Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 13.2(41)

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 Release Notes, Release 3.2(41)*, 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 Nexus 9000 ACI-Mode Switches Release Notes, Release 13.2(41)*:

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

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

Table 1 shows the online change history for this document.

Table 1. Online History Change

Date	Description
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.
July 12, 2019	13.2(41d): Release 13.2(41d) became available. The information in this document is the same as for the 13.2(4) release, with the addition of resolved bug CSCvo80686.
August 28, 2019	13.2(41d): In the Open Bugs section, added bugs CSCvq42673 and CSCvq43477.

Date	Description
September 3, 2019	13.2(41d): In the Open Bugs section, added bugs CSCvp94661.
September 20, 2019	In the Usage Guidelines section, added the following bullet: ■ A 25G link that is using the lEEE-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	In the Supported Hardware section, for the N9K-C9336C-FX2 switch, changed the port profile note to: The port profile feature supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.
March 13, 2020	13.2(41d): In the Resolved Bugs section, added bug CSCvr98827.

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-M6PQ	6-port
Pluggable module (GEM)	N9K-M6PQ-E	6-port, 40 Gigabit Ethernet expansion module
Pluggable module (GEM)	N9K-M12PQ	12-port or 8-port
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.
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

Cisco Confidential

Hardware Type	Product ID	Description
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 fan	N9K-C9300-FAN3	Port side intake fan
Spine switch fan	N9K-C9300-FAN3-B	Port side exhaust fan
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 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

Cisco Confidential

Hardware Type	Product ID	Description
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-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-B	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
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
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

Cisco Confidential

Hardware Type	Product ID	Description
Top-of-rack (ToR) 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.
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
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	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.
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.

Cisco Confidential

Hardware Type	Product ID	Description
Top-of-rack (ToR) 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 supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.
Top-of-rack (ToR) 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
Top-of-rack (ToR) 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+.
Top-of-rack (ToR) 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+.
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

Cisco Confidential

Hardware Type	Product ID	Description
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 power supply unit	N9K-PAC-650W-B	650W AC Power supply, port side exhaust pluggable
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-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
Top-of-rack (ToR) 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
Top-of-rack (ToR) 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
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 Note: 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 Note: This power supply is supported only by the Cisco Nexus 93120TX, 93128TX, and 9336PQ ACI-mode switches.

Cisco Confidential

Hardware Type	Product ID	Description
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 Note: This power supply is supported only by the Cisco Nexus 9348GC-FXP ACI-mode switch.
Top-of-rack (ToR) leaf switch power supply unit	UCS-PSU-6332-DC	930W DC power supply, reversed airflow (port side exhaust)
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 fan	NXA-FAN-30CFM-F	Port side intake fan
Top-of-rack (ToR) leaf switch fan	NXA-FAN-30CFM-B	Port side exhaust fan

Supported FEX Models

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

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

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

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

New and Changed Information

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

- New Hardware Features
- New Software Features

New Hardware Features

The following hardware features are now available:

Cisco Confidential

- The Cisco Nexus 9336C-FX2 switch (N9K-C9336C-FX2) now supports support uplink on ports 31 through 36. Use ports 35 and 36 for the minimum uplink.

New Software Features

For new software features, see the *Cisco Application Policy Infrastructure Controller Release Notes, Release 3.2(41)* 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

This release supports the following Cisco APIC servers:

Product ID	Description
APIC-L1	Cisco APIC with large CPU, hard drive, and memory configurations (more than 1000 edge ports)
APIC-L2	Cisco APIC with large CPU, hard drive, and memory configurations (more than 1000 edge ports)
APIC-M1	Cisco APIC with medium-size CPU, hard drive, and memory configurations (up to 1000 edge ports)
APIC-M2	Cisco APIC with medium-size CPU, hard drive, and memory configurations (up to 1000 edge ports)

The following list includes additional compatibility information:

- 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:

Cisco Confidential

- Directly connect the 40G FEX ports on the N2348UPQ to the 40G switch ports on the ACI leaf switches
- Break out the 40G FEX ports on the N2348UPQ to 4x10G ports and connect to the 10G ports on all other ACI leaf switches

Note: A fabric uplink port cannot be used as a FEX fabric port.

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

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.

Cisco Confidential

- A 10G copper SFP module on the peer is not supported.
- CloudSec only works with spine switches in Cisco ACI and only works between sites managed by Cisco ACI Multi-Site.
- For CloudSec to work properly, all of the spine switch links that participate in Cisco ACI Multi-Site must have MACsec/CloudSec support.

Usage Guidelines

- The current list of protocols that are allowed (and cannot be blocked through contracts) include the following. Some of the protocols have SrcPort/DstPort distinction.

  Note: See the 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
- 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

Cisco Confidential

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(41) 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.
- The dual rate BiDirectional (BiDi) transceiver QSFP-40/100-SRBD requires approximately 5 minutes for the link to come up after auto-negotiating the speed on both the local and remote end.
  - If both ends support the 40/100 combination, the link comes up quickly as 100G.
  - If one end is 40G and other end supports 40/100, then the link takes longer to negotiate to 40G.
- 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

Cisco Confidential

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 13.2(41) releases in which the bug exists.

Cisco Confidential

Table 4 Open Bugs in This Release

Bug ID	Description	Exists In
CSCvg85886	When an ARP request is generated from one endpoint to another endpoint in an isolated EPG, an ARP glean request is generated for the first endpoint.	13.2(41d) and later
CSCvg95192	Endpoint information is missing in the spine switches.	13.2(41d) 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(41d) 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(41d) 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(41d) 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(41d) 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(41d) and later
CSCvk48856	The port LED shows green when a few breakout ports lanes are down.	13.2(41d) and later
CSCvk73228	This is an enhancement to decode the binary logs offline directly from the techsupport.	13.2(41d) 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(41d) 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(41d) and later
CSCvn69340	The BFD session does not get instantiated under some circumstances in one of the VPC legs for static routes.	13.2(41d) 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(41d) 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(41d) 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(41d) and later
CSCvq98750	In Cisco ACI when using MAC pinning with a vPC, prior to reloading when you run the 'show vpc brief' command on the CLI, the command shows that the vPC is passing consistency checks. However, after reloading the leaf switch, the vPC then properly displays the consistency check as 'Not Applicable'.	13.2(41d) 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(41d) 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(41d) and later
CSCvs10395	Leaf switch downlinks all go down at one time due to FabricTrack.	13.2(41d) and later

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
CSCvi05870	On reload, vPC interfaces are temporarily brought online and then brought down. During this transition, traffic loss may be seen for 3-10 seconds.	13.2(41d)
CSCvi22143	Multi-destination traffic is not sent out the leaf switch uplinks or downlinks to other devices. This can result in ARP resolution problems or issues with the spanning tree.	13.2(41d)
CSCvj58567	When used for IPN connectivity, the fixed 10 Gb ports Eth1/65 and Eth1/66 cannot resolve ARP.	13.2(41d)
CSCvj99258	The kernel crashes with the following entry (dc3_sensor) in the dmesg-mtdoops or kernel crash logs: <0>[62761.120792] BUG: soft lockup - CPU#3 stuck for 22s! [dc3_sensor:7911]	13.2(41d)
CSCvk15151	There is a 1G port flap when 1G SFP modules are inserted into other ports on same MAC address.	13.2(41d)
CSCvk33286	The device reloads with the reset reason of "kernel-panic." This applies to kernel panics with and without PCIE issues recorded on dmesg-mtdoops.	13.2(41d)
CSCvk36215	1G QSA links are not coming up in the switches.	13.2(41d)
CSCvk41926	A leaf switch reloads and shows the following reason: reset-triggered-due-to-ha-policy-of-reset There is also an LLDP core file generated after the crash.	13.2(41d)
CSCvk50550	Cisco NX-OS shows fault F2821 under normal operations with no impact to traffic.	13.2(41d)
CSCvk75726	The bridge domain gateway IP address is not present in COOP.	13.2(41d)
CSCvm05674	The following fault is raised on the Cisco APIC for database corruption on the SQLLite database: Database got corrupted.[SQLiteCorruptionException]	13.2(41d)
CSCvm06425	The fault F3045 "IGMP Snooping fabric learning is disabled" is erroneously raised.	13.2(41d)
CSCvm09903	Following the upgrade of a Cisco ACI leaf switch, uplink port 1/51 may stay down for an extended period of time. Eventually the link will recover on its own.	13.2(41d)
CSCvm15457	The product Cisco Nexus 9000 Series Fabric Switches - ACI mode includes a version of the Linux kernel that is affected by the IP Fragment Reassembly Denial of Service Vulnerability identified by the following Common Vulnerability and Exposures (CVE) ID: CVE-2018-5391 Cisco has confirmed that this product is impacted.	13.2(41d)
CSCvm16440	PPI-COPP stats are not incrementing and traffic is not rate limited. Instead, traffic is hitting on the default COPP policy.	13.2(41d)

Cisco Confidential

Bug ID	Description	Fixed in
CSCvm26708	A 100/40G bidirectional platform sometimes latches onto a 40G link even when the peer has a 100/40G bidirectional platform between the leaf switch and spine switch.	13.2(41d)
CSCvm37645	Link detection takes up to 17 seconds when a ToR switch goes down.	13.2(41d)
CSCvm40454	After an overnight continuous interface flap test, the link fails to come up or a HAL process core might be seen on the switch (leaf or spine).	13.2(41d)
CSCvm46784	A link takes a longer time to come up when connected with 25G transceivers.	13.2(41d)
CSCvm47525	Packets are transmitted tagged with vlan-0. This can result in connectivity issues (for example, pings will time out) for connecting devices that do no support or expect tagged frames.	13.2(41d)
CSCvm48676	A Cisco N9K-C93180LC-EX switch reboots with the following reason: reset-triggered-due-to-ha-policy-of-reset	13.2(41d)
CSCvm48981	A Bidirectional Forwarding Detection (BFD) session stays in the DOWN state.	13.2(41d)
CSCvm52368	After creating a static L3Out in the same VRF instance as L3Out, the fault F0467 appears for OSPF L3Outs.	13.2(41d)
CSCvm55571	A contract with a subject that allows ssh and has reverse filter port enabled will not work after upgrading to the 13.2.3i release. The ability to use ssh to connect to the switch stops working.	13.2(41d)
CSCvm65702	If initially only a few leaf switches were configured with the Netflow policy and then the leaf switch profile was changed to include all leaf switches, Netflow will fail on some of the leaf switches.	13.2(41d)
CSCvm70324	When a peer is configured with private remove-as option in parallel, the private remove-as is not going to work in backend.	13.2(41d)
CSCvm75089	When attempting to run "show processes memory", the full output does not display, but instead the command fails early with the following error: dn "sys/procsys/proc-0/CDprocProcMem5min" could not be found Error executing command, check logs for details	13.2(41d)
CSCvm75296	Enabling the TrustSec feature on BGP causes traffic disruption.	13.2(41d)
CSCvm77485	IGMP messages are no longer flooded by a Cisco ACI leaf switch in a bridge domain with IGMP snooping disabled.	13.2(41d)
CSCvm82499	Traffic destined to a PBR anycast MAC address is dropped on the spine switch.	13.2(41d)
CSCvm88651	Traffic potentially drops with the reason of "VLAN_XLATE_MISS" or "UC_TENANT_MYTEP_BRIDGE_MISS" on changing the port scope policy from GLOBAL to LOCAL and back to GLOBAL, or changing LOCAL to GLOBAL.	13.2(41d)
CSCvn00808	When enabling "Limit IP Learning to Subnet" (limitIpLearnToSubnets) on a bridge domain, all endpoints including IP addresses within the subnet are flushed from the leaf and affect	13.2(41d)

Cisco Confidential

Bug ID	Description	Fixed in
	connectivity until the endpoints are relearned. The expectation is that only endpoints outside of the configured subnet are flushed.	
CSCvn17131	The Spanning Tree Protocol (STP) PortFast Bridge Protocol Data Units (BPDUs) are not transited across the fabric, which can lead to Layer 2 instability. You can check for tx, but not rx BPDUs, with the following command: show spanning internal info all | egrep 'bpdu_in|bpdu_out|Eth1'	13.2(41d)
CSCvo80686	A vulnerability in the SSH key management for the Cisco Nexus 9000 Series Application Centric Infrastructure (ACI) Mode Switch Software could allow an unauthenticated, remote attacker to connect to the affected system with the privileges of the root user. The vulnerability is due to the presence of a default SSH key pair that is present in all devices. An attacker could exploit this vulnerability by opening an SSH connection via IPv6 to a targeted device using the extracted key materials. An exploit could allow the attacker to access the system with the privileges of the root user. This vulnerability is only exploitable over IPv6; IPv4 is not vulnerable. Cisco has released software updates that address this vulnerability. This advisory is available at the following link: https://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20190501-nexus9k-sshkey	13.2(41d)

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(41) releases in which the known behavior exists.

Table 6 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.	13.2(41d) 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(41d) and later
CSCup65586	The show interface command shows the tunnel's Rx/Tx counters as 0.	13.2(41d) and later

Cisco Confidential

Bug ID	Description	Exists in
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(41d) and later
CSCup92534	Continuous "threshold exceeded" messages are generated from the fabric.	13.2(41d) and later
CSCuq39829	Switch rescue user ("admin") can log into fabric switches even when TACACS is selected as the default login realm.	13.2(41d) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	13.2(41d) 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(41d) 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.	13.2(41d) 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(41d) 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(41d) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	13.2(41d) and later
CSCur81822	The access-port operational status is always "trunk".	13.2(41d) 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(41d) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	13.2(41d) 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(41d) and later

Cisco Confidential

Bug ID	Description	Exists in
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(41d) and later
CSCus61748	If the TOR 1RU system is configured with the RED fan (the reverse airflow), the air will flow from front to back. The temperature sensor in the back will be defined as an inlet temperature sensor, and the temperature sensor in the front will be defined as an outlet temperature sensor. If the TOR 1RU system is configured with the BLUE fan (normal airflow), the air will flow from back to front. The temperature sensor in the front will be defined as an inlet temperature sensor, and the temperature sensor in the back will be defined as outlet temperature sensor. From the airflow perspective, the inlet sensor reading should always be less than the outlet sensor reading. However, in the TOR 1RU family, the front panel temperature sensor has some inaccurate readings due to the front panel utilization and configuration, which causes the inlet temperature sensor reading to be very close, equal, or even greater than the outlet temperature reading.	13.2(41d) 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(41d) 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(41d) 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(41d) 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(41d) 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(41d) 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(41d) and later
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.	13.2(41d) 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(41d) 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.	13.2(41d) 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(41d) and later

Cisco Confidential

Bug ID	Description	Exists in
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(41d) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	13.2(41d) and later
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	13.2(41d) 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(41d) 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(41d) 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(41d) 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(41d) and later
CSCuy96912	The debounce timer is not supported on 25G links.	13.2(41d) 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(41d) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	13.2(41d) 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(41d) and later
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	13.2(41d) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	13.2(41d) 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(41d) 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(41d) and later

Cisco Confidential

Bug ID	Description	Exists in
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	13.2(41d) 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(41d) 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(41d) 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(41d) and later
CSCvo22890	There is intermittent packet loss for some flows through FX2 leaf switches when the no-drop class is enabled.	13.2(41d) 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 multipod QoS policy in the ACI for the multipod that translates queuing class information of the packet into the DSCP value in the outer header of the iVXLAN packet.
- 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.

Cisco Confidential

- 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.

  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 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.

© 2019-2021 Cisco Systems, Inc. All rights reserved.