Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 12.1(4)

This document describes the features, bugs, and limitations for Cisco NX-OS software that runs on Cisco Nexus 9000 Series Application Centric Infrastructure (ACI) switches. Use this document in combination with the *Cisco Application Policy Infrastructure Controller Release Notes, Release 2.1(4)*, 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 12.1(4) 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
January 19, 2021	In the Known Behaviors section, changed the following sentence: The Cisco Nexus 9508 ACI-mode switch supports warm (stateless) standby where the state is not synched between the active and the standby supervisor modules. To: The modular chassis Cisco ACI spine nodes, such as the Cisco Nexus 9508, support warm (stateless) standby where the state is not synched between the active and the standby supervisor modules.
March 14, 2018	12.1(4a): Release 12.1(4a) became available.
April 24, 2018	12.1(4a): In the Open Bugs section, added bug CSCvi57920.
January 8, 2019	In the Supported Hardware section, added the Cisco N9K-C9336PQ and N9K-X9736PQ switches.
January 28, 2019	12.1(4a): In the Open Bugs section, added bug CSCvi76161.
January 30, 2019	12.1(4a): In the Open Bugs section, added bug CSCvn69340.

Date	Description
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.
August 14, 2019	12.1(4a): In the Open Bugs section, added bug CSCvp92269.
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.
March 13, 2020	12.1(4a): In the Resolved Bugs section, added bug CSCvr98827.

Contents

This document includes the following sections:

Cisco Nexus 9000 Series ACI-Mode

Cisco NX-OS Software for the Cisco Nexus 9000 Series 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

Cisco NX-OS Release 12.1 works only on Cisco Nexus 9000 Series switches in ACI Mode.

See Table 2 for a list of modules that are supported on Cisco Nexus 9000 Series switches in ACI Mode.

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
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

Hardware Type	Product ID	Description
Spine switch	N9K-C9516	Cisco Nexus 9516 switch with 16 line card slots Note: This switch supports up to 10 line cards.
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
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-X9732C-EX	Cisco Nexus 9500 32-port, 40/100 Gigabit Ethernet QSFP28 aggregation module
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
Top-of-rack (ToR) leaf switch	N9K-C93108TC-EX	Cisco Nexus 9300 with 48-port 1/10 Gigabit-T and 6-port 100 Gigabit Ethernet QSFP28 switch
Top-of-rack (ToR) leaf switch	N9K-C93120TX	Cisco Nexus 9300 with 96-port 1/10 Gigabit-T and 6-port 40 Gigabit Ethernet QSFP switch
Top-of-rack (ToR) leaf switch	N9K-C93128TX	Cisco Nexus 9300 96-port, 1-/10-Gbps BASE-T and 6-port or 8-port, 40 Gigabit Ethernet QSFP switch

Hardware Type	Product ID	Description
Top-of-rack (ToR) leaf switch	N9K-C93180YC-EX	Cisco Nexus 9300 Fixed with 48-port 10/25 Gigabit and 6-port 40/100 Gigabit QSFP28
Top-of-rack (ToR) leaf switch	N9K-C9332PQ	Cisco Nexus 9332PQ 32-port 40 Gigabit Ethernet QSFP+ Top-of-rack (ToR) Layer 3 switch
Top-of-rack (ToR) leaf switch	N9K-C9372PX	Cisco Nexus 9372PX 48-port, 10 Gigabit Ethernet SFP+ and 6-port 40 Gigabit Ethernet QSFP+ Top-of-rack (ToR) Layer 3 switch 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 48-port, 10 Gigabit Ethernet SFP+ and 6-port 40 Gigabit Ethernet QSFP+ Top-of-rack (ToR) Layer 3 switch 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 48-port, 1/10 Gbps Base-T and 6-port, 40 Gigabit Ethernet QSFP Top-of-rack (ToR) Layer 3 switch
Top-of-rack (ToR) leaf switch	N9K-C9372TX-E	Cisco Nexus 9372TX-E 48-port 1/10 Gbps Base-T and 6-port 40 Gbps Ethernet QSFP+ Top-of-rack (ToR) Layer 3 switch
Top-of-rack (ToR) leaf switch	N9K-C9396PX	Cisco Nexus 9300 48-port, 1/10 Gigabit Ethernet SFP+ and 6-port or 12-port, 40 Gigabit Ethernet QSFP switch
Top-of-rack (ToR) leaf switch	N9K-C9396TX	Cisco Nexus 9300 48-port, 1/10 Gbps Base-T and 6-port or 12-port, 40 Gigabit Ethernet QSFP switch
Top-of-rack (ToR) leaf switch power supply unit	N9K-PAC-650W-B	650W AC Power supply, port side exhaust pluggable

Hardware Type	Product ID	Description
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.

Hardware Type	Product ID	Description
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 exhaust fan
Top-of-rack (ToR) leaf switch fan	NXA-FAN-30CFM-B	Port side intake fan

Supported FEX Models

Table 3 lists the FEX models that the Cisco Nexus 9000 Series ACI Mode switches support. 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

Table 3. Supported FEX Models.

Product ID	Description
N2K-B22DELL-P	B22 FEX for Dell
N2K-B22HP-P	B22 FEX for HP
N2K-B22IBM-P	B22 FEX for IBM
N2K-C2248PQ-10GE	Cisco Nexus 2248PQ 10GE Fabric Extender, 2PS, 4 Fan Module, 48x1/10GE (req SFP/SFP+) + 4x40G QSFP+(req QSFP+), choice of airflow and power supply
N2K-C2248TP-1GE	Cisco Nexus 2248TP Series 1GE Fabric Extender, 2 AC PS, 1 Fan Module (Standard Airflow/port side exhaust), 48x100/1000Base-T + 4x10GE (req SFP+), same as N2K-C2248TP
N2K-C2248TP-E-1GE	Cisco Nexus 2248TP-E Series 1GE Fabric Extender, 2PS, 1 Fan Module, 48x100/1000Base-T + 4x10GE (req SFP+), 32MB buffer, choice of airflow and power supply
N2K-C2332TQ	Cisco Nexus 2332TQ 10G BASE T Fabric Extender, 2PS, 3 Fan Module, 48x100M/1/10GE + 4x40G QSFP+(req QSFP+), choice of airflow and power supply

Product ID	Description
N2K-C2348TQ	Cisco Nexus 2348TQ 10G BASE T Fabric Extender, 2PS, 3 Fan Module, 48x100M/1/10GE + 6x40G QSFP+(req QSFP+), choice of airflow and power supply
N2K-C2348UPQ	48 100M/1/10 Gigabit Ethernet and Unified Port host interfaces (SFP+) and up to 6 QSFP+ 10/40 Gigabit Ethernet fabric interfaces
N2K-C2232PP-10GE	Cisco Nexus 2232PP Series 10GE Fabric Extender, 2 AC PS, 1 Fan Module (Standard Airflow/port side exhaust), 32x1/10GE (req SFP/SFP+) + 8x10GE (req SFP+), same as N2K-C2232PP
N2K-C2232TM-E-10GE	Cisco Nexus 2232TM-E Series 10GBASE-T Fabric Extender, 2PS, 1 Fan Module, 32x1/10GBase-T + 8x10GE Module (req SFP+), choice of airflow and power supply

New and Changed Information

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

- New Hardware Features
- New Software Features

New Hardware Features

This release supports no new hardware features.

New Software Features

For new software features, see the *Cisco APIC 2.1(4) Release Notes* 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 hardware and software listed on the ACI Ecosystem Compatibility List and the Cisco AVS, Release 5.2(1)SV3(2.1).
- Link level flow control is not supported on ACI-mode switches.
- The breakout of 40G ports to 4x10G on the N9332PQ switch is not supported in ACI-Mode.
- 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.

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.
- 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

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 12.1(4) releases in which the bug exists. A bug might also exist in releases other than the 12.1(4) releases.

Table 4 Open Bugs in This Release

Bug ID	Description	Exists In
CSCun35596	FEX logs are missing in the output of the show fex detail command.	12.1(4a) and later
CSCun96495	The events and faults for interfaces are not updated under Ports in the GUI.	12.1(4a) and later
CSCup05629	The output of some CLI commands display very slowly. This usually occurs in a scaled environment when the switches are heavily loaded with the configuration.	12.1(4a) and later
CSCup86130	Because ibash is implemented on top of bash, when using ibash for the CLI, the bash behavior is inherited. For example, the sh mod command works in traditional Cisco switches. But when executed on N9K switches in ibash, because bash interprets sh differently, sh mod will not work. Similarly, if there is a clash in the next available options, the TAB key must be pressed twice to get the options rather than once as in other Cisco switches. In short, the CLI infra for ibash is not exactly the same as the CLI infra for the traditional Cisco switches because N9K ibash is built on top of bash.	12.1(4a) and later
CSCur32247	FEX-related diagnostic results are missing.	12.1(4a) and later

Bug ID	Description	Exists In
CSCuz82233	The server virtual Fibre Channel interface state changes to “port reinit limit reached” when an NP link is shut down.	12.1(4a) and later
CSCva27324	The virtual Fibre Channel (VFC) NP port enters the flogi-fail-retry state followed by the down state if there is a user-configured mismatch of the VSAN and VLAN mapping between the leaf switch and the Fibre Channel Forwarder (FCF). To recover the VFC NP port, delete the wrong VLAN-VSAN mapping on the leaf switch, delete and recreate the VLAN pool/encapsulation block if needed, and then create the correct VLAN-VSAN mapping.	12.1(4a) and later
CSCvb12858	With passive QSA, a GLC-SX-MMD transceiver is not detected by N9K-93108TC-EX and N9K-93180YC-EX switches.	12.1(4a) and later
CSCvb42735	A port is put into the "learn disable" state when the MAC limit is reached. When an existing endpoint on a learn-disabled port is updated with a new IP address, for example, the endpoint might get deleted erroneously. As a result, the number of dynamic endpoints on a learn-disabled port might be less than the MAC limit.	12.1(4a) and later
CSCvb54216	The permit log and glean packets share the same policer, which can cause direct BGPs to take a while to establish when one of the VPC peers is reloaded. This can occur when permit log is enabled and traffic is forwarding. There is no traffic loss, as BGP is established with other VPC peers, and traffic continues through the other peers.	12.1(4a) and later
CSCvi57920	A UCS 1225 vNIC goes down after changing the peer Cisco ACI leaf node name.	12.1(4a) and later
CSCvi76161	A version mismatch between Cisco ACI leaf switches causes the EPM process to crash and a HAP reset to occur.	12.1(4a) and later
CSCvn69340	The BFD session does not get instantiated under some circumstances in one of the VPC legs for static routes.	12.1(4a) and later
CSCvp50075	A leaf switch experiences an unexpected reload due to a HAP reset.	12.1(4a) and later
CSCvp92269	Running a Qualys security scan results in the following message: CWE - 693 Protection Mechanism Failure - "HTTP Security Header Not Detected"	12.1(4a) 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.	12.1(4a) and later
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.	12.1(4a) and later

Bug ID	Description	Exists In
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.	12.1(4a) 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.	12.1(4a) 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.

Table 5 Resolved Bugs in This Release

Bug ID	Description	Fixed in
CSCvb54482	When the hardware abstraction layer (HAL) dumps a core, running the 'show tech support' command causes a vsh core.	12.1(4a)
CSCvc43532	When an L3extSubnet with the import security scope is removed from L3extInstP, the aclqos process on a non-border leaf switch crashes.	12.1(4a)
CSCvc48434	A VRF instance gets stuck in the delete pending state.	12.1(4a)
CSCvc51529	Endpoints connected to a FEX vPC might move as traffic is load-balanced across vPC members. For example, if leaf1 and leaf2 are a vPC pair and have FEXes connected that have the vPC member ports, and if an endpoint that is learned on this vPC is communicating with an endpoint that is connected to some other leaf switch, that remote leaf will see a remote-to-remote endpoint move. This is because the source TEP being set by leaf 1 and leaf 2 is the physical leaf switch TEP rather than the vPC TEP that both own. Excessive endpoint moves could lead to other symptoms, such as learning getting disabled on a bridge domain.	12.1(4a)
CSCvc92067	IPV6 direct route installation fails in the hardware.	12.1(4a)
CSCvd54537	When upgrading from a 12.1(1) release to a 2.2(2) release, the ASIC interrupt can be hit at boot up, due to a possible issue in the order of initialization of egress multicast ASIC tables. This interrupt causes the SDK to dump a core. This issue is not consistently reproducible.	12.1(4a)
CSCvd70304	Configuring a QOS policy on the Cisco APIC for the infrastructure tenant has no effect on control plane traffic.	12.1(4a)

Bug ID	Description	Fixed in
CSCvd77155	An OSPF session between border ToR switches using the streched SVI will be down. Adjacency will point to the tunnel interface, which is incorrect.	12.1(4a)
CSCvd78419	There is a switch HAP reset due to a HAL process core.	12.1(4a)
CSCve19668	The ACLQOS or Hardware Abstraction Layer (HAL) process might fail when there are multiple access SPAN sessions configured in both directions coupled with tenant SPAN.	12.1(4a)
CSCve50164	A port-channel will be in the Failed (F) state and unrecoverable without reloading the switches.	12.1(4a)
CSCve75075	Spine switch in-band management does not work on the addition or removal of the default "allow all ip" contract.	12.1(4a)
CSCve80460	DTEP tunnels flap every 20 minutes.	12.1(4a)
CSCve82313	ARP is unable to be resolved for a host attached through Overlay Transport Virtualization (OTV) using the Common Pervasive Gateway feature.	12.1(4a)
CSCve84822	After restarting a Cisco ACI leaf switch, all statistics are 0 on one port and a Nexus 9000 switch cannot transfer packets received on that port.	12.1(4a)
CSCvf06257	In a multi-pod fabric, in-band management might not work across pods on Cisco Nexus 93128TX, 9396PX, or 9396TX-based spine switches.	12.1(4a)
CSCvf09329	The following message is consistently seen in the syslog: May 16 15:59:36 <leaf> <132> May 16 15:59:36 <leaf> %LOG_LOCAL0-4-SYSTEM_MSG [E4204936][transition][warning][sys] %IGMP-4-L3VM_LIBAPI_FAILED: VRF is NULL - failed in l3vm_if_get_context_info()	12.1(4a)
CSCvf09813	The fabric tracking feature stops working when HTTPS is disabled in the Cisco APIC.	12.1(4a)
CSCvf20789	DHCP relay cases can fail if VRF_ID is not added correctly to kic_db.	12.1(4a)
CSCvf21054	Manually running the "clear routing vrf <vrf> *" command in a leaf switch (vsh) disturbs in the Layer 3 routes.	12.1(4a)
CSCvf30762	ICMP redirects are enabled on the out-of-band management interface.	12.1(4a)
CSCvf31812	The GRE Keepalive "return" packet gets lost on the Cisco ACI fabric.	12.1(4a)
CSCvf42783	After removing PSU 1, the BFD sessions will flap.	12.1(4a)
CSCvf48161	A Cisco ACI spine switch might fail to program remote pods endpoint's information into a fabric card after a fabric/switch reload. This is a corner case with specific trigger. After that, the COOP repository shows that endpoint's MAC synthetic flag is not set as 0x5 or the endpoint's IP synthetic flag is not 0x25.	12.1(4a)
CSCvf48205	When PIM is enabled for a Layer 3 bridge domain and then disabled followed by IGMP snooping learning an mrouter port, COOP is not updated that the leaf switch is an mrouter for the bridge domain.	12.1(4a)

Bug ID	Description	Fixed in
CSCvf51513	A leaf switch crashes with the following details: Reason: reset-triggered-due-to-ha-policy-of-reset Service:eltm hap reset	12.1(4a)
CSCvf52033	A kernel crash is observed when an iping request is done with the -R option.	12.1(4a)
CSCvf61374	The "show interface" command produces the following error message on certain FEX ports: Undefined speed: 10 Error executing command, check logs for details	12.1(4a)
CSCvf65335	Fragmented packets with an IP fragment offset set to less than 0x14(20)=160 bytes are dropped on Cisco N9K-X9736C-EX, N9K-X97160YC-EX, and N9K-X9732C-EX leaf switches.	12.1(4a)
CSCvf66797	Fault F103824 "TCA: normalized temperature current value(eqptTemp5min:normalizedLast) value X% raised above threshold X%" is being triggered on /supslot-1/sup/sensor-2.	12.1(4a)
CSCvf71926	Packets selected for policy/contract logging are sampled at a faster timer (752 usec), instead of the configured slower timer (2800 miliseconds). However, this issue is not externally visible because there is a to-cpu-policier that is policing at 500 packets-per-second and thus nullifies this aggressive sampling.	12.1(4a)
CSCvf80846	The ICMPv6 process might crash even though there is no IPv6 enabled in the Cisco ACI fabric.	12.1(4a)
CSCvf89664	An SNMP core is dumped when an SNMP context name is configured as null.	12.1(4a)
CSCvf91594	In a scenario with route leaking between VRF instances using a vzAny contract and individual (non-shared) L3Outs are configured, the route map that is applied by system with prefix-list matches is not getting cleared when the configuration is reverted.	12.1(4a)
CSCvf93609	Endpoints in a Cisco ACI fabric were removed, causing a disruption.	12.1(4a)
CSCvf99198	Bidirectional forwarding detection (BFD) sessions between a Cisco ACI ToR and N7K keep flapping and never come up.	12.1(4a)
CSCvg00499	LLDP adjacency to looseNode objects is up, but the VLANs are not programmed on the ports. Run the following command in the leaf switch's CLI to see that there is no leqptRsLsNodeToIf object for the interface: moquery -c leqptRsLsNodeToIf	12.1(4a)
CSCvg02010	In a multipod environment, a multicast receiver sends the proper IGMP packets, but one of the border leaf switches does not add the mroute. When this border leaf switch is also the stripe-winner of the group, multicast does not work.	12.1(4a)
CSCvg05802	A leaf switch crashes after techsupport collection is triggered for that leaf switch.	12.1(4a)
CSCvg09582	EIGRP adjacency fails between endpoints and EIGRP routers in the same Layer 2 EPG.	12.1(4a)

Bug ID	Description	Fixed in
CSCvg16267	The dynamic tunnel that is needed for inter-pod traffic does not get setup when the stream that is sent is subjected to a policy-based redirect policy.	12.1(4a)
CSCvg25311	The "techsupport local" command might not complete collecting logs.	12.1(4a)
CSCvg38710	An IPFIB core is seen on a leaf switch. System Reset reason: Reason: reset-requested-due-to-fatal-module-error Service:Service on linecard had a hap-reset	12.1(4a)
CSCvg44842	A leaf switch reloads and the reset-reason is "epm hap reset". Furthermore, there is an EPM core file from the leaf switch.	12.1(4a)
CSCvg49121	There is an inconsistency between the Forwarding Information Base (FIB) table and hardware, which results in traffic being dropped, and there is a minor memory leak.	12.1(4a)
CSCvg58924	ICMP traffic to between endpoints might get dropped, but other traffic, including TCP and UDP, are not impacted.	12.1(4a)
CSCvg62580	Entries are not pushed to broadcom hardware level. Issue could be seen as VLANs not pushed to interfaces or ARP issues could also be seen if port is not forwarding for the VLAN in question.	12.1(4a)
CSCvg66519	Unicast DHCP packets received by a Cisco ACI leaf switch will be redirected to the CPU and dropped, but the broadcast DHCP discover message will be forwarded to the DHCP server.	12.1(4a)
CSCvg81628	When using DHCP relay in Cisco ACI, some DHCP clients randomly failed to get the IP address from the server. From the sniffer in the DHCP server side, the DHCP discovery packet was received and the DHCP server relayed the DHCP offer correctly to fabric. The following command run on the bridge domain's last reported switch shows that DHCP server entries are missing in it: cat /proc/kic_database | grep -A 40 "Routing Table"	12.1(4a)
CSCvh09693	The DHCP offer that goes back to the client has its src MAC set to 0.	12.1(4a)
CSCvh48223	A leaf switch reloads unexpectedly. The reset reason indicates a kernel panic and the "show processes log" command shows that nginx produced a stack trace.	12.1(4a)
CSCvh74146	A leaf switch resets due to an acllog process crash.	12.1(4a)
CSCvh99171	The Policy-Mgr NXOS process dumps a core when too many management VRF instance prefixes are added and deleted.	12.1(4a)
CSCvi32353	Fault F3073 is generated for a Cisco ACI switch, indicating that the SSD has reached 90% of the lifetime endurance limit. Under certain conditions this fault might be generated erroneously while the SSD is in good health.	12.1(4a)

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 12.1(4) releases in which the known behavior exists. A bug might also exist in releases other than the 12.1(4) releases.

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.	12.1(4a) 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.	12.1(4a) and later
CSCup65586	The show interface command shows the tunnel's Rx/Tx counters as 0.	12.1(4a) 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.	12.1(4a) and later
CSCup92534	Continuous "threshold exceeded" messages are generated from the fabric.	12.1(4a) and later
CSCuq39829	Switch rescue user ("admin") can log into fabric switches even when TACACS is selected as the default login realm.	12.1(4a) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) and later

Bug ID	Description	Exists In
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.	12.1(4a) 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.	12.1(4a) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	12.1(4a) and later
CSCur81822	The access-port operational status is always "trunk".	12.1(4a) 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.	12.1(4a) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) and later

Bug ID	Description	Exists In
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.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) and later
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	12.1(4a) and later

Bug ID	Description	Exists In
CSCuy16355	Transit traffic is dropped during ingress or egress when configured under the same Layer 3 Out with 0.0.0.0/0 security import subnet. This behavior is true for dynamic or static routing. To prevent this behavior, you must define more specific subnets and set the policy control enforcement preference to unenforced when configuring the associated VRF.	12.1(4a) and later
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) 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.	12.1(4a) and later
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to "0", and so traffic will still flow.	12.1(4a) and later
CSCuy96912	The debounce timer is not supported on 25G links.	12.1(4a) 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.	12.1(4a) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	12.1(4a) 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.	12.1(4a) and later
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	12.1(4a) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	12.1(4a) and later

Bug ID	Description	Exists In
CSCva21406	When nodes in the pod are running with mixed releases of the 12.0(x) release and pre-11.2(2) release, this can lead ISIS to core on the pre-11.2(2) release nodes.	12.1(4a) 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.	12.1(4a) and later
CSCvb36823	With VRF scale and 2 spine switches, reloading a spine switch will take time for the switch to re-join the fabric. During that time, the traffic will flow through the other spine switch.	12.1(4a) 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.	12.1(4a) and later
CSCvd10914	A svc_mgr core is seen on a TOR switch, and there is no space left on the device.	12.1(4a) 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.
- 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.

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