Cisco NX-OS Release 11.2(3) 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 1.2(3), 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 11.2(3) 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
April 8, 2016	Created the release notes for Release 11.2(3c).
May 15, 2016	11.2(3e): Added the content for release 1.2(3e).
June 17, 2016	11.2(3h): Added the content for release 11.2(3h).
October 28, 2016	In the Supported FEX Models section, added N2K-B22HP-P.
November 12, 2016	11.2(3m): Release 11.2(3m) became available. There are no changes to this document for this release.
November 16, 2016	11.2(3m): In the Resolved Caveats section, added bug CSCuz98339.
January 17, 2018	In the Compatibility Information section, changed "You cannot connect the APIC directly to the N9332PQ ACI spine" to "You cannot connect the APIC directly to the N9332PQ ACI leaf switch."

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 11.2 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 four slots
Chassis	N9K-C9508	Cisco Nexus 9508 chassis with 8 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
Spine switch	N9K-C9516	Cisco Nexus 9516 switch with 16 linecard slots **Note:** This switch supports up to 10 I/O modules.

Hardware Type	Product ID	Description
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
Spine switch module	N9K-C9508-FM	Fabric 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-C93120TX	Cisco Nexus 9300 with 96-port 1/10 Gigabit-T and 6-port 40 Gigabit Ethernet QSFP
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-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

Hardware Type	Product ID	Description
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
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.

Hardware Type	Product ID	Description
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
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

Product ID	Description
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 100Mï /1/10 Gigabit Ethernet and Unified Port host interfaces (SFP+) and up to 6ï 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 1.2(3) 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 4.2(1)SV2(2.3).
- 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 N9396PX, N9372PX, or N9372PX-E switches
- 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

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 11.2(3c) Release

Table 4 lists the open caveats in the 11.2(3c) release.

Table 4. Open Caveats in the 11.2(3c) 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.	
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.	
CSCuy40089	Shared L3out and spine L3out inband management do not work together.	
CSCuy47950	After completing the switch firmware upgrade from the APIC GUI, fault F1582 (firmware-version-mismatch) is raised for one of the supervisors of the switch. This fault is benign and will not impact any functionality.	
CSCuy49977	NX-OS does not boot properly and the admin login does not work.	

Open Caveats in the 11.2(3e) Release

There are no new open caveats in the 11.2(3e) release.

Open Caveats in the 11.2(3h) Release

There are no new open caveats in the 11.2(3h) release.

Open Caveats in the 11.2(3m) Release

There are no new open caveats in the 11.2(3m) 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 11.2(3c) Release

Table 5 lists the resolved caveats in the 1.2(3c) release.

Table 5 Resolved Caveats in the 1.2(3c) Release

Bug ID	Description
CSCuy18505	The GUI does not match the CLI for the Link Oper State when using LACP hot-standby.
CSCuy42951	Disk usage for /var/log is 100%.
CSCuy47950	After completing the switch firmware upgrade from the APIC GUI, fault F1582 (firmware-version-mismatch) is raised for one of the supervisors of the switch.
CSCuy65581	IP nodes on a path marked as 802.1P cannot communicate with other nodes because ARP resolution fails. On problematic nodes and their peers, ARP entries have the "Incomplete" status.
CSCuy77009	After connecting or disconnecting the cable in out-of-band management ports of ACI spines and leaves, there are some port state changes on the front ports and a link state change between the leaf and spine, which should not be affected by port link down/up of an out-of-band management port.
CSCuy79569	After upgrading to 1.2(2g), you might be unable to access hosts inside of the fabric from either outside or inside the fabric. You might see faults 1800 and 1801 raised on leaf switches that are connected to the hosts/border leafs.
CSCuy87642	A loop is formed when looping back interfaces between leafs in the same fabric. A wiring fault is raised and the physical interface is placed into the out-of-service (OOS) state, but it continues to forward a subset of traffic.
CSCuy90172	An ACI leaf does not bridge a ICMPv6 router solicitation (RS) message even when the ACI leaf does not have an IPv6 subnet (SVI) on its bridge domain. Instead, the ACI leaf always sends RS packets to its CPU in the event that the ACI leaf has an IPv6 SVI on it. As a result, if a host is waiting for an ICMPv6 router advertisement (RA) from an external IPv6 router, the host must wait until the IPv6 router sends out a periodic RA. The reason is that the ACI Leaf does not forward RS to the external IPv6 router and the external IPv6 router cannot respond to it without receiving it.

Resolved Caveats in the 11.2(3e) Release

Table 6 lists the resolved caveats in the 11.2(3e) release.

Table 6 Resolved Caveats in the 11.2(3e) Release

Bug ID	Description
CSCuw60935	FIB uses the /var/sysmgr space for binlog. If /var/sysmgr gets full for to any reason, FIB exits without leaving a core.
CSCuy08607	The pervasive subnet is missing in the endpoint manager client (EPMC).
CSCuz35582	On an ACI switch, DME log files with filenames in the format of svc_ifc_*.log.stderr can grow to a large size and eventually fill up the partition, which causes the switch to restart.
CSCuz56785	LLDP dumps a core in tg_is_timer_active () during an upgrade.

Resolved Caveats in the 11.2(3h) Release

Table 7 lists the resolved caveats in the 11.2(3h) release.

Table 7 Resolved Caveats in the 11.2(3h) Release

Bug ID	Description
CSCuy93241	A leaf node is reloaded due to an EPMC crash, which occurs due to an EPMC process infinite loop.
CSCuz55258	An issue occurs that prevents SSH connections to the leaf nodes from the APIC (Infra-Band) or externally from another device to the OOB or In-Band management addresses.
CSCuz63747	When nodes in a pod are running with mixed versions of 12.0 and pre-11.2(2) versions, this can cause ISIS to crash on the pre-11.2(2) version nodes.
CSCuz66812	Multiple N9K-9372TX interfaces become err-disabled due to a sequence timeout.
CSCva05997	A bcm_usd crash causes ACI leaf switches to reload.

Resolved Caveats in the 11.2(3m) Release

Table 8 lists the resolved caveats in the 11.2(3m) release.

Table 8 Resolved Caveats in the 11.2(3m) Release

Bug ID	Description
CSCuz98339	A switch that has been up for over 200 days sometimes crashes and reloads due to an EPM process memory leak. The memory leak occurs in cases where the MTS message type coming to the EPM was MTS_TYPE_NOTIFICATION.

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 11.2(3c) Release

Table 9 lists caveats that describe known behaviors in the 11.2(3c) release.

Table 9. Known Behaviors in the 11.2(3c) Release

Bug ID	Descriptiono
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	Descriptiono
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 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 IPs is reassigned to another server/host, the fabric detects it as an IP move and forwarding will work as expected.
CSCur01336	The PSU is not getting detected after OIR with Power input connected.
CSCur81822	The access-port operational status is 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 output incorrectly displays AOC cables as ACU cables.
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 & configuration, which causes the Inlet temperature sensor reading to be very close, equal, or even greater than the outlet temperature reading.
CSCus73592	10% to 11% traffic drops occur on Unicast Traffic Streams.
CSCut59020	If Backbone and NSSA areas are on the same leaf, and default route leak is enabled, Type-5 LSAs may not be redistributed to the Backbone area.

Bug ID	Descriptiono
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.
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 BD "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	iping 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.

- 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 (BD) is not supported. Only one EPG or Layer 2 outside for a given BD 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 BD. This knob can be set to "Optimized Flood" only for a maximum of 50 BDs per leaf.

  If "Optimized Flood" is enabled for more than the supported number of BDs on a leaf, follow these configuration steps to recover:

  - Set "unknown multicast flooding" to "Flood" for all BDs mapped to a leaf.
  - Set "unknown multicast flooding" to "Optimized Flood" on needed BDs.

Known Behaviors in the 11.2(3e) Release

There are no new known behaviors in the 11.2(3e) release.

Known Behaviors in the 11.2(3h) Release

There are no new known behaviors in the 11.2(3h) release.

Known Behaviors in the 11.2(3m) Release

There are no new known behaviors in the 11.2(3m) 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