Cisco NX-OS Release 12.1(2) Release Notes for Cisco Nexus 9000 Series ACI-Mode Switches

This document describes the features, caveats, 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 2.1(2), Release Notes*, which you can view at the following location:

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

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

Release notes are sometimes updated with new information about restrictions and caveats. See the following website for the most recent version of the *Cisco NX-OS Release 12.1(2) 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
February 18, 2017	12.1(2e): Release 12.1(2e) became available. Created this document.
April 27, 2017	12.1(2g): Release 12.1(2g) became available. Added the resolved caveats for this release.
October 23, 2017	12.1(2k): Release 12.1(2k) became available. Added the resolved caveats for this release.

Contents

This document includes the following sections:

- Cisco Nexus 9000 Series ACI-Mode
- Supported Hardware
- Supported FEX Models
- New and Changed Information
- Installation Notes
- Compatibility Information
- Usage Guidelines
- Caveats
- Related Documentation

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-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 **Note:** This switch supports up to 10 line cards.
Spine switch fan	N9K-C9300-FAN3	Port side intake fan

Hardware Type	Product ID	Description
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
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
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+.

Hardware Type	Product ID	Description
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
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

Hardware Type	Product ID	Description
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.
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

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
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 100Mî /1/10 Gigabit Ethernet and Unified Port host interfaces (SFP+) and up to 6î 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(2) 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).
- 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.

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
- Leafs and spines from two different fabrics cannot be connected regardless of whether the links are administratively kept down.

Caveats

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

- Known Limitations
- Open Caveats
- Resolved Caveats
- 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 Caveats

This section lists the open caveats. Click the bug ID to access the Bug Search tool and see additional information about the bug. If a caveat is fixed in a patch of this release, the "Fixed In" column of the tables specifies the release.

Open Caveats in the 12.1(2e) Release

Table 4 lists the open caveats in the 12.1(2e) release.

Table 4 Open Caveats in the 12.1(2e) Release

Bug ID	Description	Fixed In
CSCun35596	FEX logs are missing in the output of the **show fex detail** command.	
CSCun96495	The events and faults for interfaces are not updated under **Ports** in the GUI.	
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.	
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.	
CSCur32247	FEX-related diagnostic results are missing.	
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.	
CSCuz82233	The server virtual Fibre Channel interface state changes to "port reinit limit reached" when an NP link is shut down.	
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.	
CSCvb12858	With passive QSA, a GLC-SX-MMD transceiver is not detected by N9K-93108TC-EX and N9K-93180YC-EX switches.	

Bug ID	Description	Fixed In
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.	
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.	
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.	
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.	

Open Caveats in the 12.1(2g) Release

There are no new open caveats in the 12.1(2g) release.

Open Caveats in the 12.1(2k) Release

There are no new open caveats in the 12.1(2k) release.

Resolved Caveats

This section lists the resolved caveats. Click the bug ID to access the Bug Search tool and see additional information about the bug.

Resolved Caveats in the 12.1(2e) Release

Table 5 lists the resolved caveats in the 12.1(2e) release.

Table 5 Resolved Caveats in the 12.1(2e) Release

Bug ID	Description
CSCva60517	IPFIB crashes on the Cisco ACI leaf switches when there is a high number of IPv6 neighbor discovers present on one or more Layer 2-only bridge domains.
CSCva97057	Fault (F1186) is raised on a FEX interface due to the following port configuration failure: Port configuration failure. Reason: 2 Failed Config: l1:PhysIfautoNeg_failed_flag
CSCvb26548	A leaf switch reloads intermittently with the following reset-reason: "reset-requested-due-to-fatal-module-error." The switch also generates cores against IPFIB.
CSCvb48915	The Cisco Nexus 2348UPQ fabric extender reloads with the following reset reason: "Tiburon Hap Reset."
CSCvb57398	A Nexus 9300-EX switch drops encapsulated traffic from itself for all protocols except ICMP. This occurs with Nexus 9300-EX switches when a host in an endpoint group that is connected to leaf A sends traffic to an SVI that is also configured on leaf A, but is in a different VRF.
CSCvb64312	An ARP response through an L3Out is denied when the sender.mac is not equal to the ethernet.src.mac.

Bug ID	Description
CSCvb64809	A leaf switch responds to an ARP request with all zeroes for the source MAC address.
CSCvb67540	Priority flow control (PFC) frames sent to a fabric extender (FEX) that is connected to a leaf switch in a Cisco ACI fabric will be flooded to all host interfaces (HIF) of all FEXs that are connected to the leaf switch.
CSCvb70363	If the switches that are connected to Cisco ACI Nexus 9000 switches only have an IPv6 management address configured, the Cisco ACI switches do not show the neighbor's management address correctly.
CSCvb72043	IP endpoint programming fails with the following error messages in the epmc-trace.txt log: [2016 Oct 12 14:31:03.088915862:215149860:epmc_pd_bcm_batch_flush:1006:E] BCM_ERR: Failed TLV_TYPE_MY_L3_HOST_CREATE rv:-6 host_rv:0 egr_rv:-6 [2016 Oct 12 14:31:03.088916854:215149861:epmc_pd_bcm_ep_prog_fail:716:E] Programming failed for ep key: EPM EP IPV4 key :: VRF 2686976 IP:=10.58.12.22 [2016 Oct 12 14:31:03.088920707:215149863:epmc_pd_bcm_ep_prog_fail:728:E] BCM programming failed for EP, inserting ep to failed list
CSCvb74390	When a port channel interface goes up by reaching min-link, the LACP port is sometimes flapped.
CSCvb84796	The coop process crashes on a spine switch due to a memory leak. The core dump file has a size of almost 4GB. This occurs when there are many IPv6 moves.
CSCvb96237	Traffic drops between two endpoints in the same endpoint group across different EX leaf switches.
CSCvc14935	A manual upgrade of the EPLD software on an N9K-C93120TX switch will fail.
CSCvc18574	Local endpoints are installed without a "Bounce" flag and point to a remote tunnel instead of a local interface.
CSCvc18609	Non-fully meshed N9K-93108TC-EX and N9K-93180YC-EX border leaf switches might fail to resolve ARP to addresses in an L3Out-connected subnet.
CSCvc20190	Double-bit parity errors that occur in specific Broadcom tables are not reported.
CSCvc22366	A leaf switch reloads intermittently with the following reset-reason: "reset-requested-due-to-fatal-module-error." The switch also generates cores against IPFIB.
CSCvc33849	AVS attached endpoints are not learned locally and the AVS tunnel is incorrectly pointing to the fabric instead of a local vPC interface.
CSCvc46922	The OSPF database descriptor (DBD) unicast packet is not forwarded correctly within a Cisco APIC bridge domain when the ingress leaf switch is second generation hardware, such as the N9K-C93180YC-EX switch, and the egress leaf switch is first generation hardware, such as the N9K-C9372PX switch. The OSPF DBD unicast packet is sent out with double VLAN headers with the wrong VLAN ID from the egress leaf switch.
CSCvc66846	VMs attached to FEX are not reachable after powering off the host.

Bug ID	Description
CSCvc66860	The CPU queue fills up after pause frames are received at a high rate. The issue can manifest itself in various ways as CPU transmit frames are dropped in hardware: 1. FEX might fail to come online if this queue is full because control plane process related to FEX discovery cannot send/receive all messages responsible for the bring-up. 2. Intermittent connectivity is noticed for endpoints connected to ACI leaf switches, as ARP failed to resolve on a host. 3. Spine links to affected leaf switches might be marked as "out-of-service" and LLDP adjacency will no longer be present.
CSCvc68361	Traffic fails between an intra-EPG isolated EPG and a non-Intra-EPG isolated EPG.
CSCvc76630	SNMPGET or SNMPWALK for OIDs in the MIB CISCO-ETHERNET-FABRIC-EXTENDER-MIB fails on Cisco ACI switches.
CSCvc84123	The Forwarding Information Base (FIB) crashes due to a BGP installed route that has the next hop TEP IP address of the fabric and an RMAC of the Layer 3 EVPN services over fabric WAN Cisco Data Center Interconnect (DCI) box.
CSCvc84758	Traffic received on the FEX host-interface (hif) vPC is dropped.
CSCvc85849	An HSRP vMAC endpoint could be aged out on remote leaf switch.
CSCvc96680	In a multipod configuration with a single L3Out distributed across four border leaf switches (two in pod 1 and two in pod 2), traffic might be dropped at the transit border leaf switch if the bounce bit is set.

Resolved Caveats in the 12.1(2g) Release

Table 6 lists the resolved caveats in the 12.1(2g) release.

Table 6 Resolved Caveats in the 12.1(2g) Release

Bug ID	Description
CSCvc92841	Endpoint Manager generates core after Virtual Ethernet Module (VEM) restart on multiple hosts. Switch reloads with an EPM core file generate.
CSCvd56425	Inability to use in-band address on spine switches for any management functions.
CSCvd62018	Removing a fibre optic cable results in traffic loss, which causes the Hardware Abstraction Layer to generate a core.
CSCvd04264	Sclass is set to zero when a remote endpoint is discovered through Generic Attribute Registration Protocol (GARP) resulting in possible drop in endpoint.
CSCvd86264	Dynamic Host Configuration Protocol (DHCP) stops working through the fabric after upgrade to APIC 2.1.2e and Cisco NX-OS Release 12.1.2e.
CSCvc43532	When the l3extSubnet with import security scope is removed from the l3extInstP, aclqos process(es) on non-boarder leaf crashes.
CSCvd17158 CSCve03012	Possible traffic polarization on external leaf.
CSCvd72899	Ongoing issue with atomic counter failing in switch when the ingress and egress of the tunnel end point details are not available in the tunnel managed object bank objects.

Bug ID	Description
CSCvd96555	Leaf generates core(s) for sdkhal_log in Cisco APIC GUI.
CSCvb42851	Possible memory leakage in stats manager.

Resolved Caveats in the 12.1(2k) Release

Table 7 lists the resolved caveats in the 12.1(2k) release.

Table 7 Resolved Caveats in the 12.1(2k) Release

Bug ID	Description
CSCvd15040	Second generation Cisco ACI leaf switches (-EX models) drop the LLC frame with the Layer 2 trailer.
CSCvd47962	A core is dumped in one of the switches that has the Application Spine Engine (ASE).
CSCve40166	After rebooting the Cisco ACI C9504 switches in ACI mode, the standby supervisor engine might go into the loader prompt instead of booting up from the new image.
CSCve67147	The N9K-X9736C-FX line card goes into the failed state after a few hours.
CSCvf67718	A fabric module crashes and reloads when it encounters the PCIe uncorrectable error.
CSCvf81062	The listen socket on the supervisor engine for FTP does not work.

Known Behaviors

This section lists caveats that describe known behaviors. Click the Bug ID to access the Bug Search Tool and see additional information about the bug.

Known Behaviors in the 12.1(2e) Release

Table 8 lists caveats that describe known behaviors in the 12.1(2e) release.

Table 8 Known Behaviors in the 12.1(2e) Release

Bug ID	Description
CSCu037016	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.
CSCu050533	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.
CSCup65586	The **show interface** command shows the tunnel's Rx/Tx counters as 0.
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.
CSCup92534	Continuous "threshold exceeded" messages are generated from the fabric.

Bug ID	Description
CSCuq39829	Switch rescue user ("admin") can log into fabric switches even when TACACS is selected as the default login realm.
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.
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.
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.
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.
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.
CSCur01336	The PSU does not get detected after an OIR device with power input is connected.
CSCur81822	The access-port operational status is always "trunk".
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.
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).
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.
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.
CSCus61748	If the TOR 1RU system is configured with the RED fan (the reverse airflow), the air will flow from back to front. 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 front to back. 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.
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.
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.

Bug ID	Description
CSCuu11351	Traffic from the orphan port to the vPC pair is only updated on the destination node, so the traffic count shows as excess.
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.
CSCuv57302	Atomic counters on the border leaf do not increment for traffic from an endpoint group going to the Layer 3 out interface.
CSCuv57315	Atomic counters on the border leaf do not increment for traffic from the Layer 3 out interface to an internal remote endpoint group.
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.
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.
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.
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.
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.
CSCuy06749	Traffic is dropped between two isolated EPGs.
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.
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.
CSCuy47634	EPG statistics only count total bytes and packets. The breakdown of statistics into multicast/unicast/broadcast is not available on new hardware.
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.
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to "0", and so traffic will still flow.
CSCuy77579	For N9K-93180YC-EX switches, the unicast reachability of a multicast source is required for forwarding on Shortest-Path-Tree (SPT). However, if unicast reachability is not present for a multicast source that is outside of the fabric, receivers inside of the fabric in certain cases can still get the packets on the shared tree even after the border leaf switches have switched over to the SPT in the control plane. In other words, the PIM protocol might have switched over to SPT by installing a <S,G> route with NULL-RPF, but in the data plane, packets might still get forwarded on the shared <*,G> tree.
CSCuy96912	The debounce timer is not supported on 25G links.
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.
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.

Bug ID	Description
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.
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.
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.
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.
CSCvd10914	A svc_mgr core is seen on a TOR switch, and there is no space left on the device.

- 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 Cisco Nexus 9508 ACI-mode switch supports warm (stateless) standby where the state is not synched between the active and the standby supervisor modules. For an online insertion and removal (OIR) or reload of the active supervisor module, the standby supervisor module becomes active, but all modules in the switch are reset because the switchover is stateless. In the output of the **show system redundancy status** command, warm standby indicates stateless mode.
- When a recommissioned APIC controller rejoins the cluster, GUI and CLI commands can time out while the cluster expands to include the recommissioned APIC controller.
- If connectivity to the APIC cluster is lost while a switch is being decommissioned, the decommissioned switch may not complete a clean reboot. In this case, the fabric administrator should manually complete a clean reboot of the decommissioned switch.
- 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.

Known Behaviors in the 12.1(2g) Release

There are no new known behaviors in the 12.1(2g) release.

Known Behaviors in the 12.1(2k) Release

There are no new known behaviors in the 12.1(2k) release.

Related Documentation

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2017 Cisco Systems, Inc. All rights reserved.