP entry is always 0 seconds.	14.2(2e) and later
CSCvs60566	A vPC pair of leaf switches go into the split brain mode, causing traffic duplication.	14.2(2e) and later
CSCvs61270	A modular spine switch gets stuck during upgrade.	14.2(2e) and later
CSCvs66240	During leaf switch replacement, vPC ports (map-pinning mode) are flapped on existing switches in the vPC pair.	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvs76848	A switch SSD fails in less than two years and needs replacement. The /mnt/pss/ssd_log_amp.log file shows daily P/E cycles increasing by 10 or more each day, and fault "F3525: High SSD usage" is observed. Check the switch activity and contact Cisco Technical Support if the "High SSD usage" fault is raised on the switch.	14.2(2e) and later
CSCvs77436	Error message "No handlers could be found for logger "root"" appears when doing a moquery for certain objects.	14.2(2e) 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(2e) and later
CSCvs83366	The COPP per interface policy feature has following limitation: The TCAM entry maximum for per interface per protocol is 256. After the threshold is exceeded, a fault will be raised. For more information, see: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/security/b-Cisco-APIC-Security-Configuration-Guide-411/b-Cisco-APIC-Security-Configuration-Guide-411_chapter_01001.html This enhancement request is to have a CLI command to display the number of required ACLs/TCAM entries for the COPP policy that is applied on a per-interface level. The command should possibly display whether the configuration will succeed or not based on the current overall TCAM usage.	14.2(2e) and later
CSCvs89617	Some ARP packets get dropped across the Cisco ACI fabric.	14.2(2e) and later
CSCvs96869	A leaf switch will crash with a vntag_mgr HAP reset and generate a core file.	14.2(2e) and later
CSCvt00231	Traffic destined to a switch is policy dropped. The contracts configured on the switch look correct, but the ELAM drop reason shows a clear SECURITY_GROUP_DENY. If you dump the FPC and FPB pt.index results of the ELAM, the values are different. Specifically, the FPC index is wrong when you check the Stats Idx under the specific ACLQOS rule. FPC should be the summary of the final result. In this case, there are two hits, but there is one stable entry in TCAM and one that is not stable.	14.2(2e) and later
CSCvt08181	All routes to a particular spine switch are removed from uRIB on all leaf switches in the fabric.	14.2(2e) and later
CSCvt09775	The policy element crashes due to database space exhaustion with B22 FEX devices in the fabric.	14.2(2e) and later
CSCvt25383	The pervasive static route is missing on the spine node.	14.2(2e) and later
CSCvt35002	A link intermittently flaps on leaf switch fabric ports that are connected to a spine switch.	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvt39689	Glean ARP (0xfff2, 239.255.255.240) flood is stopped on the transit leaf switch and is not delivered toward all the leaf switches in the fabric. Thus, silent host discovery does not work.	14.2(2e) and later
CSCvt52364	A leaf switch reloaded with an NFM process core.	14.2(2e) and later
CSCvt52620	There is a stale pervasive route after a DHCP relay label is deleted.	14.2(2e) and later
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(2e) and later
CSCvt64042	The policy element crashes once during a misconfiguration.	14.2(2e) and later
CSCvt64733	An ARP request from the endpoint behind the remote leaf switch is received on the ToR 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 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.	14.2(2e) and later
CSCvt73069	A Cisco ACI fabric is not fully fit after a Cisco APIC firmware upgrade.	14.2(2e) 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(2e) and later
CSCvt94039	A leaf switch crashes and reloads due to "nfm hap reset".	14.2(2e) and later
CSCvu01639	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(2e) 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(2e) and later
CSCvu08065	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(2e) and later
CSCvu15712	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(2e) and later
CSCvu15751	The following event can be seen on the spine node: [E4204936][transition][warning][sys] %URIB-4-SYSLOG_SL_MSG_WARNING: URIB-5-RPATH_DELETE: message repeated 1 times in last 220162 sec	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvu16987	A Cisco ACI leaf switch reboots due to an ICMPv6 HAP reset.	14.2(2e) 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(2e) and later
CSCvu26947	If a rogue file grows too large, it can cause out of memory condition on a spine switch or leaf 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(2e) and later
CSCvu27617	Fault F0098 gets raised and the statistics in the Capacity dashboards do not show for the affected leaf nodes.	14.2(2e) and later
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(2e) and later
CSCvu32603	The AVAGO QSFP-100G-SR4 transceiver displays an abnormal Rx dbm value in a down port.	14.2(2e) and later
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(2e) 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(2e) and later
CSCvu53558	When walking through SNMP targets, SNMP generates a core file on a spine switch.	14.2(2e) and later
CSCvu61024	Zoning-rules are not programmed in the hardware after reloading a switch.	14.2(2e) and later
CSCvu72416	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(2e) 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(2e) 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(2e) and later
CSCvv13722	A leaf switch crashes due to a routing loop in the IPFIB process.	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvv17496	A FEX link takes a long time (5+ minutes) to come up.	14.2(2e) and later
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(2e) 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(2e) 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(2e) 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(2e) and later
CSCvv78885	A stale route map entry is causes unexpected route leaking.	14.2(2e) and later
CSCvv95800	A spine switch reloads unexpectedly due to the service on the linecard having a hap-reset.	14.2(2e) 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(2e) 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(2e) and later
CSCvw19911	After disabling unicast routing from a bridge domain, the static pervasive route is still present on the leaf switch.	14.2(2e) and later
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(2e) 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(2e) and later

Cisco Confidential

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(2e) 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(2e) 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(2e) and later
CSCwb08081	A route profile that matches on community list and sets the local pref and community is not working post upgrade to 5.2.x release. route-map imp-l3out-L3OUT_WAN-peer-2359297, permit, sequence 4201 Match clauses: community (community-list filter): peer16389-2359297-exc-ext-in-L3OUT_WAN_COMMUNITY-rgcom Set clauses: local-preference 200 community xxxxx:101 xxxxx:500 xxxxx:601 xxxxy:4 additive The match clause works as expected, but the set clause is ignored.	14.2(2e) and later
CSCwb17229	The sysmgr process crashes unexpectedly, causing the line card to reload.	14.2(2e) 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(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCwd29346	An ACI switch's console may continuously output messages similar to: svc_ifc_eventmg (*****) Ran 7911 msecs in last 7924 msecs	14.2(2e) and later
CSCvr57536	An ACI node reloads due to the Machine Check Exception error, similar to the following output: [603029.390562] sbridge: HANDLING MCE MEMORY ERROR [603029.390563] CPU 0: Machine Check Exception: 0 Bank 7: 8c00004000010091 [603029.390564] TSC 0 ADDR 2e3d3f40 MISC 140545486 PROCESSOR 0:50663 TIME 1569464793 SOCKET 0 APIC 0 [603029.390710] sbridge: HANDLING MCE MEMORY ERROR	14.2(2e)
CSCvr62348	There is a slow memory leak during line card insertion or removal, FEX reloads, or transceiver removal. This is issue occurs more frequently for copper ports.	14.2(2e)
CSCvr73276	With N9K-C9348GC-FXP top-of-rack switches, on ports where PoE is enabled in auto mode, there is a potential memory leak can be seen in fault scenarios, such as the short-ckt fault, overcurrent, or max current. Leaks can also be observed when multiple negotiations occur with a powered device (PD) or when EPG or VLAN information on a PoE interface policy is changed multiple times.	14.2(2e)
CSCvr75413	After upgrading a leaf switch, the switch brings up the front panel ports before the policies are programmed. This may cause a connectivity issue if a connected host relies on the link level state to decide whether or not it can forward traffic on a particular NIC or port. The loss duration would be proportional to the scale of configuration policies that must be programmed.	14.2(2e)
CSCvr86930	When PoE is configured on ports on doing a stateful reload--that is, manually reloading a switch with the config restore option--a PoE power adjust can happen during the same interval when APIC discovery occurs. This results in a PoE core, as it expects the objstore always to return a desired interface object that may have been updated.	14.2(2e)
CSCvv04251	A policy upgrade fails on the N9K-C9316D-GX and N9K-C93600CD-GX platforms that have the hw_changes_bit programmed as "0x03."	14.2(2e)

Resolved Issues

This section lists the resolved bugs. 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.

Table 19 Resolved Issues in This Release

Bug ID	Description	Fixed in

Cisco Confidential

Bug ID	Description	Fixed in
CSCvr63689	The ACI N93360YC-FX2 leaf switch becomes inactive.	14.2(2g)
CSCvs00854	A switch's DC PSU went to the "shut" or "failed" state, and the PSU LED was amber color.	14.2(2g)
CSCvr57536	An ACI node reloads due to the Machine Check Exception error, similar to the following output: [603029.390562] sbridge: HANDLING MCE MEMORY ERROR [603029.390563] CPU 0: Machine Check Exception: 0 Bank 7: 8c00004000010091 [603029.390564] TSC 0 ADDR 2e3d3f40 MISC 140545486 PROCESSOR 0:50663 TIME 1569464793 SOCKET 0 APIC 0 [603029.390710] sbridge: HANDLING MCE MEMORY ERROR	14.2(2f)
CSCvr62348	There is a slow memory leak during line card insertion or removal, FEX reloads, or transceiver removal. This is issue occurs more frequently for copper ports.	14.2(2f)
CSCvr73276	With N9K-C9348GC-FXP top-of-rack switches, on ports where PoE is enabled in auto mode, there is a potential memory leak can be seen in fault scenarios, such as the short-ckt fault, overcurrent, or max current. Leaks can also be observed when multiple negotiations occur with a powered device (PD) or when EPG or VLAN information on a PoE interface policy is changed multiple times.	14.2(2f)
CSCvr75413	After upgrading a leaf switch, the switch brings up the front panel ports before the policies are programmed. This may cause a connectivity issue if a connected host relies on the link level state to decide whether or not it can forward traffic on a particular NIC or port. The loss duration would be proportional to the scale of configuration policies that must be programmed.	14.2(2f)
CSCvr86930	When PoE is configured on ports on doing a stateful reload--that is, manually reloading a switch with the config restore option--a PoE power adjust can happen during the same interval when APIC discovery occurs. This results in a PoE core, as it expects the objstore always to return a desired interface object that may have been updated.	14.2(2f)
CSCvk73228	This is an enhancement to decode the binary logs offline directly from the techsupport.	14.2(2e)
CSCvq42672	The EPMC process crashes continuously.	14.2(2e)
CSCvq97092	The N2348TQ FEX randomly reboots. A crash in the 'tiburon' and/or 'ethpc' service may be observed in the syslogs immediately prior to the reload event.	14.2(2e)
CSCvq99155	On leaf or spine switches, LDAP authentication requests might get sent out of the out-of-band interface (eth0 or eth6) even though the LDAP provider is configured to use an in-band EPG.	14.2(2e)
CSCvr19396	On deleting the virtual MAC address on the APIC, the switch processes will not delete the virtual Mac address info from the database. This results in some of the switch modules using the deleted virtual mac address for host-tracking endpoints and for routing the traffic if the destination MAC address is the virtual MAC address.	14.2(2e)
CSCvr52067	A false positive lifetime endurance fault is generated for a new SSD (model SHMST064G3FECTLP51) on the Cisco APIC.	14.2(2e)

Cisco Confidential

Known Behaviors

This section lists bugs that describe known behaviors. 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(2) releases in which the known behavior exists. A bug might also exist in releases other than the 14.2(2) releases.

Table 20 Known Behaviors in This Release

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(2e) 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(2e) and later
CSCup65586	The show interface command shows the tunnel's Rx/Tx counters as 0.	14.2(2e) 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(2e) and later
CSCup92534	Continuous "threshold exceeded" messages are generated from the fabric.	14.2(2e) and later
CSCuq39829	Switch rescue user ("admin") can log into fabric switches even when TACACS is selected as the default login realm.	14.2(2e) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	14.2(2e) 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(2e) 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(2e) 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(2e) and later
CSCuq93389	If an endpoint has multiple IPs, the endpoint will not be aged until all IPs go silent. If one of the IP addresses is reassigned to another server/host, the fabric detects it as an IP address move and forwarding will work as expected.	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	14.2(2e) and later
CSCur81822	The access-port operational status is always "trunk".	14.2(2e) 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(2e) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	14.2(2e) 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(2e) 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(2e) 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(2e) 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(2e) and later
CSCuu11347	Traffic from the orphan port to the vPC pair is not recorded against the tunnel stats. Traffic from the vPC pair to the orphan port is recorded against the tunnel stats.	14.2(2e) 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(2e) 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(2e) 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(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
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(2e) and later
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.	14.2(2e) 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(2e) 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(2e) 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(2e) 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(2e) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	14.2(2e) and later
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	14.2(2e) and later
CSCuy25780	An overlapping or duplicate prefix/subnet could cause the valid prefixes not to be installed because of batching behavior on a switch. This can happen during an upgrade to the 1.2(2) release.	14.2(2e) 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(2e) 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(2e) 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(2e) and later
CSCuy96912	The debounce timer is not supported on 25G links.	14.2(2e) 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(2e) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCuz13810	When subnet check is enabled, a ToR does not learn IP addresses locally that are outside of the bridge domain subnets. However, the packet itself is not dropped and will be forwarded to the fabric. This will result in such IP addresses getting learned as remote endpoints on other ToRs.	14.2(2e) and later
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	14.2(2e) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	14.2(2e) 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(2e) 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(2e) and later
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	14.2(2e) 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(2e) 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(2e) 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(2e) 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(2e) and later
CSCvo39715	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(2e) and later
CSCvp04772	During an upgrade on a dual-SUP system, the standby SUP may go into a failed state.	14.2(2e) 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(2e) and later
CSCvq71034	There is a policy drop that occurs with L3Out transit cases.	14.2(2e) and later

Cisco Confidential

Bug ID	Description	Exists In
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(2e) and later

- 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.
- The following properties within a QoS class under "Global QoS Class policies," should not be changed from its 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)
- 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.
- 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.
- 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.
- 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.

The following list describes IGMP snooping known behaviors:

- Multicast router functionality is not supported when IGMP queries are received with VxLAN encapsulation.

Cisco Confidential

- 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.
- The rate of the number of IGMP reports sent to a leaf switch should be limited to 1000 reports per second.
- 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.
- 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.
- 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.

The following list describes IpEpg (IpCkt) known behaviors:

- An IP/MAC Ckt endpoint configuration is not supported in combination with static endpoint configurations.
- 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.
- 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.
- 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.
- An IP/MAC Ckt endpoint configuration is not supported with dynamic endpoint groups. Only static endpoint groups are supported.
- 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.
- Dynamic deployment of contracts based on instrImmedcy set to onDemand/lazy not supported; only immediate mode is supported.

The following list describes direct server return (DSR) known behaviors:

- 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

- 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.
- Direct server return is not supported for shared services. Direct server return endpoints cannot be spread around different virtual routing and forwarding (VRF) contexts.
- Configurations for a virtual IP address can only be /32 or /128 prefix.
- 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.
- 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).
- Learning through GARP will work only in ARP Flood Mode.

Related Documentation

The Cisco Application Policy Infrastructure Controller (APIC) documentation can be accessed from the following website:

https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

Cisco Confidential

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.

© 2019-2025 Cisco Systems, Inc. All rights reserved.