Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 13.2(5)

The Cisco NX-OS software for the Cisco Nexus 9000 series switches is a data center, purpose-built operating system designed with performance, resiliency, scalability, manageability, and programmability at its foundation. It provides a robust and comprehensive feature set that meets the requirements of virtualization and automation in data centers.

This Cisco NX-OS 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, 3.2(5), Release Notes*, 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 13.2(5) 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
May 16, 2022	In the Open Bugs section, added bug CSCwa47686.
August 10, 2021	In the Open Bugs section, added bug CSCvy30381.
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 Bugs section, added bug CSCvu07844.
January 22, 2021	In the Open Bugs 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.

Date	Description
May 12, 2020	13.2(5d): In the Resolved Bugs section, added bug CSCvo01105.
March 13, 2020	13.2(5d): In the Resolved Bugs section, added bug CSCvr98827.
September 27, 2019	In the Supported Hardware section, for the N9K-C9336C-FX2 switch, changed the port profile note to: The port profile feature does supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.
September 20, 2019	In the Usage Guidelines section, added the following bullet: ■ 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 11, 2019	In the Supported Hardware section, for the N9K-C9348GC-FXP, N9K-C93108TC-FX, and N9K-C93180YC-FX switches, added the following note: 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.
September 3, 2019	13.2(5d): In the Open Bugs section, added bugs CSCvp94661.
August 28, 2019	13.2(5d): In the Open Bugs section, added bugs CSCvq42673 and CSCvq43477.
August 14, 2019	13.2(5d): In the Open Bugs section, added bugs CSCvp92269, CSCvq43058, and CSCvq43477.
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.
June 5, 2019	In the Supported Hardware section, for the N9K-C9364C switch, removed the following erroneous sentence: The last 16 of the QSFP28 ports are colored green to indicate that they support wire-rate MACsec encryption.
May 15, 2019	13.2(5f): Release 13.2(5f) became available. Added the resolved bugs for this release.
April 9, 2019	13.2(5e): Release 13.2(5e) became available. Added the resolved bugs for this release.
March 20, 2019	13.2(5d): In the Open Bugs section, added bug CSCvo83991.
March 6, 2019	13.2(5d): Release 13.2(5d) became available.

Cisco Confidential

Contents

This document includes the following sections:

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

Cisco Confidential

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-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.

Cisco Confidential

Hardware Type	Product ID	Description
Spine switch	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 following PSUs are supported for the N9K-C9364C: ■ NXA-PAC-1200W-PE ■ NXA-PAC-1200W-PI ■ N9K-PUV-1200W ■ NXA-PDC-930W-PE ■ NXA-PDC-930W-PI 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. 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.
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 fan	N9K-C9300-FAN3	Port side intake fan
Spine switch fan	N9K-C9300-FAN3-B	Port side exhaust fan
Spine switch module	N9K-C9504-FM	Cisco Nexus 9504 fabric module supporting 40 Gigabit line cards

Cisco Confidential

Hardware Type	Product ID	Description
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	Cisco Nexus 9516 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 9500 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
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.

Cisco Confidential

Hardware Type	Product ID	Description
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. 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.
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.
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.
Leaf switch	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 either 24 40G ports or 12 100G ports. If 100G is connected the Port1, Port 2 will be HW disabled. Note: 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.
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

Cisco Confidential

Hardware Type	Product ID	Description
Leaf switch	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.
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.
Leaf switch	N9K-C9336C-FX2	Cisco Nexus C9336C-FX2 Top-of-rack (ToR) switch with 36 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: 1-Gigabit QSA is not supported on ports 1/1-6 and 1/33-36. The port profile feature does supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.

Cisco Confidential

Hardware Type	Product ID	Description
Leaf switch	N9K-C9348GC-FXP	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. This switch supports the following PSUs: ■ NXA-PAC-350W-PI ■ NXA-PAC-350W-PE ■ NXA-PAC-1100W-PI ■ NXA-PAC-1100W-PE 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. 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."
Leaf switch	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+.
Leaf switch	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+.
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

Cisco Confidential

Hardware Type	Product ID	Description
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
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
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
Leaf switch fan	NXA-FAN-30CFM-B	Red port side intake fan
Leaf switch fan	NXA-FAN-30CFM-F	Blue port side exhaust fan
Leaf switch fan	NXA-FAN-65CFM-PE	Blue port side exhaust fan
Leaf switch fan	NXA-SFAN-65CFM-PE	Blue port side exhaust fan
Leaf switch fan	NXA-FAN-65CFM-PI	Burgundy port side intake fan
Leaf switch fan	NXA-SFAN-65CFM-PI	Burgundy port side intake fan
Leaf switch power supply unit	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
Leaf switch power supply unit	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
Leaf switch power supply unit	NXA-PAC-1100W-PE2	1100W AC power supply, port side exhaust pluggable
Leaf switch power supply unit	NXA-PAC-1100W-PI2	1100W AC power supply, port side intake pluggable
Leaf switch power supply unit	N9K-PAC-650W	650W AC Power supply, port side intake pluggable

Cisco Confidential

Hardware Type	Product ID	Description
Leaf switch power supply unit	N9K-PAC-650W-B	650W AC Power supply, port side exhaust pluggable
Leaf switch power supply unit	NXA-PDC-1100W-PE	1100W AC power supply, port side exhaust pluggable
Leaf switch power supply unit	NXA-PDC-1100W-PI	1100W AC power supply, port side intake pluggable
Leaf switch power supply unit	NXA-PHV-1100W-PE	1100W HVAC/HVDC power supply, port-side exhaust
Leaf switch power supply unit	NXA-PHV-1100W-PI	1100W HVAC/HVDC power supply, port-side intake
Leaf switch power supply unit	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
Leaf switch power supply unit	N9K-PUV-3000W-B	3000W AC Power supply, port side exhaust pluggable
Leaf switch power supply unit	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.
Leaf switch power supply unit	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.
Leaf switch power supply unit	NXA-PAC-500W-PE	500W AC Power supply, port side exhaust pluggable

Cisco Confidential

Hardware Type	Product ID	Description
Leaf switch power supply unit	NXA-PAC-500W-PI	500W AC Power supply, port side intake pluggable
Leaf switch power supply unit	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.
Leaf switch power supply unit	UCSC-PSU-930WDC V01	Port side exhaust DC power supply compatible with all ToR leaf switches
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

There are no new hardware features in this release.

New Software Features

For new software features, see the *Cisco APIC 3.2(5) Release Notes* at the following location:

Cisco Confidential

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

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 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

- For the supported optics per device, see the Cisco Optics-to-Device Compatibility Matrix.
- 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).
- Link level flow control is not supported on ACI-mode switches.
- To connect the N2348UPQ to ACI leaf switches, the following options are available:
  - Directly connect the 40G FEX ports on the N2348UPQ to the 40G switch ports on the ACI leaf switches
  - Break out the 40G FEX ports on the N2348UPQ to 4x10G ports and connect to the 10G ports on all other ACI leaf switches

  Note: A fabric uplink port cannot be used as a FEX fabric port.
- To connect the APIC (the controller cluster) to the ACI fabric, it is required to have a 10G interface on the ACI leaf. You cannot connect the APIC directly to the 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.
- The following table provides MACsec and CloudSec compatibility information for specific hardware:

Cisco Confidential

Table 3 MACsec and CloudSec Support

Product ID	Hardware Type	MACsec Support	CloudSec Support
N9K-C93108TC-FX	Switch	Yes	No
N9K-C93180YC-FX	Switch	Yes	No
N9K-C93216TC-FX2	Switch	Yes	No
N9K-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 APIC release notes 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 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

- The Cisco APIC GUI incorrectly reports more memory used than is actually used. To calculate the appropriate amount of memory used, run the "show system internal kernel meminfo | egrep "MemT|MemA"" command on the desired switch. Divide MemAvailable by MemTotal, multiply that number by 100, then subtract that number from 100.
  - Example: 10680000 / 24499856 = 0.436 x 100 = 43.6% Free, 100% - 43.6% = 56.4% Used
- Leaf and spine switches from two different fabrics cannot be connected regardless of whether the links are administratively kept down.
- Only one instance of OSPF (or any multi-instance process using the managed object hierarchy for configurations) can have the write access to operate the database. Due to this, the operational database is limited to the default OSPF process alone and the multipodInternal instance does not store any operational data. To debug an OSPF instance ospf-multipodInternal, use the command in VSH prompt. Do not use ibash because some ibash commands depend on Operational data stored in the database.
- When you enable or disable Federal Information Processing Standards (FIPS) on a Cisco ACI fabric, you must reload each of the switches in the fabric for the change to take effect. The configured scale profile setting is lost when you issue the first reload after changing the FIPS configuration. The switch remains operational, but it uses the default port scale profile. This issue does not happen on subsequent reloads if the FIPS configuration has not changed.

  FIPS is supported on Cisco NX-OS release 13.2(5) or later. If you must downgrade the firmware from a release that supports FIPS to a release that does not support FIPS, you must first disable FIPS on the Cisco ACI fabric and reload all of the switches in the fabric.

Cisco Confidential

- 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
- 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.

Cisco Confidential

- 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 13.2(5) releases in which the bug exists. A bug might also exist in releases other than the 13.2(5) releases.

Table 4 Open Bugs in This Release

Bug ID	Description	Exists In
CSCvp79708	After a spine switch upgrade, there is traffic loss for inter-pod traffic.	13.2(5e) and later
CSCvq38040	There is a rare timing issue seen during F5 failover, which triggers a simultaneous local learn on one vPC TOR and a sync update from the peer. This sequence could end up causing an inconsistency in EPMC on one vPC peer where the endpoint ends up pointing to a bounce entry even though it was learned on the front panel.	13.2(5e) and later

Cisco Confidential

Bug ID	Description	Exists In
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"	13.2(5e) 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.	13.2(5e) 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.	13.2(5e) and later
CSCve06334	MAC and IP endpoints are not learned on the local vPC pair.	13.2(5d) and later
CSCvf09313	In the 12.2(2i) release, the BPDU filter only prevents interfaces from sending BPDUs, but does not prevent interfaces from receiving BPDUs.	13.2(5d) and later
CSCvg85886	When an ARP request is generated from one endpoint to another endpoint in an isolated EPG, an ARP glean request is generated for the first endpoint.	13.2(5d) and later
CSCvg95192	Endpoint information is missing in the spine switches.	13.2(5d) 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.	13.2(5d) and later
CSCvh14815	BGP EVPN has the tenant endpoint information, while COOP does not have the endpoint.	13.2(5d) 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.	13.2(5d) and later
CSCvj03533	When IPv6 packets are received, mab is triggered. But, only the MAC address endpoint is learned, not the IP address endpoint.	13.2(5d) 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.	13.2(5d) and later
CSCvj29908	Traffic gets dropped when a new TX SA is programmed after an old Rx SA is deleted on the peer and there are breakout ports in the link down state.	13.2(5d) and later
CSCvj50973	When the MTU settings for OSPF neighboring router interfaces do not match, the routers will be stuck in the Exstart/Exchange state. This behavior is expected. This bug is an enhancement to raise a fault to the APIC so that the routers' stuck state can be easily detected by the administrator.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvk34581	When viewing a congested interface, you do not see any drops in the output of the "show interface" command. If you type "vsh_lc" to drop into the linecard shell, and then view the platform counters for the given port, you can see Buffer Drops on output.	13.2(5d) and later
CSCvk48856	The port LED shows green when a few breakout ports lanes are down.	13.2(5d) and later
CSCvk73228	This is an enhancement to decode the binary logs offline directly from the techsupport.	13.2(5d) and later
CSCvk74561	Link down detection on the copper transceiver port takes around 1 second of time when its peer switch reloads. This issue is only with a copper transceiver.	13.2(5d) and later
CSCvm44326	COOP/IGMP crashes in the ZeroMQ library.	13.2(5d) and later
CSCvm75395	A route map is deployed even when the route profile is configured incorrectly. When upgrading to a release that includes the fixed for this defect, the incorrectly deployed route map is removed from the leaf switch, which may affect traffic that was using the route map.	13.2(5d) and later
CSCvn28108	A switch gets stuck in a bootloop with the following error raised on the console: [1041.090380] obfl_klm writing reset reason 58, LC insertion sequence failure => [Failures < MAX] : powercycle [1042.207780] write_mtd_flash_panic: successfully wrote 88 bytes at address 0xd68 to RR Iter: 0.	13.2(5d) and later
CSCvn92765	Excessive SSD writes are observed by ICMPv6, which can use up to 42GB per day.	13.2(5d) and later
CSCvo35006	Multiple N9K-X9736C-FX 40G line cards get stuck in the 'Inserted' state during a reload or reboot.	13.2(5d) 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.	13.2(5d) and later
CSCvo42234	There is high SSD utilization on the standby supervisor for a 95xx ACI spine switch.	13.2(5d) and later
CSCvo49717	After downgrading to the 13.2(5) release from a later release, an N9K-C9508-FM-E2 fabric module gets stuck in the 'Inserted' state.	13.2(5d) and later
CSCvo53218	10-20 second packet loss is observed when the designated forwarder leaf switch comes back online after a reload.	13.2(5d) and later
CSCvo74427	In a setup in which a leaf switch has 2 links to a spine switch, one link might flap a few times. The flapping seems to be triggered by a physical link flap (from the ethpm logs). After the link came up, the IS-IS update never reaches URIB. So, the leaf switch does not send any traffic on this link to the spine switch. The IS-IS database has the routes learned from this spine switch on both links.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvp00292	With contract-based L3Out QoS classification, the current implementation needs to use different filters for the QoS filter and traffic permission filters. This makes the configuration complicated, and additional TCAM cost is required.	13.2(5d) 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.	13.2(5d) and later
CSCvp50075	A leaf switch experiences an unexpected reload due to a HAP reset.	13.2(5d) and later
CSCvp59361	A kernel panic seen in some random scenarios.	13.2(5d) and later
CSCvp63213	While ACI switches are still initializing after an upgrade, TACACS requests are seen coming from the switch IP address, with the remote IP address set to 127.0.0.1 for the admin user.	13.2(5d) 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.	13.2(5d) and later
CSCvp91758	Fault F0449 gets raised and the ASIC vrm(5) status fails on the Cisco N9K-93108TC-EX or N9K-93180YC-EX switches.	13.2(5d) and later
CSCvp92269	Running a Qualys security scan results in the following message: CWE - 693 Protection Mechanism Failure - "HTTP Security Header Not Detected"	13.2(5d) and later
CSCvp94661	There is an EPM crash on a leaf switch that receives the Endpoint Announce packet with a malformed length field.	13.2(5d) and later
CSCvp98108	Traffic to be flooded in an EPG does not have fabricencap as the VNID in the IVXLAN header. Instead it has the primary VLAN that is configured for the path.	13.2(5d) and later
CSCvq20711	On a leaf switch, the "show interface description" command output in the ACI mode does not match the output of the "show int description" command output in the VSH mode.	13.2(5d) and later
CSCvq25729	Traffic is dropped when it is destined to a pervasive route and when the endpoint is not learned. This issue can be also seen on a border leaf switch when "disable remote EP learning" is set.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvq42673	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. This issue occurs when the vPC or port channel configuration is present even before the breakout is applied.	13.2(5d) and later
CSCvq43058	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	13.2(5d) and later
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.	13.2(5d) and later
CSCvq54991	Spine switches will not export flows in the absence of the controller IP address or if the controller and collectors have a different subnet.	13.2(5d) and later
CSCvq57935	A GOLF-enabled VRF instance is put into the Down state on the spine switches. This can be confirmed with the "show bgp process vrf <vrf-name>" command from the CLI of the spine switches. Behaviors that may indicate this issue include a loss of reachability to the endpoints in a GOLF-enabled VRF instance and missing routes on the leaf switch for the VRF instance in question.	13.2(5d) and later
CSCvq64803	A leaf switch crashes with the "Unknown" reset reason when the breakout ports configuration is re-applied. The reset reason for this switch is as follows: Image Version : 13.2(3o) Reset Reason (LCM): Unknown (0) at time Fri Jul 12 14:21:14 2019 Reset Reason (SW): Reset triggered due to HA policy of Reset (16) at time Fri Jul 12 14:17:40 2019 Service (Additional Info): Reset triggered due to HA policy of Reset	13.2(5d) and later
CSCvq67792	Posting the IPv6 interface configuration (including BFD enable) by using the API in an L3Out results in SVIs using the secondary IP address as the BFD source IP address. This causes the BFD session to fail.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
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'.	13.2(5d) and later
CSCvr08148	The N9K-C93180YC-EX leaf switch reboots for an unknown reason without any affected services: Last reset Reason: Unknown System version: <VERSION> Service:	13.2(5d) and later
CSCvr49904	Traffic with a UDP destination port of 8472 is dropped on ingress by the ACI fabric.	13.2(5d) and later
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.	13.2(5d) and later
CSCvr76058	A leaf switch crashes due to the following reason: Reason: reset-triggered-due-to-ha-policy-of-reset Service:device_test hap reset	13.2(5d) 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	13.2(5d) and later
CSCvr83337	A Cisco ACI leaf switch unexpectedly reloads and generates a core file.	13.2(5d) 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.	13.2(5d) and later
CSCvs10395	Leaf switch downlinks all go down at one time due to FabricTrack.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCvs18150	After a certain set of steps, it is observed that the deny-external-tag route-map used for transit routing loop prevention gets set back to the default tag 4294967295. Since routes arriving in Cisco ACI with this tag are denied from being installed in the routing table, if the VRF table that has the route-tag policy is providing transit for another VRF table in Cisco ACI (for instance and inside and outside vrf with a fw connecting them) and the non-transit VRF table has the default route-tag policy, routes from the non-transit VRF table would not be installed in the transit VRF table. This bug is also particularly impactful in scenarios where transit routing is being used and OSPF or EIGRP is used on a vPC border leaf switch pair. vPC border leaf switches peer with each other, so if member A gets a transit route from BGP, redistributes into OSPF, and then advertises to member B (since they are peers)...without a loop prevention mechanism, member B would install the route through OSPF since it has a better admin distance and would then advertise back into BGP. This VRF tag is set on redistribution of BGP > OSPF and then as a table map in OSPF that blocks routes with the tag from getting installed in the routing table. When hitting this bug, the route-map used for redistributing into OSPF still sets the tag to the correct value. However, the table map no longer matches the correct tag. Rather, it matches the default tag. As a result, member A (could be B) would install the route through OSPF pointing to B. It would then redistribute it back into BGP with the med set to 1. The rest of the fabric (including member B) would install the BGP route pointing to member A since its med is better than the original route's med.	13.2(5d) and later
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.	13.2(5d) 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:	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later
CSCvt08181	All routes to a particular spine switch are removed from uRIB on all leaf switches in the fabric.	13.2(5d) and later
CSCvt25383	The pervasive static route is missing on the spine node.	13.2(5d) and later
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.	13.2(5d) and later
CSCvt52620	There is a stale pervasive route after a DHCP relay label is deleted.	13.2(5d) and later
CSCvt82388	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. ARP/ICMPv6 adjacency updates can also contribute to many SSD writes.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later
CSCvu61024	Zoning-rules are not programmed in the hardware after reloading a switch.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later
CSCwa12763	External route import for a VRF instance fails on a leaf switch after removing a shared services contract between two EPGs.	13.2(5d) 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.	13.2(5d) and later
CSCwd29346	An ACI switch's console may continuously output messages similar to: svc_ifc_eventmg (*****) Ran 7911 msecs in last 7924 msecs	13.2(5d) and later
CSCvo83991	SNMP traps during relevant events are not being sent out of the ToR.	13.2(5d)

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.

Cisco Confidential

Table 5 Resolved Bugs in This Release

Bug ID	Description	Fixed In
CSCvo67347	There are stale entries present in the hardware forwarding tables. These entries normally do not have any impact. However, if a new VRF instance is created and is allocated the same software ID, and the VRF instance and tries to push the same prefix, then the hardware abstraction layer (HAL) will crash with an abort during the l3::L3RtgRouteTable::alloc process.	13.2(5f)
CSCvp32562	The link with active copper cables do not come up under certain conditions.	13.2(5f)
CSCvo83991	SNMP traps during relevant events are not being sent out of the ToR.	13.2(5e)
CSCvq17451	The "show int tunnel x" command shows a tunnel as down on the switch with a hw-prog failure, but in the hardware, the tunnel shows as up and operational.	13.2(5d)
CSCvk72867	EX and FX leaf switches crash with an SNMPD HAP reset.	13.2(5d)
CSCvm01561	DOM statistics in "show interface ethX/Y transceiver details" may differ for the same SFP between NX-OS standalone and ACI mode.	13.2(5d)
CSCvm52368	After creating a static L3Out in the same VRF instance as an OSPF L3Out, fault F0467 appears for the OSPF L3Out.	13.2(5d)
CSCvm71009	There is a dbgrelem memory leak every 30 seconds if PTP is enabled.	13.2(5d)
CSCvm72130	In a rare condition, a link with a 40/100G transceiver stays down. An example of such a transceiver is the QSFP-100G-PSM4.	13.2(5d)
CSCvm77150	After an upgrade, leaf switch front-panel ports are set into the admin-down (blacklist) state for 90 minutes.	13.2(5d)
CSCvm96594	Configuration import fails with an incorrect IP address duplicated across SVIs.	13.2(5d)
CSCvn01640	After adding a new L3Out in the same VRF instance (without a default route leak policy) to an existent tenant with the default route leak policy already configured, the default route is removed from the database. If the newly added L3Out includes a default route leak policy, the default route is kept in the database. If the most recently added L3Out has a default route leak policy configured, the default route leak policy remains in the database; otherwise, the default route leak policy is removed.	13.2(5d)
CSCvn18770	New IP addresses are unable to be learned in an ACI fabric.	13.2(5d)
CSCvn26568	A leaf switch does not boot into the OS. Because of this issue, there is no console output and the switch is not responsive.	13.2(5d)
CSCvn38855	When the First Hop Router has already a (*,G) entry for the group, a PIM data register does not get sent to the rendezvous point. This introduces a delay between 5 and 60 seconds until the First Hop Router sends the NULL register to the rendezvous point.	13.2(5d)
CSCvn41955	When Rogue Endpoint Detection is enabled, you may see endpoints that are attached by way of a vPC get marked as being Rogue on remote switches.	13.2(5d)

Cisco Confidential

Bug ID	Description	Fixed In
CSCvn44959	Traffic to GOLF routes are being dropped due to an inconsistency between the Forwarding Information Base (FIB) table and hardware.	13.2(5d)
CSCvn45911	Some packet loss can be seen when an OSPF neighbor goes into the restart state.	13.2(5d)
CSCvn55841	Traffic is dropped due to CRC errors when pinging from an isolated intra-EPG to another EPG in the same BD and the frame size is 5300 or greater.	13.2(5d)
CSCvn69340	The BFD session does not get instantiated under some circumstances in one of the VPC legs for static routes.	13.2(5d)
CSCvn69340	The BFD session does not get instantiated under some circumstances in one of the vPC legs for static routes.	13.2(5d)
CSCvn71211	A fabric port flaps without any related configuration or physical changes.	13.2(5d)
CSCvn72145	An ACI leaf switch reboots with the following reason: *************** module reset reason (1) ************* 0) At 2018-12-14T14:25:15.050+00:00 Reason: reset-triggered-due-to-ha-policy-of-reset Service:epm hap reset Version: 13.1(2m)	13.2(5d)
CSCvn77885	When spine switches are connected by multiple links to multiple external IPN switches, reloading of one of the IPN switches sometimes results in multi-destination traffic loss across PODs. This is the result of the IS-IS and multicast code getting out of sync with respect to the external interfaces. This issue can be detected by dumping the GIPO object in the SDK.	13.2(5d)
CSCvn80904	A leaf switch crashes due to a sysmgr crash.	13.2(5d)
CSCvn98080	The span_mgr sometimes dumps a core.	13.2(5d)
CSCvn99354	When an endpoint moves from POD2 to POD1, the bounce to the proxy entry points to the MAC address proxy in the POD2 leaf switches instead of pointing to the IPv4 proxy in the spine switch.	13.2(5d)
CSCvo01105	The SSD lifetime on a switch may be reduced significantly due to constant ARP/ICMPv6 adjacency updates. Rather than process all adjacency updates in real-time, this enhancement is raised to report changes every minute to reduce early SSD wear out.	13.2(5d)
CSCvo12091	The "show platform internal hal tunnel rtep apd" command output does not include "rwencapidx" when run on N9K-C9336C-FX2, N9K-C93240YC-FX2, or N9K-C93240YC-FX2Z TOR switches.	13.2(5d)
CSCvo16431	When negotiation is enabled and a remote device is a Red Hat server version 7.5, the bond0 log file of an N9K-C93108TC-EX switch shows a link failure count of 1. The count normally is 0. If the remote device is instead an N9KN9K-C93108TC-FX switch in standalone mode, the physical link flaps once, when it should not flap.	13.2(5d)
CSCvo17792	The port-client periodically crashes on N9K-93180LC-EX switches.	13.2(5d)

Cisco Confidential

Bug ID	Description	Fixed In
CSCvo31537	A static route configured with BFD may result in a null route when the route is installed after a link up even after connect next hops have been resolved.	13.2(5d)

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 13.2(5) releases in which the known behavior exists. A bug might also exist in releases other than the 13.2(5) releases.

Table 6 Known Behaviors in the 13.2(5f) 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.	13.2(5d) 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.	13.2(5d) and later
CSCup65586	The show interface command shows the tunnel's Rx/Tx counters as 0.	13.2(5d) 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.	13.2(5d) and later
CSCup92534	Continuous "threshold exceeded" messages are generated from the fabric.	13.2(5d) and later
CSCuq39829	Switch rescue user ("admin") can log into fabric switches even when TACACS is selected as the default login realm.	13.2(5d) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	13.2(5d) 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.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	13.2(5d) and later
CSCur81822	The access-port operational status is always "trunk".	13.2(5d) 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.	13.2(5d) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.	13.2(5d) 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.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	13.2(5d) and later
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to "0", and so traffic will still flow.	13.2(5d) and later
CSCuy96912	The debounce timer is not supported on 25G links.	13.2(5d) 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.	13.2(5d) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	13.2(5d) 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.	13.2(5d) and later

Cisco Confidential

Bug ID	Description	Exists In
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	13.2(5d) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) 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.	13.2(5d) and later
CSCvo22890	There is intermittent packet loss for some flows through FX2 leaf switches when the no-drop class is enabled.	13.2(5d) and later
CSCvo27881	Ping stops working between a VM behind a non-FIE EPG and a VM behind an FIE-enabled EPG.	13.2(5d) 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.	13.2(5d) 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, which you can find on the following URL:

Cisco Confidential

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 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.

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 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 switch, follow these configuration steps to recover:
  - Set "unknown multicast flooding" to "Flood" for all bridge domains mapped to a leaf switch.
  - Set "unknown multicast flooding" to "Optimized Flood" on needed bridge domains.
- 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.

Cisco Confidential

- 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.
- When IGMP snoop is enabled on UCSM and VLAN encapsulation endpoints subscribe to any multicast groups, the FI will rewrite the source MAC address of the IGMP report with its own MAC address. The TOR switch will learn from the IGMP report, which causes the endpoints' IP addresses to move under the FI uplink MAC instead of actual MAC, which will cause traffic loss for those endpoints. We recommend that you disable IGMP snoop under UCSM or deploy the endpoints behind vxLAN mode to avoid unicast traffic losses.

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.

© 2019-2024 Cisco Systems, Inc. All rights reserved.