



Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 12.2(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 2.2(4), Release Notes*, which you can view at the following location:

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

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

Release notes are sometimes updated with new information about restrictions and bugs. See the following website for the most recent version of the *Cisco NX-OS Release 12.2(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
July 6, 2021	In the Supported Hardware section, added the NXA-PAC-500W-PI and NXA-PAC-500W-PE PSUs.
June 24, 2021	Added open issue CSCvu07844.
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 13, 2020	12.2(4f): In the Resolved Bugs section, added bug CSCvr98827.
September 20, 2019	In the Usage Guidelines section, added the following bullet: <ul style="list-style-type: none">■ 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.

Contents

Date	Description
September 11, 2019	<p>In the Supported Hardware section, for the N9K-C93108TC-FX and N9K-C93180YC-FX switches, added the following note:</p> <p>Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.</p>
August 14, 2019	12.2(4f): In the Open Bugs section, added bug CSCvp92269.
July 31, 2019	<p>In the Compatibility Information section, added the following bullet:</p> <ul style="list-style-type: none"> ■ 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.
January 30, 2019	12.2(1n): In the Open Bugs section, added bug CSCvn69340.
January 28, 2019	<p>12.2(4f): In the Open Bugs section, added bug CSCvi76161.</p> <p>12.2(4p): In the Resolved Bugs section, added bug CSCvi76161.</p>
October 24, 2018	12.2(4f): In the Open Bugs section, added bug CSCvm66008.
August 7, 2018	12.2(4r): Release 12.2(4r) became available; there are no changes to this document for this release.
July 2, 2018	12.2(4q): Release 12.2(4q) became available; there are no changes to this document for this release.
June 7, 2018	12.2(4f): In the Open Bugs section, added bug CSCvg35892.
May 18, 2018	12.2(4p): Release 12.2(4p) became available. Added the resolved bugs for this release.
April 24, 2018	12.2(4f): In the Open Bugs section, added bug CSCvi57920.
April 6, 2018	12.2(4f): Release 12.2(4f) became available.

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
- Bugs
- 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 12.3 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-M12PQ	12-port or 8-port
Pluggable module (GEM)	N9K-M6PQ	6-port
Pluggable module (GEM)	N9K-M6PQ-E	6-port, 40 Gigabit Ethernet expansion module
Spine switch	N9K-C9336PQ	Cisco Nexus 9336PQ switch, 36-port 40 Gigabit Ethernet QSFP
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

Supported Hardware

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 platform switch with 48 1/10GBASE-T (copper) APIC-facing ports and 6 100-Gigabit Ethernet QSFP28 spine-facing ports

Supported Hardware

Hardware Type	Product ID	Description
Top-of-rack (ToR) leaf switch	N9K-C93108TC-FX	<p>Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) APIC-facing ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports.</p> <p>Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.</p>
Top-of-rack (ToR) leaf switch	N9K-C93120TX	<p>Cisco Nexus 9300 platform switch with 96 1/10GBASE-T (copper) APIC-facing ports and 6-port 40-Gigabit Ethernet QSFP spine-facing ports.</p>
Top-of-rack (ToR) leaf switch	N9K-C93128TX	<p>Cisco Nexus 9300 platform switch with 96 1/10GBASE-T (copper) APIC-facing ports and 6 or 8 40-Gigabit Ethernet QSFP spine-facing ports</p>
Top-of-rack (ToR) leaf switch	N9K-C93180LC-EX	<p>Cisco Nexus 9300 platform switch with 24 40-Gigabit APIC-facing ports and 6 10/40/100-Gigabit QSFP28 spine-facing ports</p> <p>Note: This switch has the following limitations:</p> <ul style="list-style-type: none"> • This release does not support 1 Gbps for QSA. • The top and bottom ports must use the same speed. If there is a speed mismatch, the top port takes precedence and bottom port will be error disabled. Both ports both must be used in either the 40 Gbps or 10 Gbps mode. • Ports 26 and 28 are hardware disabled. • This release supports 40 and 100 Gbps for the front panel ports. The uplink ports can be used at the 100 Gbps speed.

Supported Hardware

Hardware Type	Product ID	Description
Top-of-rack (ToR) leaf switch	N9K-C93180YC-EX	Cisco Nexus 9300 platform switch with 48 10/25-Gigabit APIC-facing ports and 6-port 40/100 Gigabit QSFP28 spine-facing ports.
Top-of-rack (ToR) leaf switch	N9K-C93180YC-FX	Cisco Nexus 9300 platform switch with 48 10/25-Gigabit Ethernet SFP28 APIC-facing ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. The SFP28 ports support 1-, 10-, and 25-Gigabit Ethernet connections and 8-, 16-, and 32-Gigabit Fibre Channel connections. Note: Incoming FCOE packets are redirected by the supervisor module. The data plane-forwarded packets are dropped and are counted as forward drops instead of as supervisor module drops.
Top-of-rack (ToR) leaf switch	N9K-C9332PQ	Cisco Nexus 9332PQ Top-of-rack (ToR) Layer 3 switch with 32 40-Gigabit Ethernet QSFP+ APIC-facing ports
Top-of-rack (ToR) leaf switch	N9K-C9372PX	Cisco Nexus 9372PX Top-of-rack (ToR) Layer 3 switch with 48 10GBASE-T (copper) Ethernet SFP+ APIC-facing ports and 6 40-Gbps Ethernet QSFP+ spine-facing ports Note: Only the downlink ports 1-16 and 33-48 are capable of supporting SFP1-10G-ZR SFP+.
Top-of-rack (ToR) leaf switch	N9K-C9372PX-E	Cisco Nexus 9372PX-E Top-of-rack (ToR) Layer 3 switch with 48 10GBASE-T (copper) Ethernet SFP+ APIC-facing ports and 6 40-Gbps Ethernet QSFP+ spine-facing ports Note: Only the downlink ports 1-16 and 33-48 are capable of supporting SFP1-10G-ZR SFP+.
Top-of-rack (ToR) leaf switch	N9K-C9372TX	Cisco Nexus 9372TX Top-of-rack (ToR) Layer 3 switch with 48 1/10GBASE-T (copper) APIC-facing ports and 6 40-Gbps Ethernet QSFP spine-facing ports

Supported Hardware

Hardware Type	Product ID	Description
Top-of-rack (ToR) leaf switch	N9K-C9372TX-E	Cisco Nexus 9372TX-E Top-of-rack (ToR) Layer 3 switch with 48 10GBASE-T (copper) APIC-facing ports and 6 40-Gbps Ethernet QSFP+ spine-facing ports
Top-of-rack (ToR) leaf switch	N9K-C9396PX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) SFP+ APIC-facing ports and 6 or 12 40-Gigabit Ethernet QSFP spine-facing ports
Top-of-rack (ToR) leaf switch	N9K-C9396TX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) APIC-facing ports and 6 or 12 40-Gigabit Ethernet QSFP spine-facing ports
Top-of-rack (ToR) leaf switch fan	NXA-FAN-30CFM-B	Port side intake fan
Top-of-rack (ToR) leaf switch fan	NXA-FAN-30CFM-F	Port side exhaust fan
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-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-650W	650W AC Power supply, port side intake pluggable
Top-of-rack (ToR) leaf switch power supply unit	N9K-PAC-650W-B	650W AC Power supply, port side exhaust pluggable

Supported Hardware

Hardware Type	Product ID	Description
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.
Top-of-rack (ToR) leaf switch power supply unit	NXA-PAC-500W-PE	500W AC Power supply, port side exhaust pluggable
Top-of-rack (ToR) leaf switch power supply unit	NXA-PAC-500W-PI	500W AC Power supply, port side intake pluggable
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 power supply unit	UCS-PSU-6332-DC	930W DC power supply, reversed airflow (port side exhaust)

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
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-C2348UPO	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 (required SFP/SFP+) + 8x10GE (required 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 (required 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

There are no new hardware features in this release.

New Software Features

For new software features, see the *Cisco APIC 2.2(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>

Compatibility Information

- This release supports the hardware and software listed on the Cisco ACI Ecosystem Compatibility List, and supports the Cisco AVS, Release 5.2(1)SV3(2.14).
- Link level flow control is not supported on ACI-mode switches.
- To connect the N2348UPQ to Cisco ACI leaf switches, the following options are available:
 - Directly connect the 40G FEX ports on the N2348UPQ to the 40G switch ports on the Cisco ACI leaf switches
 - Break out the 40G FEX ports on the N2348UPQ to 4x10G ports and connect to the 10G ports on all other Cisco 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 Cisco ACI fabric, it is required to have a 10G interface on the ACI leaf. You cannot connect the Cisco APIC directly to the N9332PQ Cisco 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

Bugs

- 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](#)
- [Open Bugs](#)
- [Resolved Bugs](#)
- [Known Behaviors](#)

Known Limitations

The following list describes IpEpg (IpCkt) known limitations in this release:

- An IP/MAC Ckt endpoint configuration is not supported in combination with static endpoint configurations.
- An IP/MAC Ckt endpoint configuration is not supported with Layer 2-only bridge domains. Such a configuration will not be blocked, but the configuration will not take effect as there is no Layer 3 learning in these bridge domains.
- An IP/MAC Ckt endpoint configuration is not supported with external and Infra bridge domains because there is no Layer 3 learning in these bridge domains.
- An IP/MAC Ckt endpoint configuration is not supported with a shared services provider configuration. The same or overlapping prefix cannot be used for a shared services provider and IP Ckt endpoint. However, this configuration can be applied in bridge domains having shared services consumer endpoint groups.
- 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

Bugs

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 12.2(4) releases in which the bug exists. A bug might also exist in releases other than the 12.2(4) releases.

Table 4 Open Bugs in This Release

Bug ID	Description	Exists In
CSCvj35940	In multi-pod, broadcast/multicast/unknown unicast traffic gets dropped when a transit TOR boots up after a reload.	12.2(4p) and later
CSCvj35961	When there is a high bridge domain scale on a heavily loaded spine switch, when a leaf switch reload is triggered, depending on the amount of reprogramming needed, there might be a loss of 3-5 seconds seen in some instances.	12.2(4p) and later
CSCvc17015	Quick changes in the RPF interface trigger RPF lookups in MRIB and PIM for all the routes.	12.2(4f) and later
CSCvd34047	Many MTS messages are seen in the MTS buffer for SAP 213.	12.2(4f) and later
CSCve06334	MAC and IP endpoints are not learned on the local vPC pair.	12.2(4f) and later

Bugs

Bug ID	Description	Exists In
CSCvf09313	In the 12.2(2i) release, the BPDU filter only prevents interfaces from sending BPDUs, but does not prevent interfaces from receiving BPDUs.	12.2(4f) and later
CSCvf85014	Multicast traffic loss might occur during spine switch upgrade/reload when at least one fabric link port from spine switch to ToR is in "admin down" state.	12.2(4f) and later
CSCvg35892	There is a delay in internal process communication of over a minute due to a timeout (timeout period is 60 seconds). This can manifest itself in different ways. In one way, the Policy Element sent a delete/create request for a VLAN to the VLAN manager, but the VLAN manager did not process the request until the initial request timed out. This resulted in an outage that lasted a little over a minute while changing a BD from Unknown Unicast Flood to Hardware Proxy.	12.2(4f) and later
CSCvi57920	A UCS 1225 vNIC goes down after changing the peer Cisco ACI leaf node name.	12.2(4f) and later
CSCvi76161	A version mismatch between Cisco ACI leaf switches causes the EPM process to crash and a HAP reset to occur.	12.2(4f) and later
CSCvm66008	When you have two NSSA areas deployed on a vPC leaf switch pair that also see each other as OSPF peers and both L3Outs have the default route leak policy configured, then if you delete one of the L3Outs, the following things occur: OSPF generates the default route without a tag The peer installs that route in RIB The peer redistributes the route into BGP The peer advertises the route to rest of the leaf switches This issue blackholes all traffic that uses the default route.	12.2(4f) and later
CSCvn69340	The BFD session does not get instantiated under some circumstances in one of the VPC legs for static routes.	12.2(4f) and later
CSCvp50075	A leaf switch experiences an unexpected reload due to a HAP reset.	12.2(4f) and later
CSCvp92269	Running a Qualys security scan results in the following message: CWE - 693 Protection Mechanism Failure - "HTTP Security Header Not Detected"	12.2(4f) 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.	12.2(4f) and later

Bugs

Bug ID	Description	Exists In
CSCvs18150	<p>After a certain set of steps, it is observed that the deny-external-tag route-map used for transit routing loop prevention gets set back to the default tag 4294967295. Since routes arriving in Cisco ACI with this tag are denied from being installed in the routing table, if the VRF table that has the route-tag policy is providing transit for another VRF table in Cisco ACI (for instance and inside and outside VRF with a firewall connecting them) and the non-transit VRF table has the default route-tag policy, routes from the non-transit VRF table would not be installed in the transit VRF table.</p> <p>This bug is also particularly impactful in scenarios where transit routing is being used and OSPF or EIGRP is used on a vPC border leaf switch pair. vPC border leaf switches peer with each other, so if member A gets a transit route from BGP, redistributes into OSPF, and then advertises to member B (since they are peers)...without a loop prevention mechanism, member B would install the route through OSPF since it has a better admin distance and would then advertise back into BGP. This VRF tag is set on redistribution of BGP > OSPF and then as a table map in OSPF that blocks routes with the tag from getting installed in the routing table. When hitting this bug, the route-map used for redistributing into OSPF still sets the tag to the correct value. However, the table map no longer matches the correct tag. Rather, it matches the default tag. As a result, member A (could be B) would install the route through OSPF pointing to B. It would then redistribute it back into BGP with the med set to 1. The rest of the fabric (including member B) would install the BGP route pointing to member A since its med is better than the original route's med.</p>	12.2(4f) 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.	12.2(4f) and later
CSCvt82388	A switch SSD fails in less than two years and needs replacement. The /mnt/pss/ssd_log_amp.log file shows daily P/E cycles increasing by 10 or more each day, and fault "F3525: High SSD usage" is observed. ARP/ICMPv6 adjacency updates can also contribute to many SSD writes.	12.2(4f) 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.	12.2(4f) and later
CSCvu07844	When a Cisco N9K-C93180LC-EX, N9K-93180YC-EX, or N9K-C93108TC-EX leaf switch receives control, data, or BUM traffic from the front panel ports with the storm policer configured for BUM traffic, the storm policer will not get enforced. As such, the switch will let all such traffic through the system.	12.2(4f) 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.

Bugs

Table 5 Resolved Bugs in This Release

Bug ID	Description	Fixed In
CSCvd77159	Multicast or flood traffic is not sent out from the correct ports on a Cisco ACI leaf switch.	12.2(4p)
CSCve40174	The IP address is missing in COOP after performing some fast IP address moves across a pod or site.	12.2(4p)
CSCve75260	BGP and COOP are out of sync after endpoint flaps.	12.2(4p)
CSCvg05802	A leaf switch crashes after techsupport collection is triggered for that leaf switch.	12.2(4p)
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.	12.2(4p)
CSCvh84058	While configuring inband mgmt for Cisco Tetration Analytics, 3 out of 6 leaf switches reload due to the sdkhal process dumping a core.	12.2(4p)
CSCvi05703	There can be a loss of traffic when a new node comes up.	12.2(4p)
CSCvi11291	XR IP learning occurs on border leaf switches even when 'disabled remote EP learn' is configured for the fabric. XR IP learns should only occur on border leaf switches for routed multicast packets with PIM enabled.	12.2(4p)
CSCvi73383	A leaf switch reloads during techsupport collection. The reload reason is " Service on linecard had a hap-reset" and there is a core file generated for the ipfib process.	12.2(4p)
CSCvi75081	Hardware is mis-programmed for an endpoint that is learned on a vPC interface while its peer is down. As an example, if leaf101 and leaf102 are a vPC pair, the problem is that leaf102 has programmed the host MAC behind the vPC interface without the isvpc flag, which changes the hardware macdainfo interface value.	12.2(4p)
CSCvi76161	The EPM process sometimes crashes when the TOR is booting up.	12.2(4p)
CSCvi97670	Endpoints on an AVS host are intermittently or entirely unreachable. When running the "vemcmd show opflex" command on the AVS host, both Opflex tunnels are up/active. When running the "show endpoint ip X.X.X.X" command for the VTEP IP of the AVS host on the leaf switches where the AVS host connects, no entries are seen. Running the same command for the endpoint IP addresses results in no entries.	12.2(4p)
CSCvc67766	A prefix check is not performed when EP learning is triggered on a remote leaf switch through ARP.	12.2(4f)
CSCvc72792	There is a memory leak in the policy-mgr NX-OS process.	12.2(4f)
CSCvc80035	There is an issue that causes the SysRq signal to be received. Some of the possible symptoms are as follows: system crash system reload syslog level change	12.2(4f)

Bugs

Bug ID	Description	Fixed In
CSCvd94149	After upgrading or downgrading from the NX-OS 12.2(1n) release or 12.2(1o) release, the virtual IP address cannot be reached. A ping to the vIP will return a Destination Host Unreachable error and the ARP table on the source device will show that the MAC address for the vIP address is incomplete.	12.2(4f)
CSCve75932	A switch resets when 'ctrl+Break' and 'b' keys are pressed at the login prompt.	12.2(4f)
CSCve80460	DTEP tunnels flap every 20 minutes.	12.2(4f)
CSCvf11037	The 'show cdp neighbor' command run from ibash gives a different number of neighbors from same command on vsh.	12.2(4f)
CSCvf35156	<p>When allow-micro-segmentation is selected for a DVS VMM domain, proxy-arp is automatically enabled for the EPG VLAN. All intra-EPG ARP messages must be punted to the CPU and the switch must generate a proxy response indicating its own MAC address to get to other endpoints in the EPG.</p> <p>For FX switches that are running a Cisco NX-OS release prior to the 13.0 release, the ARP messages are not punted to the CPU, and so the leaf switch is never able to generate a proxy response. As a result, communication between endpoints in the same subnet and microsegmentation EPG will fail.</p>	12.2(4f)
CSCvf46351	A Cisco ACI leaf port with LACP enabled cannot come up. LACP is configured, but is inactive on the port.	12.2(4f)
CSCvf49746	<p>The contents of a USB drive formatted with the FAT file system from a switch in ACI-mode cannot be viewed. Because of this, files cannot be copied from the USB drive. Attempting to boot using an image on the USB drive from the loader fails with the following output:</p> <pre>Copying image /usb/usb1/aci-n9000-dk9.12.3.1f.bin to /bootflash /bin/cp: cannot stat `/usb/usb1/aci-n9000-dk9.12.3.1f.bin': No such file or directory Failed to open input file /bootflash/aci-n9000-dk9.12.3.1f.bin *** Extraction of system image failed ***</pre>	12.2(4f)
CSCvf57563	<p>Cisco Nexus 9000 Series Fabric Switches - ACI mode includes a version of the Linux Kernel that is affected by the vulnerabilities identified by the following Common Vulnerability and Exposures (CVE) IDs:</p> <p>CVE-2017-1000366, CVE-2017-1000379, CVE-2017-1000370, CVE-2017-1000365, CVE-2017-8831</p> <p>This bug was opened to address the potential impact on this product.</p>	12.2(4f)
CSCvf71926	Packets selected for policy/contract logging are sampled at a faster timer (752 usec), instead of the configured slower timer (2800 milliseconds). 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.	12.2(4f)

Bugs

Bug ID	Description	Fixed In
CSCvf73694	When using OpenStack to deploy a VM and the setup contains a FEX and N9K-C93180YC-EX switch, two NICs of one node go down. In addition, the Cisco OpFlex loses connection to the other nodes in the same FEX. As a result, the host loses infra VLAN connectivity to the fabric.	12.2(4f)
CSCvf91594	In a scenario with route leaking between VRF instances using a vzAny contract and individual (non-shared) L3Outs that are configured, the route map that is applied by the system with prefix-list matches is not getting cleared when the configuration is reverted. In the NX-OS 12.2(2o) release, the configuration was not reverted. In 12.2(2k), the configuration is sometimes reverted. The issue is unpredictable. Even deleting the VRF instances, contracts, and EPGs might not revert the internal route-map configuration.	12.2(4f)
CSCvf93609	Endpoints in a Cisco ACI fabric were removed, causing a disruption.	12.2(4f)
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.	12.2(4f)
CSCvg35224	While both top and bottom ports are up, if some down event occurs on the bottom port and the bottom port has a copper SFP (admin down or otherwise), the corresponding top port will also go down with it. It will show as the "link down not-connected" behavior for the top port: Ethernet1/15 is down (notconnect) admin state is up, Dedicated Interface This occurs even though the QSFP is still inserted and was otherwise previously up and functional.	12.2(4f)
CSCvg44842	A leaf switch reloads and the reset-reason is "epm hap reset". Furthermore, there is an EPM core file from the leaf switch.	12.2(4f)
CSCvg62580	Entries are not pushed to the broadcom hardware level. As a result, VLANs might not get pushed to the interfaces or there might be ARP issues if the port is not forwarding for the VLAN in question.	12.2(4f)
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.	12.2(4f)
CSCvg66700	Communication through a shared L3Out is broken when an additional external network is added to the shared L3Out. This occurs because the prefix list that controls redistribution of the bridge domain static route into EIGRP is deleted. Communication to and from an EIGRP L3Out will be lost due to the route-map being deleted.	12.2(4f)
CSCvg73430	The shared secondary IPv6 address state on an external SVI might become "DUP" (duplicate state) and become unusable as a result of the Duplicate Address Detection (DAD) process on the address provisioning. Use the "show ipv6 interface brief vrf <name>" command to check the state of the address on the external SVI to see if this issue is hit. The shared secondary address is the global IPv6 subnet that is shared across multiple border leaf switches to provide border leaf switch redundancy to externally connected Layer 3 devices in the L3Out policy.	12.2(4f)

Bugs

Bug ID	Description	Fixed In
CSCvg81299	There is an Onboard Failure Logging (OBFL) issue while reloading a switch that is caused by a directory not being free during bootup. As a result, the bootstrap process is not able to unmount the directory.	12.2(4f)
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"	12.2(4f)
CSCvg91161	In a GOLF setup, if there are multiple instP managed objects configured under a GOLF L3Out, deleting one of the instP managed objects clears the nexthop object (bgpPfxLeakCtrlIP). This can also happen if there are multiple GOLF L3Outs for a given VRF instance (context) and one of the L3Outs is deleted. In a GOLF setup on Cisco ACI, when moving a regular L3Out from one GOLF VRF instance to another GOLF VRF instance, the bridge domain subnet might stop being advertised.	12.2(4f)
CSCvh09693	The DHCP offer that goes back to the client has its source MAC address set to 0.	12.2(4f)
CSCvh21652	A leaf switch reloads unexpectedly and generates a core file for the device_test process.	12.2(4f)
CSCvh28182	When using microsegmentation in Cisco ACI, if 2 IP addresses are configured in an IP filter, both IP addresses can be detected. However, when using a MAC filter, only 1 MAC address can be detected even if 2 MAC addresses are configured in the filter.	12.2(4f)
CSCvh29461	BGP between spine switches in different pods in a multipod environment goes down and is unable to be re-established. After upgrading to a Cisco APIC 3.1 release, BGP between spine switches in different pods goes down if a QoS CoS translation policy is enabled.	12.2(4f)
CSCvh32225	An iBGP session goes down when iBGP breaks in stretched fabric if QoS Dot1p is enabled on Cisco -EX and -FX switches.	12.2(4f)
CSCvh51474	A leaf switch reloads while collecting tech support files, and there is a core file from the unicast Routing Information Base (RIB) process.	12.2(4f)
CSCvh60441	A crash occurs in vntag_mgr in a scale environment.	12.2(4f)
CSCvh64426	The RJ45 link does not come up when connecting to an N9K-C93108TC-EX switch at 100 Mbps.	12.2(4f)
CSCvh68419	Layer 2 switched MPLS traffic is dropped in the Cisco ACI fabric.	12.2(4f)
CSCvh70146	Two routers are not able to form eBGP neighborhood with ttl=1. The routers hit different sup-TCAM rules and become flooded in the EPG.	12.2(4f)
CSCvh74146	A leaf switch resets due to an acllog process crash.	12.2(4f)
CSCvh76066	Cisco ACI leaf switches hit the policy_mgr high availability policy reset.	12.2(4f)

Bugs

Bug ID	Description	Fixed In
CSCvh83687	There is a memory leak when a static route add is followed by a static route modify, because the same BFD entry gets installed multiple times for that BFD session.	12.2(4f)
CSCvh89176	The loss of prefix list entries for a shared L3Out causes the border leaf switch to stop advertising its subnets after an unrelated configuration change.	12.2(4f)
CSCvh91642	If the VLAN scale limit is hit, the VLAN translation table will not be cleaned up even after removing the VLANs that exceed the scale limit. This can introduce unexpected connectivity issues for the VLANs that do not have a translation entry.	12.2(4f)
CSCvh99171	The Policy-Mgr NX-OS process dumps a core when too many management VRF instance prefixes are added and deleted.	12.2(4f)
CSCvi05426	Timestamps are 0 in Cisco Tetration Analytics exports from -EX switches in the ACI mode.	12.2(4f)
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.	12.2(4f)
CSCvi42886	After upgrading several leaf switches, using the show module command on each switch shows some as having the " pending" status, while others have the " pass" status.	12.2(4f)

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 12.2(4) releases in which the known behavior exists. A bug might also exist in releases other than the 12.2(4) releases.

Table 6 Known Behaviors in the 12.2(4f) 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.	12.2(4f) 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.	12.2(4f) and later
CSCup65586	The show interface command shows the tunnel's Rx/Tx counters as 0.	12.2(4f) 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.	12.2(4f) and later

Bugs

Bug ID	Description	Exists In
CSCup92534	Continuous "threshold exceeded" messages are generated from the fabric.	12.2(4f) and later
CSCuq39829	Switch rescue user ("admin") can log into fabric switches even when TACACS is selected as the default login realm.	12.2(4f) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) and later
CSCuq92447	When modifying the L2Unknown Unicast parameter on a Bridge Domain (BD), interfaces on externally connected devices may bounce. Additionally, the endpoint cache for the BD is flushed and all endpoints will have to be re-learned.	12.2(4f) 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.	12.2(4f) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	12.2(4f) and later
CSCur81822	The access-port operational status is always "trunk" .	12.2(4f) 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.	12.2(4f) and later
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	12.2(4f) and later

Bugs

Bug ID	Description	Exists In
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.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) and later
CSCut59020	If Backbone and NSSA areas are on the same leaf, and default route leak is enabled, Type-5 LSAs cannot be redistributed to the Backbone area.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) and later

Bugs

Bug ID	Description	Exists In
CSCu57316	TEP counters from the border leaf to remote leaf nodes do not increment.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	12.2(4f) and later
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) 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.	12.2(4f) and later
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to "0", and so traffic will still flow.	12.2(4f) and later
CSCuy96912	The debounce timer is not supported on 25G links.	12.2(4f) and later

Bugs

Bug ID	Description	Exists In
CSCuz12913	An ACI leaf switch sends ARP to a device (such as a router or host) that belongs to directly connected subnets for an L3Out. After ARP is resolved, devices in directly connected subnets on two different L3Outs can talk each other without any contracts.	12.2(4f) 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.	12.2(4f) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrIRule will not increment.	12.2(4f) 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.	12.2(4f) and later
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	12.2(4f) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	12.2(4f) 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.	12.2(4f) and later
CSCvb36823	If you have only one spine switch that is part of the infra WAN and you reload that switch, there can be drops in traffic. You should deploy the infra WAN on more than one spine switch to avoid this issue.	12.2(4f) and later
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	12.2(4f) 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.	12.2(4f) and later
CSCvd11146	Bridge domain subnet routes advertised out of the Cisco ACI fabric through an OSPF L3Out can be relearned in another node belonging to another OSPF L3Out on a different area.	12.2(4f) and later
CSCvd63567	After upgrading a switch, Layer 2 multicast traffic flowing across PODs gets affected for some of the bridge domain Global IP Outsides.	12.2(4f) and later

Bugs

- 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 outsiders (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.

Bugs

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.

© 2017–2021 Cisco Systems, Inc. All rights reserved.