



Cisco NX-OS Release Notes for Cisco Nexus 9000 Series ACI-Mode Switches, Release 14.1(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 4.1(2)*, which you can view at the following location:

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

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

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

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

Table 1 shows the online change history for this document.

Table 1. Online History Change

Date	Description
June 11, 2019	14.1(2g): Release 14.1(2g) became available.
July 2, 2019	14.1(2g): In the Open Bugs section, added bug CSCvq24680.
July 16, 2019	14.1(2g): In the Open Bugs section, added bug CSCvq53300. In the Supported Hardware section, for the N9K-C93360YC-FX2 switch, added the following note: Note: The supported total number of fabric ports and port profile converted fabric links is 64.
July 31, 2019	In the Compatibility Information section, added the following bullet: ■ 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.

Date	Description
August 12, 2019	14.1(2m): Release 14.1(2m) became available. Added the resolved bugs for this release.
August 28, 2019	14.2(2g): In the Open Bugs section, added bugs CSCvq42673 and CSCvq43477.
August 31, 2019	14.1(2o): Release 14.1(2o) became available. Added the resolved bugs for this release.
September 3, 2019	14.1(2g): In the Open Bugs section, added bugs CSCvp94661.
September 11, 2019	<p>In the Supported Hardware section, for the N9K-C9348GC-FXP, N9K-C93108TC-FX, and N9K-C93180YC-FX switches, added the following note:</p> <p>Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.</p>
September 20, 2019	<p>In the Usage Guidelines section, added the following bullet:</p> <ul style="list-style-type: none"> ■ 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.
September 27, 2019	<p>In the Supported Hardware section, for the N9K-C9336C-FX2 switch, changed the port profile note to:</p> <p>The port profile feature does supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.</p>
October 7, 2019	14.1(2g): In the Open Bugs section, added bug CSCvr14376.
October 16, 2019	14.1(2s): Release 14.1(2s) became available. Added the open and resolved bugs for this release.
October 18, 2019	14.1(2s): In the Resolved Bugs section, added bug CSCvq28541.
October 24, 2019	14.1(2s): In the Open Bugs section, added bug CSCvr73276.
October 30, 2019	14.1(2u): Release 14.1(2u) became available. Added the resolved bugs for this release.

Contents

This document includes the following sections:

- [Contents](#)
- [Supported Hardware](#)
- [Supported FEX Models](#)
- [New and Changed Information](#)
- [Installation Notes](#)
- [Compatibility Information](#)
- [Usage Guidelines](#)
- [Bugs](#)
- [Related Documentation](#)

Supported Hardware

Table 2 lists the hardware that the Cisco Nexus 9000 Series ACI Mode switches support.

Table 2 Cisco Nexus 9000 Series Hardware

Hardware Type	Product ID	Description
Chassis	N9K-C9504	Cisco Nexus 9504 chassis with 4 I/O slots
Chassis	N9K-C9508	Cisco Nexus 9508 chassis with 8 I/O slots
Chassis component	N9K-C9508-FAN	Fan tray
Chassis component	N9K-PAC-3000W-B	Cisco Nexus 9500 3000W AC power supply, port side intake
Pluggable module (GEM)	N9K-M12PQ	12-port or 8-port
Pluggable module (GEM)	N9K-M6PQ	6-port
Pluggable module (GEM)	N9K-M6PQ-E	6-port, 40 Gigabit Ethernet expansion module
Spine switch	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.
Spine switch	N9K-C9336PQ	Cisco Nexus 9336PQ switch, 36-port 40 Gigabit Ethernet QSFP Note: 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.

Supported Hardware

Hardware Type	Product ID	Description
Spine switch	N9K-C9364C	<p>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.</p> <p>The following PSUs are supported for the N9K-C9364C:</p> <ul style="list-style-type: none"> ■ NXA-PAC-1200W-PE ■ NXA-PAC-1200W-PI ■ N9K-PUV-1200W ■ NXA-PDC-930W-PE ■ NXA-PDC-930W-PI <p>Note: 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.</p> <p>A 930W-DC PSU (NXA-PDC-930W-PE or NXA-PDC-930W-PI) is supported in redundancy mode if 3.5W QSFP+ modules or passive QSFP cables are used and the system is used in 40C ambient temperature or less; for other optics or a higher ambient temperature, a 930W-DC PSU is supported only with 2 PSUs in non-redundancy mode.</p> <p>1-Gigabit QSA is not supported on ports 1/49-64.</p>
Spine switch	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
Spine switch	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
Spine switch	N9K-C9516	Cisco Nexus 9516 switch with 16 line card slots
Spine switch	N9K-X9736Q-FX	The Cisco Nexus 9736Q-FX is a 36-port, 40 Gigabit Ethernet QSFP28 spine line card.
Spine switch fan	N9K-C9300-FAN3	Port side intake fan
Spine switch fan	N9K-C9300-FAN3-B	Port side exhaust fan

Supported Hardware

Hardware Type	Product ID	Description
Spine switch module	N9K-C9504-FM	Cisco Nexus 9504 fabric module supporting 40 Gigabit line cards
Spine switch module	N9K-C9504-FM-E	Cisco Nexus 9504 fabric module supporting 100 Gigabit line cards
Spine switch module	N9K-C9508-FM	Cisco Nexus 9508 fabric module supporting 40 Gigabit line cards
Spine switch module	N9K-C9508-FM-E	Cisco Nexus 9508 Fabric module supporting 100 Gigabit line cards
Spine switch module	N9K-C9508-FM-E2	Cisco Nexus 9508 Fabric module supporting 100 Gigabit line cards
Spine switch module	N9K-C9516-FM-E2	Cisco Nexus 9516 Fabric module supporting 100 Gigabit line cards
Spine switch module	N9K-X9732C-EX	Cisco Nexus 9500 32-port, 40/100 Gigabit Ethernet QSFP28 aggregation module Note: The N9K-X9732C-EX line card cannot be used when a fabric module is installed in FM slot 25.
Spine switch module	N9K-X9736C-FX	Cisco Nexus 9736 36-port, 40/100 Gigabit Ethernet QSFP28 aggregation module Note: 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.
Spine switch module	N9K-X9736PQ	Cisco Nexus 9500 36-port, 40 Gigabit Ethernet QSFP aggregation module
Switch module	N9K-SC-A	Cisco Nexus 9500 Series system controller
Switch module	N9K-SUP-A	Cisco Nexus 9500 Series supervisor module
Switch module	N9K-SUP-A+	Cisco Nexus 9500 Series supervisor module
Switch module	N9K-SUP-B	Cisco Nexus 9500 Series supervisor module
Switch module	N9K-SUP-B+	Cisco Nexus 9500 Series supervisor module
Top-of-rack (ToR) leaf switch	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

Supported Hardware

Hardware Type	Product ID	Description
Top-of-rack (ToR) leaf switch	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. <i>Note:</i> The supported total number of fabric ports and port profile converted fabric links is 64.
Top-of-rack (ToR) leaf switch	N9K-93240YC-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-93240YC-FX2 is a 1.2-RU switch. <i>Note:</i> 10/25G-LR-S with QSA is not supported.
Top-of-rack (ToR) leaf switch	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.
Top-of-rack (ToR) leaf switch	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. <i>Note:</i> 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.
Top-of-rack (ToR) leaf switch	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
Top-of-rack (ToR) leaf switch	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

Hardware Type	Product ID	Description
Top-of-rack (ToR) leaf switch	N9K-C93180LC-EX	<p>Cisco Nexus 9300 platform switch with 24 40-Gigabit front panel ports and 6 40/100-Gigabit QSFP28 spine-facing ports</p> <p>The switch can be used either 24 40G ports or 12 100G ports. If 100G is connected the Port1, Port 2 will be HW disabled.</p> <p>Note: This switch has the following limitations:</p> <ul style="list-style-type: none"> ■ The top and bottom ports must use the same speed. If there is a speed mismatch, the top port takes precedence and bottom port will be error disabled. Both ports both must be used in either the 40 Gbps or 10 Gbps mode. ■ Ports 26 and 28 are hardware disabled. ■ This release supports 40 and 100 Gbps for the front panel ports. The uplink ports can be used at the 100 Gbps speed. ■ Port profiles and breakout ports are not supported on the same port.
Top-of-rack (ToR) leaf switch	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
Top-of-rack (ToR) leaf switch	N9K-C93180YC-FX	<p>Cisco Nexus 9300 platform switch with 48 1/10/25-Gigabit Ethernet SFP28 front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. The SFP28 ports support 1-, 10-, and 25-Gigabit Ethernet connections and 8-, 16-, and 32-Gigabit Fibre Channel connections.</p> <p>Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.</p>
Top-of-rack (ToR) leaf switch	N9K-C9332PQ	Cisco Nexus 9332PQ Top-of-rack (ToR) Layer 3 switch with 26 APIC-facing ports and 6 fixed-Gigabit spine facing ports.

Supported Hardware

Hardware Type	Product ID	Description
Top-of-rack (ToR) leaf switch	N9K-C9336C-FX2	<p>Cisco Nexus C9336C-FX2 Top-of-rack (ToR) switch with 36 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports.</p> <p>Note: 1-Gigabit QSA is not supported on ports 1/1-6 and 1/33-36. The port profile feature does supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.</p>
Top-of-rack (ToR) leaf switch	N9K-C9348GC-FXP	<p>The Cisco Nexus 9348GC-FXP switch (N9K-C9348GC-FXP) is a 1-RU fixed-port, L2/L3 switch, designed for ACI deployments. This switch has 48 100/1000-Megabit 1GBASE-T downlink ports, 4 10-/25-Gigabit SFP28 downlink ports, and 2 40-/100-Gigabit QSFP28 uplink ports.</p> <p>This switch supports the following PSUs:</p> <ul style="list-style-type: none"> ■ NXA-PAC-350W-PI ■ NXA-PAC-350W-PE <p>Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.</p>
Top-of-rack (ToR) leaf switch	N9K-C9358GY-FXP	<p>The Cisco Nexus 9358GY-FXP switch (N9K-C9358GY-FXP) is a 2-RU, fixed-port switch designed for spine-leaf-APIC deployment in data centers. This switch has eight 100-Megabit/1-Gigabit BASE-T ports and two 1/10/25-Gigabit SFP28 ports. Ports 1-8 support PoE/PoE+ and are software configurable.</p>
Top-of-rack (ToR) leaf switch	N9K-C9372PX	<p>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</p> <p>Note: Only the downlink ports 1-16 and 33-48 are capable of supporting SFP1-10G-ZR SFP+.</p>
Top-of-rack (ToR) leaf switch	N9K-C9372PX-E	<p>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</p> <p>Note: Only the downlink ports 1-16 and 33-48 are capable of supporting SFP1-10G-ZR SFP+.</p>

Supported Hardware

Hardware Type	Product ID	Description
Top-of-rack (ToR) leaf switch	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
Top-of-rack (ToR) leaf switch	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
Top-of-rack (ToR) leaf switch	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
Top-of-rack (ToR) leaf switch	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
Top-of-rack (ToR) leaf switch fan	NXA-FAN-30CFM-B	Port side exhaust fan
Top-of-rack (ToR) leaf switch fan	NXA-FAN-30CFM-F	Port side intake fan
Top-of-rack (ToR) leaf switch power supply unit	N9K-PAC-1200W	1200W AC Power supply, port side intake pluggable <i>Note:</i> This power supply is supported only by the Cisco Nexus 93120TX, 93128TX, and 9336PQ ACI-mode switches
Top-of-rack (ToR) leaf switch power supply unit	N9K-PAC-1200W-B	1200W AC Power supply, port side exhaust pluggable <i>Note:</i> This power supply is supported only by the Cisco Nexus 93120TX, 93128TX, and 9336PQ ACI-mode switches
Top-of-rack (ToR) leaf switch power supply unit	N9K-PAC-650W	650W AC Power supply, port side intake pluggable
Top-of-rack (ToR) leaf switch power supply unit	N9K-PAC-650W-B	650W AC Power supply, port side exhaust pluggable

Supported FEX Models

Hardware Type	Product ID	Description
Top-of-rack (ToR) leaf switch power supply unit	N9K-PUV-1200W	1200W HVAC/HVDC dual-direction airflow power supply <i>Note:</i> This power supply is supported only by the Cisco Nexus 93120TX, 93128TX, and 9336PQ ACI-mode switches
Top-of-rack (ToR) leaf switch power supply unit	N9K-PUV-3000W-B	3000W AC Power supply, port side exhaust pluggable
Top-of-rack (ToR) leaf switch power supply unit	NXA-PAC-1200W-PE	1200W AC Power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance <i>Note:</i> This power supply is supported only by the Cisco Nexus 93120TX, 93128TX, and 9336PQ ACI-mode switches.
Top-of-rack (ToR) leaf switch power supply unit	NXA-PAC-1200W-PI	1200W AC Power supply, port side intake pluggable, with higher fan speeds for NEBS compliance <i>Note:</i> This power supply is supported only by the Cisco Nexus 93120TX, 93128TX, and 9336PQ ACI-mode switches.
Top-of-rack (ToR) leaf switch power supply unit	NXA-PDC-440W-PI	440W DC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance <i>Note:</i> This power supply is supported only by the Cisco Nexus 9348GC-FXP ACI-mode switch.
Top-of-rack (ToR) leaf switch power supply unit	UCSC-PSU-930WDC V01	Port side exhaust DC power supply compatible with all ToR leaf switches
Top-of-rack (ToR) leaf switch power supply unit	UCS-PSU-6332-DC	930W DC power supply, reversed airflow (port side exhaust)

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
- New Software Features

New Hardware Features

The following hardware features are now available:

- 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

New Software Features

For new software features, see the *Cisco Application Policy Infrastructure Controller Release Notes, Release 4.1(2)* at the following location:

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

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 ACI Fabric Hardware Installation Guide* at the following location:

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

Compatibility Information

- 100mb optics, such as the GLC-TE, are supported in 100mb speed only on -EX and -FX switches, such as the N9K-C93180YC-EX and N9K-C93180YC-FX switches, 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.

Usage Guidelines

- 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 N2348UPO to ACI leaf switches, the following options are available:
 - Directly connect the 40G FEX ports on the N2348UPO to the 40G switch ports on the ACI leaf switches
 - Break out the 40G FEX ports on the N2348UPO 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 N9332PQ 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.

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 4.1(2)* for policy information: <https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html>

- UDP DstPort 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

Bugs

- UDP SrcPort 123: NTP
- UDP DstPort 123: NTP

- 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 14.1(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.

- Link-level flow control is not supported on leaf switches that are running in ACI mode.
- 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.
- 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.

Bugs

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

- [Known Limitations](#)

Bugs

- [Open Bugs](#)
- [Resolved Bugs](#)
- [Known Behaviors](#)

Known Limitations

The following list describes IpEpg (IpCkt) known limitations in this release:

- 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 limitations in this release:

- 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.

Open Bugs

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 14.1(2) releases in which the bug exists. A bug might also exist in releases other than the 14.1(2) releases.

Table 3 Open Bugs in This Release

Bug ID	Description	Exists In
CSCvp92121	<p>A vulnerability in the Link Layer Discovery Protocol (LLDP) subsystem of Cisco Nexus 9000 Series Application Centric Infrastructure (ACI) Mode Switch Software could allow an adjacent, unauthenticated attacker to cause a denial of service (DoS) condition or execute arbitrary code with root privileges.</p> <p>The vulnerability is due to improper input validation of certain type, length, value (TLV) fields of the LLDP frame header. An attacker could exploit this vulnerability by sending a crafted LLDP packet to the targeted device. A successful exploit may lead to a buffer overflow condition that could either cause a DoS condition or allow the attacker to execute arbitrary code with root privileges.</p> <p>Note: This vulnerability cannot be exploited by transit traffic through the device; the crafted packet must be targeted to a directly connected interface.</p> <p>Cisco has released software updates that address this vulnerability. There are no workarounds that address this vulnerability.</p> <p>This advisory is available at the following link:</p> <p>https://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20190731-nxos-bo</p>	14.1(2g)
CSCvp92269	<p>Running a Qualys security scan results in the following message:</p> <p>CWE - 693 Protection Mechanism Failure -</p> <p>"HTTP Security Header Not Detected"</p>	14.1(2g)
CSCvq24680	<p>After a reboot of a leaf switch that is operating as a tier-1 leaf in a multi-tier ACI fabric, the leaf switch will be stuck in a reboot loop. The reason for the loop is because during boot, the sdkhal process crashes. You can see the crash when running "show cores" on the switch.</p>	14.1(2g)
CSCvq43058	<p>A spine switch fabric module or line card is reloaded unexpectedly due to a kernel panic. The stack trace includes the following statement:</p> <p>Kernel panic - not syncing: Out of memory: system-wide panic_on_oom is enabled</p>	14.1(2g)

Bugs

Bug ID	Description	Exists In
CSCvq43477	In the IPv6 options, for the source-link layer address field, IPv6 traffic is blackholed because the leaf switch sets the incorrect MAC address in the router advertisement's (RA's) source link-layer address. This happens only with RAs that are sent as a reply to the router solicitation from the host. Unsolicited RAs from the leaf switch have the correct MAC address of the leaf switch itself. The border leaf switch sends out unsolicited RA messages correctly with its link MAC address (0022.bdf8.19ff) in the source link-layer address field.	14.1(2g)
CSCvq53300	Tier 2 switches reboot every 10 minutes. Spine switches also reload once with the same IS-IS core file backtrace.	14.1(2g)
CSCvq65546	The PIM process crashes randomly.	14.1(2g)
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.	14.1(2g) 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.	14.1(2g) and later
CSCvh18100	If Cisco ACI Virtual Edge or AVS is operating in VxLAN non-switching mode behind a FEX, the traffic across the intra-EPG endpoints will fail when the bridge domain has ARP flooding enabled.	14.1(2g) and later
CSCvj23046	In Cisco ACI Multi-Site plus multi-pod topologies, there could be multicast traffic loss for about 30 seconds on the remote-site. If only one LC has fabric links, there are other LCs with no fabric links and the LC with fabric links is reloaded.	14.1(2g) and later
CSCvo57063	10G AOC shows up as 10G ACU when using passive QSA. There is no functionality impact; this is only a display issue.	14.1(2g) 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.	14.1(2g) and later
CSCvp90131	The Hardware Abstraction Layer (HAL) generates a core file.	14.1(2g) and later
CSCvp94661	There is an EPM crash on a leaf switch that receives the Endpoint Announce packet with a malformed length field.	14.1(2g) and later
CSCvp97665	Traffic loss may be observed for flows from local leaf switches to remote leaf switches when one of the remote leaf switches in a vPC is de-commissioned and commissioned again.	14.1(2g) and later
CSCvq01755	There is a traffic drop of approximately 30 to 116 seconds during the decommission and re-commission of a remote leaf switch vPC peer after RLD is enabled.	14.1(2g) and later
CSCvq05545	Downgrading from the 14.1(2) release to an earlier release from the APIC does not complete, and the status in the APIC firmware tab shows as unknown. Reload the node to recover it.	14.1(2g) and later
CSCvq08723	Remote leaf switch shared services local switching traffic drops occur from an orphan endpoint to a vPC endpoint.	14.1(2g) and later

Bugs

Bug ID	Description	Exists In
CSCvq64803	<p>A leaf switch crashes with the " Unknown" reset reason when the breakout ports configuration is re-applied.</p> <p>The reset reason for this switch is as follows:</p> <p style="padding-left: 40px;">Image Version : 13.2(3o)</p> <p style="padding-left: 40px;">Reset Reason (LCM): Unknown (0) at time Fri Jul 12 14:21:14 2019</p> <p style="padding-left: 40px;">Reset Reason (SW): Reset triggered due to HA policy of Reset (16) at time Fri Jul 12 14:17:40 2019</p> <p style="padding-left: 40px;">Service (Additional Info): Reset triggered due to HA policy of Reset</p>	14.1(2g) and later
CSCvr14376	The Cisco N9K-C93180YC-FX leaf switch sometimes unexpectedly reloads if you run the " show system internal epmc bltrace" command in the vsh_lc mode.	14.1(2g) and later
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.	14.1(2g) through 14.1(2s)
CSCvq40849	Some 100 Gbps uplink ports between a spine switch and leaf switch do not come up.	14.1(2g) through 14.1(2m)
CSCvq42673	<ol style="list-style-type: none"> 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 causes a port channel bringup failure, or causes the port channel manager or eth_port_manager to crash on the switch. <p>This issue occurs when the vPC or port channel configuration is present even before the breakout is applied.</p>	14.1(2g) through 14.1(2m)
CSCvr03623	When using a non-default SSM range, PIM and NGMVPN processes can be out of sync with respect to which border leaf switch is the stripe winner. This leads to the border leaf switch no longer sending PIM join upstream, and traffic within the SSM range is affected.	14.1(2g) through 14.1(2m)
CSCvq79143	When inter-VRF multicast is configured, chaining of the receiver VRFs is not allowed. Because of this restriction, if there is a source VRF A and receiver VRF B, the source VRF A cannot be a receiver VRF for the same group range. A configuration that chains the receiver VRFs, such as during bring up, can result in an MRIB crash.	14.1(2g) through 14.1(2o)
CSCvr62348	There is a slow memory leak during line card insertion or removal, FEX reloads, or transceiver removal. This issue occurs more frequently for copper ports.	14.1(2g) through 14.1(2o)

Bugs

Bug ID	Description	Exists In
CSCvn57265	When providing and consuming a deny contract with the 'log' or 'no stats' directive enabled between an in-band EPG and a separate EPG, the F1259 fault is raised on the fabric nodes. The "Rule failed due to software programming error" error displays, even though the contract is programmed and functioning.	14.1(2g) through 14.1(2s)
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	14.1(2g) through 14.1(2s)
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.	14.1(2g) through 14.1(2s)
CSCvr35108	Fibre Channel links do not come up when connecting Cisco Nexus 9000 switches in ACI mode to MDS 9396T switches on ports above 63.	14.1(2m) through 14.1(2o)
CSCvr12002	Traffic is lost after deleting and adding a tenant when using a non-default SSM range.	14.1(2o)
CSCvr73276	With N9K-C9348GC-FXP and N9K-C9358GY-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.	14.1(2s)

Resolved Bugs

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 4 Resolved Bugs in This Release

Bug ID	Description	Fixed in
CSCvk61394	The policy element (PE) keeps running at 100%.	14.1(2g)
CSCvn57265	When providing and consuming a Deny contract with the 'log' or 'no stats' directive enabled between an in-band EPG and an EPG, the F1259 fault is raised on the fabric nodes. The "Rule Failed due to software programming error" error displays, even though the contract is	14.1(2g)

Bugs

Bug ID	Description	Fixed in
	programmed and functioning.	
CSCvo43220	Leaf switches repeatedly crash after configuring more than 510 VMACs (common perversive gateway) on one leaf switch.	14.1(2g)
CSCvo46419	If Host Based route (HBR) is enabled on the same Bridge Domain (BD) where L3 Multicast source is present, multicast traffic will be dropped. Enabling HBR on any other BD has no impact.	14.1(2g)
CSCvo58703	ISIS show commands fail in ibash with errors saying that the time data format doesn't match format: d-leaf103# show isis database l1 vrf all time data '2019-02-27T17:21:' does not match format '%Y-%m-%dT%H:%M:%S.%f' time data '2019-02-27T17:26:' does not match format '%Y-%m-%dT%H:%M:%S.%f' time data '2019-02-27T17:31:' does not match format '%Y-%m-%dT%H:%M:%S.%f' time data '2019-02-27T17:23:' does not match format '%Y-%m-%dT%H:%M:%S.%f' time data '2019-02-27T17:23:' does not match format '%Y-%m-%dT%H:%M:%S.%f'	14.1(2g)
CSCvo58815	Policyelem is using all available switch memory. Leaf switch reloads with reason kernel-panic.	14.1(2g)
CSCvo65024	The " show ip sla statistics," " show ip sla statistics details," and " show ip sla statistics aggregate" commands have the same output.	14.1(2g)
CSCvo67919	You may see that after enabling BGP MD5 Authentication for multipod in the Fabric External Connection Policies under the POD Peering Profile that the bgpPeerEntry objects on the spine show that a password is not set.	14.1(2g)
CSCvo71335	ACI leaf switch running 13.2(2o) and 14.0(2c) is crashing and generating an NTP core file when collecting a techsupport file. This can be from apic generated or locally generated on the switch.	14.1(2g)
CSCvo84002	Link flaps, route add/delete. will cause l2mcast mem-leak. When the leak reached to around 1.5G, it will crash.	14.1(2g)
CSCvo86121	A link intermittently flaps on leaf switch fabric ports that are connected to a spine switch.	14.1(2g)
CSCvp05770	Recomissioning, decomissioning a spine in an environment, causes another SPINE to crash.	14.1(2g)
CSCvp19404	A vPC pair of leaf switches reboot due to an EPM HAP reset.	14.1(2g)
CSCvp47696	A port-client crash is seen on FX2 leaf switches when a large number of breakouts are configured.	14.1(2g)
CSCvp52500	In the broken flow the traffic from the source leaf was pointing to the destination TEP of MAC Spine-Proxy when it should be sent to the dest TEP of the destination leaf. The issue is that we don't have an XR learn for the destination MAC on the source leaf and therefore do not know	14.1(2g)

Bugs

Bug ID	Description	Fixed in
	<p>the dclass is for the destination EP. The problem seems to be in the placement of the rule that is getting installed. It looks like we are hitting the src_any_any rule before the BD rule. We see two rule priorities, 15 for vzAny redirect rule and 16 for any_dest_any with the dest pcTAG of the BD. Since priority 15 is preferred over 16 we hit the redirect rule. Subsequently, the traffic is hitting a PBR rule pointing to the MAC proxy.</p> <pre> 4184 0 49169 implicit enabled 2785280 permit any_dest_any(16) ... 4414 32809 0 default enabled 2785280 redir(destgrp-9) src_any_any(15) </pre>	
CSCvp65188	<p>Excessive GARP messages when Common Pervasive Gateway vMAC is configured.</p> <p>While with Common Pervasive Gateway vMAC, GARPs are expected to be sent out, a scale set up seems to be a problem as the packets seem to get multiplied.</p>	14.1(2g)
CSCvp65207	<p>Leaf crash with the following reset reason running 4.0(1h).</p> <pre> LO-L11# show system reset-reason ***** module reset reason (1) ***** 0) At 2019-04-30T08:46:05.182-05:00 Reason: reset-triggered-due-to-ha-policy-of-reset Service:nfm hap reset Version: 14.0(1h) </pre>	14.1(2g)
CSCvp66009	Unable to Login to Switches when using LDAP as Login domain.	14.1(2g)
CSCvp77034	The hardware abstraction layer (HAL) generates a core file when all of the non-fabric ports are converted into breakout ports.	14.1(2g)
CSCvp92121	<p>A vulnerability in the Link Layer Discovery Protocol (LLDP) subsystem of Cisco Nexus 9000 Series Application Centric Infrastructure (ACI) Mode Switch Software could allow an adjacent, unauthenticated attacker to cause a denial of service (DoS) condition or execute arbitrary code with root privileges.</p> <p>The vulnerability is due to improper input validation of certain type, length, value (TLV) fields of the LLDP frame header. An attacker could exploit this vulnerability by sending a crafted LLDP packet to the targeted device. A successful exploit may lead to a buffer overflow condition that could either cause a DoS condition or allow the attacker to execute arbitrary code with root privileges.</p> <p>Note: This vulnerability cannot be exploited by transit traffic through the device; the crafted packet must be targeted to a directly connected interface.</p> <p>Cisco has released software updates that address this vulnerability. There are no workarounds</p>	14.1(2m)

Bugs

Bug ID	Description	Fixed in
	<p>that address this vulnerability.</p> <p>This advisory is available at the following link:</p> <p>https://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20190731-nxos-bo</p>	
CSCvp92269	<p>Running a Qualys security scan results in the following message:</p> <p>CWE - 693 Protection Mechanism Failure -</p> <p>" HTTP Security Header Not Detected"</p>	14.1(2m)
CSCvq24680	<p>After a reboot of a leaf switch that is operating as a tier-1 leaf in a multi-tier ACI fabric, the leaf switch will be stuck in a reboot loop. The reason for the loop is because during boot, the sdkhal process crashes. You can see the crash when running " show cores" on the switch.</p>	14.1(2m)
CSCvq43058	<p>A spine switch fabric module or line card is reloaded unexpectedly due to a kernel panic. The stack trace includes the following statement:</p> <p>Kernel panic - not syncing: Out of memory: system-wide panic_on_oom is enabled</p>	14.1(2m)
CSCvq43477	<p>In the IPv6 options, for the source-link layer address field, IPv6 traffic is blackholed because the leaf switch sets the incorrect MAC address in the router advertisement's (RA's) source link-layer address. This happens only with RAs that are sent as a reply to the router solicitation from the host. Unsolicited RAs from the leaf switch have the correct MAC address of the leaf switch itself.</p> <p>The border leaf switch sends out unsolicited RA messages correctly with its link MAC address (0022.bdf8.19ff) in the source link-layer address field.</p>	14.1(2m)
CSCvq53300	<p>Tier 2 switches reboot every 10 minutes. Spine switches also reload once with the same IS-IS core file backtrace.</p>	14.1(2m)
CSCvq65546	<p>The PIM process crashes randomly.</p>	14.1(2m)
CSCvq40849	<p>Some 100 Gbps uplink ports between a spine switch and leaf switch do not come up.</p>	14.1(2o)
CSCvq42673	<ol style="list-style-type: none"> 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 causes a port channel bringup failure, or causes the port channel manager or eth_port_manager to crash on the switch. <p>This issue occurs when the vPC or port channel configuration is present even before the breakout is applied.</p>	14.1(2o)

Bugs

Bug ID	Description	Fixed in
CSCvr03623	When using a non-default SSM range, PIM and NGMVPN processes can be out of sync with respect to which border leaf switch is the stripe winner. This leads to the border leaf switch no longer sending PIM join upstream, and traffic within the SSM range is affected.	14.1(2o)
CSCvq06174	Multicast traffic is dropped when it is sent with COS 6 from a tenant.	14.1(2s)
CSCvq28541	Ports with power-over-Ethernet auto mode enabled go into the errdisable mode when a non-PoE-capable device is connected to it.	14.1(2s)
CSCvq64586	When the next join comes into the receiver VRF before the route expires in the source VRF, the receiver VRF tries to locate an extranet route in the source VRF. The receiver VRF either creates an extranet route if no route exists, or passes back a pointer. Because the previous route is still present in the Protocol Independent Multicast (PIM), there is no re-add for the new join. However, MRIB and NGMVPN updates rely on pim_add_route, which is called only when a new route needs to be created. Therefore, MRIB and NGMVPN will not be updated. No joins will be forwarded in the source VRF. Traffic will not flow.	14.1(2s)
CSCvq79143	When inter-VRF multicast is configured, chaining of the receiver VRFs is not allowed. Because of this restriction, if there is a source VRF A and receiver VRF B, the source VRF A cannot be a receiver VRF for the same group range. A configuration that chains the receiver VRFs, such as during bring up, can result in an MRIB crash.	14.1(2s)
CSCvr26434	There is a memory leak with the "show system internal epmc bltrace" CLI command.	14.1(2s)
CSCvr33980	The Oper VLAN is not programmed on a Power-over-Ethernet phone when using LLDP.	14.1(2s)
CSCvr35108	Fibre Channel links do not come up when connecting Cisco Nexus 9000 switches in ACI mode to MDS 9396T switches on ports above 63.	14.1(2s)
CSCvr62348	There is a slow memory leak during line card insertion or removal, FEX reloads, or transceiver removal. This issue occurs more frequently for copper ports.	14.1(2s)
CSCvn57265	With N9K-C9348GC-FXP and N9K-C9358GY-FXP top-of-rack switches, on ports where PoE is enabled in auto mode, there is a potential memory leak that 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.	14.1(2u)
CSCvr57536	An ACI node reloads due to the Machine Check Exception error, similar to the following output: <pre>[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</pre>	14.1(2u)
CSCvr73276	With N9K-C9348GC-FXP and N9K-C9358GY-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	14.1(2u)

Bugs

Bug ID	Description	Fixed in
	negotiations occur with a powered device (PD) or when EPG or VLAN information on a PoE interface policy is changed multiple times.	
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.	14.1(2u)
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.	14.1(2u)

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 14.1(2) releases in which the known behavior exists. A bug might also exist in releases other than the 14.1(2) releases.

Table 5 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.	14.1(2g) 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.	14.1(2g) and later
CSCup65586	The show interface command shows the tunnel's Rx/Tx counters as 0.	14.1(2g) 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.	14.1(2g) and later
CSCup92534	Continuous "threshold exceeded" messages are generated from the fabric.	14.1(2g) and later
CSCuq39829	Switch rescue user ("admin") can log into fabric switches even when TACACS is selected as the default login realm.	14.1(2g) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	14.1(2g) and later

Bugs

Bug ID	Description	Exists In
CSCuq77095	When the command <code>show ip ospf vrf <vrf_name></code> is run from bash on the border leaf, the checksum field in the output always shows a zero value.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	14.1(2g) and later
CSCur81822	The access-port operational status is always "trunk".	14.1(2g) 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.	14.1(2g) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) and later

Bugs

Bug ID	Description	Exists In
CSCus61748	<p>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.</p> <p>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.</p> <p>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.</p>	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) and later
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) and later

Bugs

Bug ID	Description	Exists In
CSCuy06749	Traffic is dropped between two isolated EPGs.	14.1(2g) and later
CSCuy22288	The <code>iping</code> command's replies get dropped by the QOS ingress policer.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) and later
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to "0", and so traffic will still flow.	14.1(2g) and later
CSCuy96912	The debounce timer is not supported on 25G links.	14.1(2g) 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.	14.1(2g) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	14.1(2g) 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.	14.1(2g) and later
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	14.1(2g) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) and later
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	14.1(2g) and later

Bugs

Bug ID	Description	Exists In
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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) 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.	14.1(2g) and later
CSCvp04772	During an upgrade on a dual-SUP system, the standby SUP may go into a failed state.	14.1(2g) 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 iVLAN packet. You can also embed CoS by enabling CoS preservation. For more information, see the *Cisco APIC and QoS* KB article, which you can find on the following URL:

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

- 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 Cisco Nexus 9508 ACI-mode switch supports 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.

Related Documentation

- 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.

IGMP Snooping Known Behaviors:

- Multicast router functionality is not supported when IGMP queries are received with VxLAN encapsulation.
- IGMP Querier election across multiple Endpoint Groups (EPGs) or Layer 2 outsiders (External Bridged Network) in a given bridge domain is not supported. Only one EPG or Layer 2 outside for a given bridge domain should be extended to multiple multicast routers if any.
- 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**.

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**.
- 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.

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 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 Cisco Systems, Inc. All rights reserved.