Cisco Nexus 9000 ACI-Mode Switches Release Notes, Release 13.2(10)

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Introduction

The Cisco NX-OS software for the Cisco Nexus 9000 series switches is a data center, purpose-built operating system designed with performance, resiliency, scalability, manageability, and programmability at its foundation. It provides a robust and comprehensive feature set that meets the requirements of virtualization and automation in data centers.

This release works only on Cisco Nexus 9000 Series switches in ACI mode.

This document describes the features, issues, and limitations for the Cisco NX-OS software. For the features, issues, and limitations for the Cisco Application Policy Infrastructure Controller (APIC), see the Cisco Application Policy Infrastructure Controller Release Notes, Release 3.2(10).

For more information about this product, see " Related Content."

Date	Description
July 8, 2022	Release 13.2(10g) became available. Added a resolved bug for this release.
May 16, 2022	In the Open Issues section, added bug CSCwa47686.
August 10, 2021	In the Open Issues section, added bug CSCvy30381.
August 6, 2021	Release 13.2(10f) became available. There are no changes to this document for this release. See the Cisco Application Policy Infrastructure Controller Release Notes, Release 3.2(10) for the changes in this release.
July 6, 2021	In the Supported Hardware section, added the NXA-PAC-500W-PI and NXA-PAC-500W-PE PSUs.
June 24, 2021	In the Open Issues section, added bug CSCvu07844.
March 3, 2021	Added bug CSCvx19640 as a resolved issue.
February 27, 2021	Release 13.2(10e) became available.

Supported Hardware

Table 1. Modular Spine Switches

Product ID	Description
N9K-C9504	Cisco Nexus 9504 switch chassis
N9K-C9508	Cisco Nexus 9508 switch chassis
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
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
N9K-C9516	Cisco Nexus 9516 switch chassis

Table 2. Modular Spine Switch Line Cards

© 2024 Cisco and/or its affiliates. All rights reserved.

Product ID	Description	Maximum Quantity		
		Cisco Nexus 9504	Cisco Nexus 9508	Cisco Nexus 9516
N9K-X9736C-FX	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet Cloud Scale line card	4	8	16
N9K-X9732C-EX	Cisco Nexus 9500 32-port, 40/100 Gigabit Ethernet Cloud Scale line card Note: The N9K-X9732C-EX line card cannot be used when a fabric module is installed in FM slot 25.	4	8	16
N9K-X9736PQ	Cisco Nexus 9500 36-port 40 Gigabit Ethernet line card	4	8	16

Table 3. Modular Spine Switch Fabric Modules

Product ID	Description	Minimum	Maximum
N9K-C9504-FM-E	Cisco Nexus 9504 cloud scale fabric module	4	5
N9K-C9508-FM-E	Cisco Nexus 9508 cloud scale fabric module	4	5
N9K-C9508-FM-E2	Cisco Nexus 9508 cloud scale fabric module	4	5
N9K-C9516-FM-E2	Cisco Nexus 9516 cloud scale fabric module	4	5
N9K-C9504-FM	Cisco Nexus 9504 classic fabric module Note: This fabric module is not supported in slot 21 nor 25.	3	4
N9K-C9508-FM	Cisco Nexus 9508 classic fabric module Note: This fabric module is not supported in slot 21 nor 25.	3	4
N9K-C9516-FM	Cisco Nexus 9516 classic fabric module Note: This fabric module is not supported in slot 21 nor 25.	3	4

Table 4. Modular Spine Switch Supervisor and System Controller Modules

Product ID	Description
N9K-SUP-A+	Cisco Nexus 9500 Series supervisor module
N9K-SUP-B+	Cisco Nexus 9500 Series supervisor module
N9K-SUP-A	Cisco Nexus 9500 Series supervisor module
N9K-SUP-B	Cisco Nexus 9500 Series supervisor module
N9K-SC-A	Cisco Nexus 9500 Series system controller

Table 5. Fixed Spine Switches

© 2024 Cisco and/or its affiliates. All rights reserved.

Product ID	Description
N9K-C9336PQ	Cisco Nexus 9336PQ switch, 36-port 40 Gigabit Ethernet QSFP
N9K-C9364C	Cisco Nexus 9300 platform switch with 64 40/100-Gigabit QSFP28 ports and two 1/10-Gigabit SFP+ ports. The last 16 of the QSFP28 ports are colored green to indicate that they support wire-rate MACsec encryption.

Table 6. Fixed Spine Switch Power Supply Units

Product ID	Description
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
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
NXA-PAC-1200W-PE	1200W AC power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance
NXA-PAC-1200W-PI	1200W AC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance
NXA-PAC-1100W-PE2	1100W AC power supply, port side exhaust pluggable
NXA-PAC-1100W-PI2	1100W AC power supply, port side intake pluggable
NXA-PDC-1100W-PE	1100W AC power supply, port side exhaust pluggable
NXA-PDC-1100W-PI	1100W AC power supply, port side intake pluggable
NXA-PHV-1100W-PE	1100W HVAC/HVDC power supply, port-side exhaust
NXA-PHV-1100W-PI	1100W HVAC/HVDC power supply, port-side intake
NXA-PDC-930W-PE	930W AC power supply, port side exhaust pluggable
NXA-PDC-930W-PI	930W AC power supply, port side intake pluggable
N9K-PUV-1200W	1200W HVAC/HVDC dual-direction airflow power supply

Table 7. Fixed Spine Switch Fans

Product ID	Description
N9K-C9300-FAN3	Burgundy port side intake fan
N9K-C9300-FAN3-B	Blue port side exhaust fan
N9K-C9508-FAN	Fan tray for Cisco Nexus 9508 chassis

Table 8. Fixed Leaf Switches

© 2024 Cisco and/or its affiliates. All rights reserved.

Product ID	Description
N9K-C9336C-FX2	Cisco Nexus C9336C-FX2 Top-of-rack (ToR) switch with 36 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. Note: 1-Gigabit QSA is not supported on ports 1/1-6 and 1/33-36. The port profile feature supports downlink conversion of ports 31 through 34. Ports 35 and 36 can only be used as uplinks.
N9K-C93108TC-FX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. 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.
N9K-C93180YC-FX	Cisco Nexus 9300 platform switch with 48 1/10/25-Gigabit Ethernet SFP28 front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports. The SFP28 ports support 1-, 10-, and 25-Gigabit Ethernet connections and 8-, 16-, and 32-Gigabit Fibre Channel connections. 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.
N9K-C9348GC-FXP	Cisco Nexus 9348GC-FXP switch with 48 100/1000-Megabit 1GBASE-T downlink ports, 4 10-/25-Gigabit SFP28 downlink ports, and 2 40-/100-Gigabit QSFP28 uplink ports.
N9K-C93108TC-EX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) front panel ports and 6 40/100-Gigabit QSFP28 spine facing ports.
N9K-C93180LC-EX	Cisco Nexus 9300 platform switch with 24 40-Gigabit front panel ports and 6 40/100-Gigabit QSFP28 spine-facing ports. The switch can be used as either a 24 40G port switch or a 12 100G port switch. If 100G is connected the Port1, Port 2 will be HW disabled.
N9K-C93180YC-EX	Cisco Nexus 9300 platform switch with 48 1/10/25-Gigabit front panel ports and 6-port 40/100 Gigabit QSFP28 spine-facing ports.
N9K-C9372PX-E	Cisco Nexus 9372PX-E Top-of-rack (ToR) Layer 3 switch with 48 Port 1/10-Gigabit APIC-facing ports Ethernet SFP+ front panel ports and 6 40-Gbps Ethernet QSFP+ spine-facing ports Note: Only the downlink ports 1-16 and 33-48 are capable of supporting SFP1-10G-ZR SFP+.
N9K-C9372TX-E	Cisco Nexus 9372TX-E Top-of-rack (ToR) Layer 3 switch with 48 10GBASE-T (copper) front panel ports and 6 40-Gbps Ethernet QSFP+ spine-facing ports
N9K-C93120TX	Cisco Nexus 9300 platform switch with 96 1/10GBASE-T (copper) front panel ports and 6-port 40-Gigabit Ethernet QSFP spine-facing ports.
N9K-C93128TX	Cisco Nexus 9300 platform switch with 96 1/10GBASE-T (copper) front panel ports and 6 or 8 40-Gigabit Ethernet QSFP spine-facing ports.
N9K-C9332PQ	Cisco Nexus 9332PQ Top-of-rack (ToR) Layer 3 switch with 26 APIC-facing ports and 6 fixed-Gigabit spine facing ports.
N9K-C9372PX	Cisco Nexus 9372PX Top-of-rack (ToR) Layer 3 switch with 48 Port 1/10-Gigabit APIC-facing ports Ethernet SFP+ front panel ports and 6 40-Gbps Ethernet QSFP+ spine-facing ports Note: Only the downlink ports 1-16 and 33-48 are capable of supporting SFP1-10G-ZR

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	Description
	SFP+.
N9K-C9372TX	Cisco Nexus 9372TX Top-of-rack (ToR) Layer 3 switch with 48 1/10GBASE-T (copper) front panel ports and 6 40-Gbps Ethernet QSFP spine-facing ports
N9K-C9396PX	Cisco Nexus 9300 platform switch with 48 1/10-Gigabit SFP+ front panel ports and 6 or 12 40-Gigabit Ethernet QSFP spine-facing ports
N9K-C9396TX	Cisco Nexus 9300 platform switch with 48 1/10GBASE-T (copper) front panel ports and 6 or 12 40-Gigabit Ethernet QSFP spine-facing ports

Table 9. Expansion Modules

Product ID	Description
N9K-M12PQ	12-port or 8-port Gigabit Ethernet expansion module
N9K-M6PQ	6-port Gigabit Ethernet expansion module
N9K-M6PQ-E	6-port, 40 Gigabit Ethernet expansion module

Table 10. Fixed Leaf Switch Power Supply Units

Product ID	Description
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
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
N9k-PAC-3000W-B	3000W AC power supply, port side intake
N9K-PAC-650W	650W AC power supply, port side intake pluggable
N9K-PAC-650W-B	650W AC power supply, port side exhaust pluggable
NXA-PAC-1200W-PE	1200W AC power supply, port side exhaust pluggable, with higher fan speeds for NEBS compliance
NXA-PAC-1200W-PI	1200W AC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance
NXA-PAC-1100W-PE2	1100W AC power supply, port side exhaust pluggable
NXA-PAC-1100W-PI2	1100W AC power supply, port side intake pluggable
NXA-PDC-1100W-PE	1100W AC power supply, port side exhaust pluggable
NXA-PDC-1100W-PI	1100W AC power supply, port side intake pluggable
NXA-PHV-1100W-PE	1100W HVAC/HVDC power supply, port-side exhaust

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	Description
NXA-PHV-1100W-PI	1100W HVAC/HVDC power supply, port-side intake
NXA-PAC-500W-PE	500W AC Power supply, port side exhaust pluggable
NXA-PAC-500W-PI	500W AC Power supply, port side intake pluggable
NXA-PAC-350W-PE	350W AC power supply, port side exhaust pluggable
NXA-PAC-350W-PI	350W AC power supply, port side intake pluggable
NXA-PDC-930W-PE	930W AC power supply, port side exhaust pluggable
NXA-PDC-930W-PI	930W AC power supply, port side intake pluggable
NXA-PDC-440W-PI	440W DC power supply, port side intake pluggable, with higher fan speeds for NEBS compliance Note: This power supply is supported only by the Cisco Nexus 9348GC-FXP ACI-mode switch.
N9K-PUV-1200W	1200W HVAC/HVDC dual-direction airflow power supply
N9K-PUV-3000W-B	3000W AC power supply, port side exhaust pluggable
UCSC-PSU-930WDC V01	Port side exhaust DC power supply compatible with all ToR leaf switches
UCS-PSU-6332-DC	930W DC power supply, reversed airflow (port side exhaust)

Table 11. Fixed Leaf Switch Fans

Product ID	Description
N9K-C9300-FAN3	Burgundy port side intake fan
N9K-C9300-FAN3-B	Blue port side exhaust fan
NXA-FAN-30CFM-B	Burgundy port side intake fan
NXA-FAN-30CFM-F	Blue port side exhaust fan
NXA-FAN-65CFM-PE	Blue port side exhaust fan
NXA-SFAN-65CFM-PE	Blue port side exhaust fan
NXA-FAN-65CFM-PI	Burgundy port side intake fan
NXA-SFAN-65CFM-PI	Burgundy port side intake fan

Supported FEX Models

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

© 2024 Cisco and/or its affiliates. All rights reserved.

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

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

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

New Hardware Features

There are no new hardware features in this release.

New Software Features

For new software features, see the Cisco Application Policy Infrastructure Controller Release Notes, Release 3.2(10).

Changes in Behavior

For the changes in behavior, see the Cisco ACI Releases Changes in Behavior document.

Open Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug. The " Exists In" column of the table specifies the 13.2(10) releases in which the bug exists. A bug might also exist in releases other than the 13.2(10) releases.

Bug ID	Description	Exists in
CSCvg85886	When an ARP request is generated from one endpoint to another endpoint in an isolated EPG, an ARP glean request is generated for the first endpoint.	13.2(10e) and later
CSCvg95192	Endpoint information is missing in the spine switches.	13.2(10e) and later
CSCvh11299	In COOP, the MAC IP address route has the wrong VNID, and endpoints are missing from the IP address DB of COOP.	13.2(10e) and later
CSCvh18100	If Cisco ACI Virtual Edge or AVS is operating in VxLAN non-switching mode behind a FEX, the traffic across the intra-EPG endpoints will fail when the bridge domain has ARP flooding enabled.	13.2(10e) and later
CSCvj03533	When IPv6 packets are received, mab is triggered. But, only the MAC address endpoint is learned, not the IP address endpoint.	13.2(10e) and later
CSCvj23046	In Cisco ACI Multi-Site plus multi-pod topologies, there could be multicast traffic loss for about 30 seconds on the remote-site. If only one LC has fabric links, there are other LCs with no fabric links and the LC with fabric links is reloaded.	13.2(10e) and later
CSCvj29908	Traffic gets dropped when a new TX SA is programmed after an old Rx SA is deleted on the peer and there are breakout ports in the link down state.	13.2(10e) and later
CSCvk48856	The port LED shows green when a few breakout ports lanes are down.	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Exists in
CSCvk73228	This is an enhancement to decode the binary logs offline directly from the techsupport.	13.2(10e) and later
CSCvk74561	Link down detection on the copper transceiver port takes around 1 second of time when its peer switch reloads. This issue is only with a copper transceiver.	13.2(10e) and later
CSCvm75395	A route map is deployed even when the route profile is configured incorrectly. When upgrading to a release that includes the fixed for this defect, the incorrectly deployed route map is removed from the leaf switch, which may affect traffic that was using the route map.	13.2(10e) and later
CSCvn28108	A switch gets stuck in a bootloop with the following error raised on the console: [1041.090380] obfl_klm writing reset reason 58, LC insertion sequence failure => [Failures < MAX] : powercycle [1042.207780] write_mtd_flash_panic: successfully wrote 88 bytes at address 0xd68 to RR Iter: 0.	13.2(10e) and later
CSCvo35006	Multiple N9K-X9736C-FX 40G line cards get stuck in the 'Inserted' state during a reload or reboot.	13.2(10e) and later
CSCvo39715	When downgrading a Cisco ACI fabric, the OSPF neighbors go down after downgrading the Cisco APICs from a 3.2 or later release to a pre-3.2 release. After the upgrade, the switches are still running a 13.2 or later release.	13.2(10e) and later
CSCvo49717	After downgrading to the 13.2(9) release from a later release, an N9K-C9508-FM-E2 fabric module gets stuck in the 'Inserted' state.	13.2(10e) and later
CSCvp09949	Copy service traffic will fail to reach the TEP where the copy devices are connected. Traffic will not be seen on the spine switches.	13.2(10e) and later
CSCvp47696	A port-client crash is seen on FX2 leaf switches when a large number of breakouts are configured.	13.2(10e) and later
CSCvp50075	A leaf switch experiences an unexpected reload due to a HAP reset.	13.2(10e) and later
CSCvp72312	A contract that is provided by an EPG using a bridge domain with subnet X and that is consumed by an L3Out EPG causes a leak of subnet X from VRF B to VRF A. The existing non-pervasive static route in VRF A is replaced by a pervasive route in pointing to spine switch V4 proxy. After the contract leaking subnet A is removed, the pervasive static route persists.	13.2(10e) and later
CSCvp77034	The hardware abstraction layer (HAL) generates a core file when all of the non-fabric ports are converted into breakout ports.	13.2(10e) and later
CSCvq20215	Some ECMP paths may be flap between "multipath" and "non-multipath." For example, if the configured EBGP MAX ECMP number is 10 and there are 16 BGP ECMP paths for a prefix in the BGP routing table, then 5 paths change between multipath and non-multipath whenever the BGP bestpath calculation is run.	13.2(10e) and later
CSCvq64803	A leaf switch crashes with the "Unknown" reset reason when the breakout ports configuration is re-applied. The reset reason for this switch is as follows: Image Version : 13.2(3o) Reset Reason (LCM): Unknown (0) at time Fri Jul 12 14:21:14 2019	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
	Reset Reason (SW): Reset triggered due to HA policy of Reset (16) at time Fri Jul 12 14:17:40 2019 Service (Additional Info): Reset triggered due to HA policy of Reset	
CSCvq98750	In Cisco ACI when using MAC pinning with a vPC, prior to reloading when you run the 'show vpc brief' command on the CLI, the command shows that the vPC is passing consistency checks. However, after reloading the leaf switch, the vPC then properly displays the consistency check as 'Not Applicable'.	13.2(10e) and later
CSCvr09108	An interface does not come up when a new link is connected. However, from the DOM data, the signals are present.	13.2(10e) and later
CSCvr46867	A Cisco ACI modular spine switch (N9504 chassis) with redundant supervisor modules (N9K-SUP-A) had an unexpected series of switchovers during a 6 minute period.	13.2(10e) and later
CSCvr47042	After removing a transceiver or cable from the interface, the port LED remains green. A port is physically down, but the " show interface" command says that the port is still up.	13.2(10e) and later
CSCvr49904	Traffic with a UDP destination port of 8472 is dropped on ingress by the ACI fabric.	13.2(10e) and later
CSCvr79911	An LLDP/CDP MAC address entry gets stuck in the blade switch table on a leaf switch in a vPC. The entry can get stuck if the MAC address flaps and hits the move detection interval, which stops all learning for the address. Use the following command to verify if a switch has a stale MAC address entry: module-1# show system internal epmc bladeswitch_mac all	13.2(10e) and later
CSCvr83337	A Cisco ACI leaf switch unexpectedly reloads and generates a core file.	13.2(10e) and later
CSCvr88009	The Netflow (nfm) process crashes during configuration changes.	13.2(10e) and later
CSCvr91674	Whenever a switch hits a burst of PCIe, DRAM, or MCE errors, sometimes the device_test process crashes, which can cause the switch to reload.	13.2(10e) and later
CSCvr98827	Some of the control plane packets are incorrectly classified as the user class and are reported as dropped in single chip spine switches. The statistics are incorrect because the packets are not actually dropped.	13.2(10e) and later
CSCvs02955	When running " show system internal epm endpoint all summary" on an FX leaf, the command output is cut short.	13.2(10e) and later
CSCvs08304	The spine outerdstip, which indicates that the egress TEP is connecting to the Tetration network, is not updated when an egress L3Out in the mgmt:inb VRF fails over to a redundant L3Out on another leaf switch.	13.2(10e) and later
CSCvs18150	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 fw 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.	13.2(10e) and later
CSCvs34065	The " get_bkout_cfg failed" error displays when the following vsh_lc cli command is	13.2(10e) and

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
	executed: vsh_lc -c " show system internal port-client event-history all"	later
CSCvs40299	The policy_mgr process on an ACI leaf switch has a memory leak and results in an unexpected reload. The problem can happen over a long period of time, such as a year. Depending on when individual switches were last rebooted, multiple devices could experience the reload at around the same time.	13.2(10e) and later
CSCvs41818	Port 1/2 on N9k-C9364C flaps continuously and does not come up.	13.2(10e) and later
CSCvs45414	A N9K-X9736PQ linecard in an ACI mode Nexus 9500 spine switch unexpectedly reloads. The following output is seen in the command " show system reset-reason module 1" : `show system reset-reason module 1` ********** module reset reason (1) ********** 0) At 2019-12-01T00:00:00.00 Reason: line-card-not-responding Service:Line card not responding => [Failures < MAX] : powercycle Version:	13.2(10e) and later
CSCvs49377	After a virtual machine is vMotioned, traffic begins to drop the source from that endpoint. When running " show logging ip access-list internal packet-log deny" on the leaf switch, you can see policy drops for the endpoint.	13.2(10e) and later
CSCvs56978	Connectivity between a server EPG and external L3Out EPG can be broken for some subnets that are configured with an external subnet for an external EPG.	13.2(10e) and later
CSCvs89617	Some ARP packets get dropped across the Cisco ACI fabric.	13.2(10e) and later
CSCvt00231	Traffic destined to a switch is policy dropped. The contracts configured on the switch look correct, but the ELAM drop reason shows a clear SECURITY_GROUP_DENY. If you dump the FPC and FPB pt.index results of the ELAM, the values are different. Specifically, the FPC index is wrong when you check the Stats Idx under the specific ACLQOS rule. FPC should be the summary of the final result. In this case, there are two hits, but there is one stable entry in TCAM and one that is not stable.	13.2(10e) and later
CSCvt08181	All routes to a particular spine switch are removed from uRIB on all leaf switches in the fabric.	13.2(10e) and later
CSCvt25383	The pervasive static route is missing on the spine node.	13.2(10e) and later
CSCvt35002	A link intermittently flaps on leaf switch fabric ports that are connected to a spine switch.	13.2(10e) and later
CSCvt39689	Glean ARP (0xfff2, 239.255.255.240) flood is stopped on the transit leaf switch and is not delivered toward all the leaf switches in the fabric. Thus, silent host discovery does not work.	13.2(10e) and later
CSCvt52620	There is a stale pervasive route after a DHCP relay label is deleted.	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Exists in
CSCvt57119	A Cisco ACI leaf switch sends traffic that is untagged for a particular VLAN even though it is configured as trunk (tagged).	13.2(10e) and later
CSCvt73069	A Cisco ACI fabric is not fully fit after a Cisco APIC firmware upgrade.	13.2(10e) and later
CSCvt94039	A leaf switch crashes and reloads due to " nfm hap reset" .	13.2(10e) 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.	13.2(10e) 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.	13.2(10e) and later
CSCvu08065	If inter-VRF DHCP relay is used, it may be observed that DHCP breaks after performing any activity that causes the client VRF to get removed and re-deployed on the client leaf nodes.	13.2(10e) and later
CSCvu15712	If a spine switch's PTEP is configured as the multipod L3Out router ID and the router ID is later changed, the spine switch's PTEP loopback gets deleted and the MP BGP session goes down.	13.2(10e) and later
CSCvu15751	The following event can be seen on the spine node: [E4204936][transition][warning][sys] %URIB-4-SYSLOG_SL_MSG_WARNING: URIB-5-RPATH_DELETE: message repeated 1 times in last 220162 sec	13.2(10e) and later
CSCvu22736	There is an event in which the syslog message is masked and does not provide details about the issue. The main syslog message is not seen, but rate-throttled syslog messages are seen.	13.2(10e) and later
CSCvu26947	If a rogue file grows too large, it can cause out of memory condition on a spine switch or leaf switch line card or fabric module without proactively alerting the user to the memory leak, and the line card or fabric module will reload.	13.2(10e) and later
CSCvu40050	The spine node KIC database is missing the v4 default route from RIB. This causes in-band return traffic to drop on the way back to the border leaf nodes.	13.2(10e) and later
CSCvu48811	When a Cisco ACI switch is configured in a " maintenance mode" (mmode), a banner is displayed to the user indicating the operating mode of the switch.	13.2(10e) and later
CSCvu61024	Zoning-rules are not programmed in the hardware after reloading a switch.	13.2(10e) and later
CSCvu72416	Triggered by a physical layer issue, such as fiber or a bad transceiver, a link flap may happen every now and then. However, it is uncommon to have continuous flaps when the node is left unattended over an extended period, such as having 688,000 flaps over a year. Each time after the fabric link flaps, one dbgRemotePort managed object is added to the policyElement database. After a long time flapping like this, unexpected memory allocation and access can be triggered for the Nexus OS process, such as policy_mgr or ethpm. This defect is to enhance the object-store to reduce the impact for such scenarios.	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Exists in
CSCvu84587	VTEP endpoints are learned and set to bounce on some leaf switches. A single VTEP IP address could be seen as local on one vPC pair, but as an IP XR with bounce on another leaf switch pair.	13.2(10e) and later
CSCvv27817	DHCP unicast renewal ACKs are NOT forwarded across the fabric to clients. This traffic is sourced from port 67 destined to port 68. The regular Discover, Offer, Request, Acknowledge (DORA) process and unicast ACKs function correctly. This traffic is sourced from port 67 destined to port 67. The DHCP renewals are incorrectly being punted to the CPU as ISTACK_SUP_CODE_DHCP_SNOOP on the ingress leaf switch.	13.2(10e) and later
CSCvv33100	The IPS port is not down when an RX cable is removed on a Cisco ACI leaf switch 1G port. An ACI switch with 1G fiber would signal a peer IOS device, such as a Catalyst 6000 series switch, with flow control auto/desired to turn on the flow control.	13.2(10e) and later
CSCvv39277	After an upgrade, for one of the VRF tables, the BGP route map is missing on the spine switch, which results in bridge domain prefixes not being advertised.	13.2(10e) and later
CSCvv75224	IPv6 BGP route with recursive next-hop is programmed in the software, but not programmed in the hardware. Traffic destined to this route is blackholed.	13.2(10e) and later
CSCvv78885	A stale route map entry is causes unexpected route leaking.	13.2(10e) and later
CSCvv95800	A spine switch reloads unexpectedly due to the service on the linecard having a hap-reset.	13.2(10e) and later
CSCvw07282	On a modular spine switch, an unconnected port's switching state is disabled, which means it is out of service. The issue is that after reloading a line card, all of the ports on that line card change to switching state enabled, even if the port is not connected to anything. This issue is mostly cosmetic; there is no real impact if an unconnected port has switching state enabled.	13.2(10e) and later
CSCvy30381	After replacing the hardware for a leaf switch, the leaf switch front-panel ports are set to the admin-down state for 45 minutes.	13.2(10e) and later
CSCwa12763	External route import for a VRF instance fails on a leaf switch after removing a shared services contract between two EPGs.	13.2(10e) and later
CSCwa47686	For a Cisco ACI fabric with more than 128 leaf switches in a given pod, such as 210 leaf switches in a single pod deployment, after enabling PTP globally, only 128 leaf switches are able to enable PTP. The remaining 82 leaf switches fail to enable PTP due to the error F2728 latency-enable-failed.	13.2(10e) and later
CSCwb08081	A route profile that matches on community list and sets the local pref and community is not working post upgrade to 5.2.x release. route-map imp-l3out-L3OUT_WAN-peer-2359297, permit, sequence 4201 Match clauses: community (community-list filter): peer16389-2359297-exc-ext-in-L3OUT_WAN_COMMUNITY-rgcom Set clauses: local-preference 200 community xxxxx:101 xxxxx:500 xxxxx:601 xxxxy:4 additive	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
	The match clause works as expected, but the set clause is ignored.	
CSCwd29346	An ACI switch's console may continuously output messages similar to: svc_ifc_eventmg (*****) Ran 7911 msecs in last 7924 msecs	13.2(10e) and later
CSCvz89270	CRC/Output errors are observed in the Cisco ACI fabric or end devices regarding performance issues with jumbo MTU in use. There are no Rx FCS errors on spine switch ports, although Tx Frm or output errors are observed. In some cases, there may be Rx FCS errors due to some other problem, although they will not match the output/Tx Frm errors. Leaf switches have Rx CRC/stomp errors on fabric links from the affected spine switches and will forward such stomped CRCs out of the access ports. End devices report CRC/FCS errors and the application side will not receive these packets, leading to performance or packet loss	13.2(10e) and 13.2(10f)

Resolved Issues

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.

Bug ID	Description	Fixed in
CSCvz89270	CRC/Output errors are observed in the Cisco ACI fabric or end devices regarding performance issues with jumbo MTU in use. There are no Rx FCS errors on spine switch ports, although Tx Frm or output errors are observed. In some cases, there may be Rx FCS errors due to some other problem, although they will not match the output/Tx Frm errors. Leaf switches have Rx CRC/stomp errors on fabric links from the affected spine switches and will forward such stomped CRCs out of the access ports. End devices report CRC/FCS errors and the application side will not receive these packets, leading to performance or packet loss	13.2(10g)
CSCvq82478	Cisco APICs cannot ping each other through inband management. LLDP neighbors on the Cisco APIC sees the directly-connected leaf switch as a neighbor. The 'show lldp neighbor' command on the leaf switch sees the directly-attached Cisco APIC as a neighbor. Pinging the infra bridge domain gateway 10.0.0.30 from the Cisco APIC does not work.	13.2(10e)
CSCvr75413	After upgrading a leaf switch, the switch brings up the front panel ports before the policies are programmed. This may cause a connectivity issue if a connected host relies on the link level state to decide whether or not it can forward traffic on a particular NIC or port. The loss duration would be proportional to the scale of configuration policies that must be programmed.	13.2(10e)
CSCvt07621	The fabric module on a modular spine switch crashes with sdkhal-hap-reset and generates an sdkhal core file.	13.2(10e)
CSCvv13722	A leaf switch crashes due to a routing loop in the IPFIB process.	13.2(10e)
CSCvx19640	Nexus 9000 switch /bootflash goes into read-only mode with Micron M500IT SSD drive. Fault 1222 is raised for ssd-acc diagnostic test with reason " Failed to write to file" .	13.2(10e)

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Known Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug. The " Exists In" column of the table specifies the 13.2(10) releases in which the bug exists. A bug might also exist in releases other than the 13.2(10) releases.

Bug ID	Description	Exists in
CSCuo37016	When configuring the output span on a FEX Hif interface, all the layer 3 switched packets going out of that FEX Hif interface are not spanned. Only layer 2 switched packets going out of that FEX Hif are spanned.	13.2(10e) and later
CSCuo50533	When output span is enabled on a port where the filter is VLAN, multicast traffic in the VLAN that goes out of that port is not spanned.	13.2(10e) and later
CSCup65586	The **show interface** command shows the tunnel's Rx/Tx counters as 0.	13.2(10e) 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.	13.2(10e) and later
CSCup92534	Continuous " threshold exceeded" messages are generated from the fabric.	13.2(10e) and later
CSCuq39829	Switch rescue user (" admin") can log into fabric switches even when TACACS is selected as the default login realm.	13.2(10e) and later
CSCuq46369	An extra 4 bytes is added to the untagged packet with Egress local and remote SPAN.	13.2(10e) and later
CSCuq77095	When the command **show ip ospf vrf <vrf_name>** is run from bash on the border leaf, the checksum field in the output always shows a zero value.	13.2(10e) and later
CSCuq83910	When an IP address moves from one MAC behind one ToR to another MAC behind another ToR, even though the VM sends a GARP packet, in ARP unicast mode, this GARP packet is not flooded. As a result, any other host with the original MAC to IP binding sending an L2 packet will send to the original ToR where the IP was in the beginning (based on MAC lookup), and the packet will be sent out on the old port (location). Without flooding the GARP packet in the network, all hosts will not update the MAC-to-IP binding.	13.2(10e) and later
CSCuq92447	When modifying the L2Unknown Unicast parameter on a Bridge Domain (BD), interfaces on externally connected devices may bounce. Additionally, the endpoint cache for the BD is flushed and all endpoints will have to be re-learned.	13.2(10e) and later
CSCuq93389	If an endpoint has multiple IPs, the endpoint will not be aged until all IPs go silent. If one of the IP addresses is reassigned to another server/host, the fabric detects it as an IP address move and forwarding will work as expected.	13.2(10e) and later
CSCur01336	The power supply will not be detected after performing a PSU online insertion and removal (OIR).	13.2(10e) and later
CSCur81822	The access-port operational status is always " trunk" .	13.2(10e) and later
CSCus18541	An MSTP topology change notification (TCN) on a flood domain (FD) VLAN may not flush endpoints learned as remote where the FD is not deployed.	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description	Exists in
CSCus29623	The transceiver type for some Cisco AOC (active optical) cables is displayed as ACU (active copper).	13.2(10e) and later
CSCus43167	Any TCAM that is full, or nearly full, will raise the usage threshold fault. Because the faults for all TCAMs on leaf switches are grouped together, the fault will appear even on those with low usage. Workaround: Review the leaf switch scale and reduce the TCAM usage. Contact TAC to isolate further which TCAM is full.	13.2(10e) 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.	13.2(10e) and later
CSCus61748	If the TOR 1RU system is configured with the RED fan (the reverse airflow), the air will flow from front to back. The temperature sensor in the back will be defined as an inlet temperature sensor, and the temperature sensor in the front will be defined as an outlet temperature sensor. If the TOR 1RU system is configured with the BLUE fan (normal airflow), the air will flow from back to front. The temperature sensor in the front will be defined as an inlet temperature sensor, and the temperature sensor in the back will be defined as outlet temperature sensor. From the airflow perspective, the inlet sensor reading should always be less than the outlet sensor reading. However, in the TOR 1RU family, the front panel temperature sensor has some inaccurate readings due to the front panel utilization and configuration, which causes the inlet temperature sensor reading to be very close, equal, or even greater than the outlet temperature reading.	13.2(10e) and later
CSCut59020	If Backbone and NSSA areas are on the same leaf, and default route leak is enabled, Type-5 LSAs cannot be redistributed to the Backbone area.	13.2(10e) and later
CSCuu11347	Traffic from the orphan port to the vPC pair is not recorded against the tunnel stats. Traffic from the vPC pair to the orphan port is recorded against the tunnel stats.	13.2(10e) and later
CSCuu11351	Traffic from the orphan port to the vPC pair is only updated on the destination node, so the traffic count shows as excess.	13.2(10e) and later
CSCuu66310	If a bridge domain "Multi Destination Flood" mode is configured as "Drop", the ISIS PDU from the tenant space will get dropped in the fabric.	13.2(10e) and later
CSCuv57302	Atomic counters on the border leaf do not increment for traffic from an endpoint group going to the Layer 3 out interface.	13.2(10e) and later
CSCuv57315	Atomic counters on the border leaf do not increment for traffic from the Layer 3 out interface to an internal remote endpoint group.	13.2(10e) and later
CSCuv57316	TEP counters from the border leaf to remote leaf nodes do not increment.	13.2(10e) and later
CSCuw09389	For direct server return operations, if the client is behind the Layer 3 out, the server-to-client response will not be forwarded through the fabric.	13.2(10e) and later
CSCux97329	With the common pervasive gateway, only the packet destination to the virtual MAC is being properly Layer 3 forwarded. The packet destination to the bridge domain custom MAC fails to be forwarded. This is causing issues with certain appliances that rely on the incoming packets' source MAC to set the return packet destination MAC.	13.2(10e) and later
CSCuy00084	BCM does not have a stats option for yellow packets/bytes, and so BCM does not show in the switch or APIC GUI stats/observer.	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
CSCuy02543	Bidirectional Forwarding Detection (BFD) echo mode is not supported on IPv6 BFD sessions carrying link-local as the source and destination IP address. BFD echo mode also is not supported on IPv4 BFD sessions over multihop or VPC peer links.	13.2(10e) and later
CSCuy06749	Traffic is dropped between two isolated EPGs.	13.2(10e) and later
CSCuy22288	The iping command's replies get dropped by the QOS ingress policer.	13.2(10e) and later
CSCuy25780	An overlapping or duplicate prefix/subnet could cause the valid prefixes not to be installed because of batching behavior on a switch. This can happen during an upgrade to the 1.2(2) release.	13.2(10e) and later
CSCuy47634	EPG statistics only count total bytes and packets. The breakdown of statistics into multicast/unicast/broadcast is not available on new hardware.	13.2(10e) and later
CSCuy56975	You must configure different router MACs for SVI on each border leaf if L3out is deployed over port-channels/ports with STP and OSPF/OSPFv3/eBGP protocols are used. There is no need to configure different router MACs if you use VPC.	13.2(10e) and later
CSCuy61018	The default minimum bandwidth is used if the BW parameter is set to " 0" , and so traffic will still flow.	13.2(10e) and later
CSCuy96912	The debounce timer is not supported on 25G links.	13.2(10e) and later
CSCuz13529	With the N9K-C93180YC-EX switch, drop packets, such as MTU or storm control drops, are not accounted for in the input rate calculation.	13.2(10e) and later
CSCuz13614	For traffic coming out of an L3out to an internal EPG, stats for the actrlRule will not increment.	13.2(10e) and later
CSCuz13810	When subnet check is enabled, a ToR does not learn IP addresses locally that are outside of the bridge domain subnets. However, the packet itself is not dropped and will be forwarded to the fabric. This will result in such IP addresses getting learned as remote endpoints on other ToRs.	13.2(10e) and later
CSCuz47058	SAN boot over a virtual Port Channel or traditional Port Channel does not work.	13.2(10e) and later
CSCuz65221	A policy-based redirect (PBR) policy to redirect IP traffic also redirects IPv6 neighbor solicitation and neighbor advertisement packets.	13.2(10e) and later
CSCva98767	The front port of the QSA and GLC-T 1G module has a 10 to 15-second delay as it comes up from the insertion process.	13.2(10e) and later
CSCvb36823	If you have only one spine switch that is part of the infra WAN and you reload that switch, there can be drops in traffic. You should deploy the infra WAN on more than one spine switch to avoid this issue.	13.2(10e) and later
CSCvb39965	Slow drain is not supported on FEX Host Interface (HIF) ports.	13.2(10e) and later
CSCvb49451	In the case of endpoints in two different TOR pairs across a spine switch that are trying to communicate, an endpoint does not get relearned after being deleted on the local TOR pair. However, the endpoint still has its entries on the remote TOR pair.	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
CSCvd11146	Bridge domain subnet routes advertised out of the Cisco ACI fabric through an OSPF L3Out can be relearned in another node belonging to another OSPF L3Out on a different area.	13.2(10e) and later
CSCvd63567	After upgrading a switch, Layer 2 multicast traffic flowing across PODs gets affected for some of the bridge domain Global IP Outsides.	13.2(10e) and later
CSCvn94400	There is a traffic blackhole that lasts anywhere from a few seconds to a few mins after a border leaf switch is restored.	13.2(10e) and later
CSCvp04772	During an upgrade on a dual-SUP system, the standby SUP may go into a failed state.	13.2(10e) and later
CSCvq56811	Output packets that are ERSPAN'd still have the PTP header. Wireshark might not be able to decode the packets, and instead shows frames with ethertype 0x8988.	13.2(10e) and later
CSCvq71034	There is a policy drop that occurs with L3Out transit cases.	13.2(10e) and later
CSCvr12912	A switch reloads due to a sysmgr heartbeat failure and sysmgr HAP reset.	13.2(10e) and later
CSCvr61096	In a port group that has ports of mixed speeds, the first port in the port group that has valid optics present and is not in the admin down state is processed. The ports that come up later are brought up if they are using the same speed; otherwise, they are put in the hw-disabled state. For example, if ports 14 and 15 are up and are using the 100G speed, then if ports 13 and 16 are using the 40G speed, these ports will be put in the hw-disabled state. After reloading or upgrading, you might not have the same interfaces in the port group in the UP state and in the hw-disabled state as you did before the reload or upgrade.	13.2(10e) and later
CSCvt53089	If a Cisco UCS fabric interconnect is deployed in the end host mode and is a peer to a Cisco ACI ToR switch, and CDP is enabled without LLDP, Blade switch MAC address move tracking is not feasible because CDP does not advertise the peer's MAC address. The blade switch MAC address entry for the fabric interconnect port MAC addresses is not seen in the output of the " show system internal epmc bladeswitch_mac all" command.	13.2(10e) and later
CSCvv16647	A minor traffic outage is seen with a Cisco APIC downgrade.	13.2(10e) and later
CSCvw20049	A switch allows more storm traffic than the configured storm policer rate.	13.2(10e) and later
N/A	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.	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
N/A	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 DSCP-CoS translation policy in the APIC for the pod that translates queuing class information of the packet into the DSCP value in the outer header of the iVXLAN packet. You can also embed CoS by enabling CoS preservation. For more information, see the **Cisco APIC and QoS** KB article.	13.2(10e) and later
N/A	The following properties within a QoS class under " Global QoS Class policies" should not be changed from their 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)	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	Multicast router functionality is not supported when IGMP queries are received with VxLAN encapsulation.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	The rate of the number of IGMP reports sent to a leaf switch should be limited to 1000 reports per second.	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
N/A	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 switch. 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 switch. • Set " unknown multicast flooding" to " Optimized Flood" on needed bridge domains.	13.2(10e) and later
N/A	Traffic destined to Static Route EP VIPs sourced from N9000 switches (switches with names that end in -EX) might not function properly because proxy route is not programmed.	13.2(10e) and later
N/A	An iVXLAN header of 50 bytes is added for traffic ingressing into the fabric. A bandwidth allowance of (50/50 + ingress_packet_size) needs to be made to prevent oversubscription from happening. If the allowance is not made, oversubscription might happen resulting in buffer drops.	13.2(10e) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported in combination with static endpoint configurations.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	An IP/MAC Ckt endpoint configuration is not supported with dynamic endpoint groups. Only static endpoint groups are supported.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	Dynamic deployment of contracts based on instrImmedcy set to onDemand/lazy not supported; only immediate mode is supported.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	Direct server return is not supported for shared services. Direct server return endpoints cannot be spread around different virtual routing and forwarding (VRF) contexts.	13.2(10e) and later

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Bug ID	Description	Exists in
N/A	Configurations for a virtual IP address can only be /32 or /128 prefix.	13.2(10e) and later
N/A	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.	13.2(10e) and later
N/A	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).	13.2(10e) and later
N/A	Learning through GARP will work only in ARP Flood Mode.	13.2(10e) and later

Compatibility Information

- For the supported optics per device, see the Cisco Optics-to-Device Compatibility Matrix.
- Link level flow control is not supported on ACI-mode switches.
- 100mb optics, such as the GLC-TE, are supported in 100mb speed only on -EX and -FX switches, such as the N9K-C93180YC-EX and N9K-C93180YC-FX, and only on front panel ports 1/1-48. 100mb optics are not supported any other switches. 100mb optics cannot be used on EX or FX leaf switches on port profile converted downlink ports (1/49-52) using QSA.
- This release supports the hardware and software listed on the ACI Ecosystem Compatibility List, and supports the Cisco AVS, Release 5.2(1)SV3(3.10).
- 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 C9332PQ 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.
- The following tables provide compatibility information for specific hardware:

Table 12. Modular Spine Switch Compatibility Information

Product ID	Compatibility Information
N9K-C9336PQ	The Cisco N9K-C9336PQ switch is supported for multipod.

© 2024 Cisco and/or its affiliates. All rights reserved.

Product ID	Compatibility Information
	The N9K-9336PQ switch is not supported for inter-site connectivity with Cisco ACI Multi-Site, but is supported for leaf switch-to-spine switch connectivity within a site. The N9K-9336PQ switch is not supported when multipod and Cisco ACI Multi-Site are deployed together.

Table 13. Modular Spine Switch Line Card Compatibility Information

Product ID	Compatibility Information
N9K-X9736C-FX	1-Gigabit QSA is not supported on ports 1/29-36. This line card supports the ability to add a fifth Fabric Module to the Cisco N9K-C9504 and N9K-C9508 switches. The fifth Fabric Module can only be inserted into slot 25.

Table 14. Modular Spine Switch Line Card Compatibility Information

Product ID	Compatibility Information
N9K-C9348GC-FXP	This switch supports the following PSUs: • NXA-PAC-350W-PI • NXA-PAC-350W-PE • NXA-PAC-1100W-PI • NXA-PAC-1100W-PE The following information applies to this switch: • 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. • This switch does not support the 10G GLC-T optic. • The PSU SPROM is not readable when the PSU is not connected. The model displays as " UNKNOWN" and status of the module displays as " shutdown."
N9K-C93180LC-EX	This switch has the following limitations: • The top and bottom ports must use the same speed. If there is a speed mismatch, the top port takes precedence and bottom port will be error disabled. Both ports both must be used in either the 40 Gbps or 10 Gbps mode. • Ports 26 and 28 are hardware disabled. • This release supports 40 and 100 Gbps for the front panel ports. The uplink ports can be used at the 100 Gbps speed. • Port profiles and breakout ports are not supported on the same port.

Table 15. Fixed Spine Switches Compatibility Information

Product ID	Compatibility Information
N9K-C9364C	You can deploy multipod or Cisco ACI Multi-Site separately (but not together) on the Cisco N9K-9364C switch starting in the 3.1 release. You can deploy multipod and Cisco ACI Multi-Site together on the Cisco N9K-9364C switch starting in the 3.2 release. A 930W-DC PSU (NXA-PDC-930W-PE or NXA-PDC-930W-PI) is supported in redundancy mode if 3.5W QSFP+ modules or passive QSFP cables are used and the system is used in 40C ambient temperature or less; for other optics or a higher ambient temperature, a 930W-DC PSU is supported only with 2 PSUs in non-redundancy mode. 1-Gigabit QSA is not supported on ports 1/49-64.

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	Compatibility Information
	This switch supports the following PSUs: • NXA-PAC-1200W-PE • NXA-PAC-1200W-PI • N9K-PUV-1200W • NXA-PDC-930W-PE • NXA-PDC-930W-PI

Table 16. Fixed Leaf Switches Compatibility Information

Product ID	Compatibility Information
N9K-C93180YC-EX	The following FEC modes are not supported on N9K-C93180YC-EX ports 1 through 48 when running in 25G speed: • cl91-rs-fec • cons16-rs-fec • ieee-rs-fec
N9K-C9336C-FX2	The following information applies to this switch: • On older N9K-C9336C-FX2 switches, auto-negotiation does not work on port eth1/4. You can check whether your switch is older by using the following command: ifav124-leaf5# **cat /sys/kernel/cisco_board_info/hw_change_bits** 0x0 The output of "0x0" indicates an older switch that has this limitation. • You can apply a breakout configuration on ports 1 through 34, which can give up to 136 (34*4) server or downlink ports. • Port profiles and breakouts are not supported on the same port. However, you can apply a port profile to convert a fabric port to a downlink, and then apply a breakout configuration. • If you apply a breakout configuration on 34 ports, you must configure a port profile on the ports first, which requires you to reboot the leaf switch. • If you apply a breakout configuration to a leaf switch for multiple ports at the same time, it can take up to 10 minutes for the hardware of 34 ports to be programmed. The ports remain down until the programming completes. The delay can occur for a new configuration, after a clean reboot, or during switch discovery. • Ports 7 through 32 have a link bring up time of less than 2 seconds with QSFP-100G-LR4 and QSFP-40/100G-SRBD optics. For all other ports, the link up time for these optics is between 5 to 14 seconds. In the following situations, the link bring up time will also be greater than 2 seconds: ○ After reloading the Top-of-Rack (ToR) switch ○ When using port optical insertion and removal (OIR) ○ When performing bulk flaps of ports on the ToR switch
N9K-C9332PQ	To connect the Cisco APIC to the Cisco ACI fabric, you must have a 10G interface on the ACI leaf switch. You cannot connect the APIC directly to the N9332PQ ACI leaf switch.

- The following table provides MACsec and CloudSec compatibility information for specific hardware:

© 2024 Cisco and/or its affiliates. All rights reserved.

Table 17. MACsec and CloudSec Support

Product ID	Hardware Type	MACsec Support	CloudSec Support
N9K-C93108TC-FX	Switch	Yes	No
N9K-C93180YC-FX	Switch	Yes	No
N9K-C93216TC-FX2	Switch	Yes	No
N9K-C93360YC-FX2	Switch	Yes	No
N9K-C9336C-FX2	Switch	Yes	No
N9K-C9348GC-FXP	Switch	Yes, only with 10G+	No
N9K-C9364C	Switch	Yes	Yes, only on the last 16 ports
N9K-X9736C-FX	Line Card	Yes	Yes, only on the last 8 ports

- The following additional MACsec and CloudSec compatibility restrictions apply:
    - MACsec is not supported with 1G speed on Cisco ACI leaf switch.
    - MACsec is supported only on the leaf switch ports where an L3Out is enabled. For example, MACsec between a Cisco ACI leaf switch and any computer host is not supported. Only switch-to-switch mode is supported.
    - When using copper ports, the copper cables must be connected directly the peer device (standalone N9k) in 10G mode.
    - A 10G copper SFP module on the peer is not supported.
    - CloudSec only works with spine switches in Cisco ACI and only works between sites managed by Cisco ACI Multi-Site.
    - For CloudSec to work properly, all of the spine switch links that participate in Cisco ACI Multi-Site must have MACsec/CloudSec support.

Usage Guidelines

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

  Note: See the Cisco Application Policy Infrastructure Controller Release Notes, Release 3.2(10) for policy information.

    - 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

© 2024 Cisco and/or its affiliates. All rights reserved.

    - OSPF
    - UDP DstPort 67: BOOTP/DHCP
    - UDP DstPort 68: BOOTP/DHCP
    - IGMP
    - PIM
    - UDP SrcPort 53: DNS replies
    - TCP SrcPort 25: SMTP replies
    - TCP DstPort 443: HTTPS
    - UDP SrcPort 123: NTP
    - UDP DstPort 123: NTP
- Leaf switches and spine switches typically have memory utilization of approximately 70% to 75%, even in a new deployment where no configuration has been pushed. This amount of memory utilization is due to the Cisco ACI-specific processes, which take up more memory compared to a standalone Nexus deployment. The memory utilization is not a problem unless it exceeds 90%. You can open a Cisco TAC case to troubleshoot proactively when memory utilization is more than 85%.
    - Example: 10680000 / 24499856 = 0.436 x 100 = 43.6% Free, 100% - 43.6% = 56.4% Used
-
- Leaf and spine switches from two different fabrics cannot be connected regardless of whether the links are administratively kept down.
- Only one instance of OSPF (or any multi-instance process using the managed object hierarchy for configurations) can have the write access to operate the database. Due to this, the operational database is limited to the default OSPF process alone and the multipodInternal instance does not store any operational data. To debug an OSPF instance ospf-multipodInternal, use the command in VSH prompt. Do not use ibash because some ibash commands depend on Operational data stored in the database.
- When you enable or disable Federal Information Processing Standards (FIPS) on a Cisco ACI fabric, you must reload each of the switches in the fabric for the change to take effect. The configured scale profile setting is lost when you issue the first reload after changing the FIPS configuration. The switch remains operational, but it uses the default port scale profile. This issue does not happen on subsequent reloads if the FIPS configuration has not changed.
    - FIPS is supported on Cisco NX-OS release 13.2(10) or later. If you must downgrade the firmware from a release that supports FIPS to a release that does not support FIPS, you must first disable FIPS on the Cisco ACI fabric and reload all of the switches in the fabric.
- Link-level flow control is not supported on leaf switches that are running in ACI mode.
- You cannot use the breakout feature on a port that has a port profile configured on a Cisco N9K-C93180LC-EX switch. With a port profile on an access port, the port is converted to an uplink, and breakout is not supported on an uplink. With a port profile on a fabric port, the port is converted to a downlink. Breakout is currently supported only on ports 1 through 24.

© 2024 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

- On Cisco 93180LC-EX Switches, ports 25 and 27 are the native uplink ports. Using a port profile, if you convert ports 25 and 27 to downlink ports, ports 29, 30, 31, and 32 are still available as four native uplink ports. Because of the threshold on the number of ports (which is maximum of 12 ports) that can be converted, you can convert 8 more downlink ports to uplink ports. For example, ports 1, 3, 5, 7, 9, 13, 15, 17 are converted to uplink ports and ports 29, 30, 31 and 32 are the 4 native uplink ports, which is the maximum uplink port limit on Cisco 93180LC-EX switches.
    - When the switch is in this state and if the port profile configuration is deleted on ports 25 and 27, ports 25 and 27 are converted back to uplink ports, but there are already 12 uplink ports on the switch in the example. To accommodate ports 25 and 27 as uplink ports, 2 random ports from the port range 1, 3, 5, 7, 9, 13, 15, 17 are denied the uplink conversion; the chosen ports cannot be controlled by the user. Therefore, it is mandatory to clear all the faults before reloading the leaf node to avoid any unexpected behavior regarding the port type. If a node is reloaded without clearing the port profile faults, especially when there is a fault related to limit-exceed, the ports might be in an unexpected mode.
- When using a 25G Mellanox cable that is connected to a Mellanox NIC, you can set the ACI leaf switch port to run at a speed of 25G or 10G.
- A 25G link that is using the IEEE-RS-FEC mode can communicate with a link that is using the CL16-RS-FEC mode. There will not be a FEC mismatch and the link will not be impacted.

Related Content

See the Cisco Application Policy Infrastructure Controller (APIC) page for the documentation.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, send your comments to apic-docfeedback@cisco.com. We appreciate your feedback.

Legal Information

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: http://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.

© 2021-2024 Cisco Systems, Inc. All rights reserved.