Cisco Nexus 9000 Series NX-OS Release Notes, Release 10.2(6)M

© 2023 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Introduction

This document describes the features, issues, and exceptions of Cisco NX-OS Release 10.2(6)M software for use on Cisco Nexus 9000 Series switches.

The new Cisco NX-OS Software Release and Image-naming Convention information is available here — Cisco NX-OS Software Strategy and Lifecycle Guide.

Note: The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

The following table lists the changes to this document:

Date	Description
June 28, 2024	Added CSCwi92006 to the Open Issues section.
April 25, 2024	Added CSCwh50989 to the Open Issues section.
September 23, 2023	Added CSCwb53265 to the Resolved Issues section.
September 01, 2023	Cisco NX-OS Release 10.2(6)M became available.

New and Enhanced Software Features

There are no new or enhanced features introduced in Cisco NX-OS Release 10.2(6)M.

Hardware Features

There are no new hardware features introduced in Cisco NX-OS Release 10.2(6)M.

For details on transceivers and cables that are supported by a switch, see the Transceiver Module (TMG) Compatibility Matrix.

Unsupported Features on N9K-C92348GC

Beginning with Cisco NX-OS Release 10.1(1), the following features are not supported on N9K-C92348GC:

- VXLAN
- SW/HW Telemetry
- NetFlow/Analytics
- iCAM
- PTP
- NX-SDK

© 2023 Cisco and/or its affiliates. All rights reserved.

- DME, Device YANG, OpenConfig YANG, gRPC, NETCONF, RESTCONF

Note: NXAPI CLI and XML Agent (NETCONF over SSH) are supported on this platform.

Release Image

In Cisco NX-OS Release 10.2(6)M, the following two 64-bit images are supported:

- The 64-bit Cisco NX-OS image filename that begins with " nxos64-cs" (for example, nxos64-cs.10.2.6.M.bin). This image is supported on all Cisco Nexus 9000 series fixed switches as well as 9000 Modular switches with FM-E/FM-E2/FM-G.
- The 64-bit Cisco NX-OS image filename that begins with " nxos64-msll" (for example, nxos64-msll.10.2.6.M.bin). This image is supported on Cisco Nexus 9000 -R and -R2 series modular switches.

The 32-bit image is no longer supported.

Open Issues

Bug ID	Description

CSCwh18918	Headline: Nexus 9000- python3 crashes being observed after upgrade
	Symptoms: Several python3 core files are created after upgrading to 10.2.4.M. You can use the command show cores to see the cores generated.
	
	`Nexus-switch# show cores`
	`VDC  Module  Instance  Process-name     PID       Date(Year-Month-Day Time)`
	`---  ------  --------  ---------------  --------  -------------------------`
	`1    1       1         python3          --`
	
	Workarounds: Configure clock timezone from the device:
	
	`conf t`
	`clock timezone PST -8 0`
	`end`
	
	After having completed this configuration, you need to check from shell that localtime file has been populated.
	
	`bash-4.4# ls -al /etc/localtime`
	`-rw-r--r-- 1 root root 132 Jul  5 11:49 /etc/localtime`
	`bash-4.4# cat /etc/localtime`
	`TZif2PSTTZif2t`
	`T`
	
	`bash-4.4#`
	`bash-4.4# strings /etc/localtime`
	`TZif2`
	`TZif2`
	
	`bash-4.4#`

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwh19743	Headline: Nexus 9000: 'flowcontrol send on' is configured on Port-channel after removing FEX HIF member ports.
	
	Symptoms: 'flowcontrol send on' configuration is added automatically to Port-channel interface.
	`interface port-channel63`
	`  switchport`
	`  switchport mode trunk`
	`  switchport trunk allowed vlan 2506`
	`  spanning-tree port type edge trunk`
	
	`   flowcontrol send on   <<<Flowcontrol configuration added this way cannot be modified.`
	
	Workarounds: Perform 'no interface port-channel X' to manually remove the port-channel interface and then recreate the port-channel.

CSCwh26488	Headline: Nexus 9000: Policies Applied to NVE On Border Spine Affect BUM Multicast for non-NVE OIFs Symptoms: Multicast traffic for the shared mcast-groups is affected by the NVE Service-Policy even for replicated flows that exist non-NVE OIFs. Workarounds: No workarounds have been identified for the requirement to have the interface service-policy and NVE service-policy to co-exist. The best option would be to only use one of them so that consistent behavior is observed.
CSCwh27705	Headline: N9K NBM: OIL is missing form "show ip mroute" for a static NBM flows. Symptoms: The egress interface is not seen for some random flows in the outgoing interface list in multicast routing table. Workarounds: None
CSCwh27987	Headline: N9K: command "power inline auto" does not seem to work for the class 4 powered devices. Symptoms: On a Nexus 9000 switches running in POE mode facing issues with auto negotiation on PoE power to be used on powered device. Workarounds: Powered device works with "power inline static max xxxx".
CSCwh28126	Headline: Mac moves seen after flapping a member of the peer-link Symptoms: In vPC back-to-back mac moves are seen on N9300-3 and N9300-4 during flapping a member of the peer-link between N9300-1 and N9300-2 creating a loop. Workarounds: None
CSCwh30325	Headline: VTEP flooding packet from fabric back to fabric in L3VNI with host dst mac Symptoms: VTEP flooded packet recieved from fabric back to other VTEPs with L3VNI and host MAC as destination MAC of encapsulated packet. Workarounds: None

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwh30858	Headline: Executing Control Plane Policy stats takes longer than expected on 93360YC-FX2.
	
	Symptoms: Most of the ports on these switches are populated with GLC-TE Transceivers. The following are the symptoms:
	`ottprdoobsw01a# sho clock ; show policy-map int control ; show clo`
	`22:31:39.884 UTC Wed Aug 09 2023`
	`Time source is NTP`
	`Control Plane`
	`<snip>`
	`22:32:20.513 UTC Wed Aug 09 2023`
	`Time source is NTP`
	
	As we can observe the qos statistics takes from 50 seconds to Service not responding -
	
	`ottprdoobsw01a(config)# qos sta`
	`ottprdoobsw01a(config)# end`
	`ottprdoobsw01a# show ver intern bui ; show cloc ; show policy-map int contro | no ; sho clo`
	`nxos image file: bootflash:///nxos64-cs.10.2.4.M-statsfinal.bin : S0`
	`03:41:31.393 UTC Sat Aug 19 2023`
	`Time source is NTP`
	`Control Plane`
	`<snip>`
	`03:42:20.690 UTC Sat Aug 19 2023`
	`Time source is NTP`
	
	`ottprdoobsw01a# show ver intern bui ; show cloc ; show policy-map int contro | no ; sho clo`
	`nxos image file: bootflash:///nxos64-cs.10.2.4.M-statsfinal.bin : S0`
	`03:43:58.104 UTC Sat Aug 19 2023`
	`Time source is NTP`
	`Control Plane`
	`<snip>`
	`03:44:37.623 UTC Sat Aug 19 2023`
	`Time source is NTP`
	
	`ottprdoobsw01a# show ver intern bui ; show cloc ; show policy-map int contro | no ; sho clo`
	`nxos image file: bootflash:///nxos64-cs.10.2.4.M-statsfinal.bin : S0`
	`03:45:03.357 UTC Sat Aug 19 2023`
	`Time source is NTP`
	`Service not responding`
	`03:46:03.793 UTC Sat Aug 19 2023`
	
	`Time source is NTP`
	
	Workarounds: Use no qos statistics.

CSCwh50989	Headline: Custom COPP causing transit traffic to be punted to the CPU on Nexus 9300-GX2 Symptoms: When custom-COPP policy contains ACL rules which match on Layer 4 destination or source port, transit traffic also hits the COPP and the packets are copied to CPU. This causes duplication of traffic as CPU also routes the copied packets to the destination. Workarounds: Custom COPP policy using src/dst match mitigates punt for transit traffic.
CSCwi92006	Headline: all interfaces shown input/output rate as 0pps on N9K FX2 platform Symptoms: all interfaces shown input/output rate as 0pps on N9K-FX2 platform without any operations. Impact is on N9K-FX2 switches and does not impact the N3K. Workarounds: None.

© 2023 Cisco and/or its affiliates. All rights reserved.

Resolved Issues

Bug ID	Description
CSCwb53265	Headline: Netconf core seen when performing netconf merge operation for ospf auth. Symptoms: The symptoms are as follows: 1. NX-API Sandbox freezes and configuration does not take place on the device. After waiting for approximately 10 minutes, it unfreezes; else, clear the netconf session. 2. The policyelem process crashes. 3. The device hangs. 4. Netconf process crashes. Workarounds: None
CSCwd89402	Headline: BFD auth config not displayed in "show run" after upgrade to 9.3.10, without any functional impact Symptoms: BFD authentication config gets lost from running-config after upgrading N9k to 9.3.10, however it's present in the startup-config. However, sessions remain up using authentication. Workarounds: The "no BFD echo" command needs to be configured before the upgrade. This works as both preventive and reactive workaround.

CSCwe55237	Headline: Nexus: Configuring overlapping IPv6 address with different mask on L3 intf does not throw error
	Symptoms: Observed that if you configure an IPv6 address which is overlapping on two different L3 interfaces of the same box using different subnets, no warning message is seen. This is not the case for IPv4.
	`N9K-2-382(config-if)# int loop29`
	`N9K-2-382(config-if)# ipv6 add 2001:db0::1/72  <<< 2001:db0::1/72 configured on loop29`
	`N9K-2-382(config-if)# no shut`
	`N9K-2-382(config-if)# int loop31`
	`N9K-2-382(config-if)# ipv6 address 2001:db0::1/73`
	`Address already configured on interface loopback29 <<< warning seen when the IP is duplicated`
	`N9K-2-382(config-if)# ipv6 add 2001:db0::1/68`
	`Address already configured on interface loopback29  <<< warning seen when the IP is duplicated`
	`N9K-2-382(config-if)# ipv6 add 2001:db0::3/68   <<< accepted without warning`
	Workarounds: Use a subnet calculator to ensure your addresses are not overlapping.

CSCwf34708	Headline: N9K OS installation with no-reload option can cause BFD down Symptoms: BFD cannot go up after upgrading OS as follows.1. Run "install all nxos [os image] no-reload" on all switches.2. Reload one of switches.3. BFD neighbor down and cannot come up until reload the opposite switch. Besides, OSPF neighbor with BFD is still in full state even though BFD is down. Workarounds: Reload other switches.
CSCwf52916	Headline: Improve telemetry trace tm-error "failure reason:Value too large" Symptoms: " failed reason:Value too large " from "show system internal telemetry trace tm-errors" Workarounds: None

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf57548	Headline: N9300 sends icmp Frag needed and DF set msg to NAT Inside global instead of inside local IP Symptoms: Application failure due icmp msg "Frag needed and DF set" sent to wrong NAT inside global IP instead of NAT inside local IP. Workarounds: Set equal MTU on both NAT inside and outside interfaces.
CSCwf72985	Headline: Telemetry subscription of DME path sys/ptp/correction as an event is not working Symptoms: Telemetry subscription with DME path sys/ptp/correction as an event is not working. The Nexus 9K is not generating push packets for the subscription of DME path sys/ptp/correction as an event while the PTP clock correction falls into a configured correction range. DME does not list the event driven telemetry subscription. Workarounds: 1. Under the affected telemetry sensor-group, reconfigure the DME path. EXAMPLE - sensor group 302 with DME path sys/ptp/correction conf t sensor-group 302 no path sys/ptp/correction depth unbounded path sys/ptp/correction depth unbounded end. This would have to be done for each new subscription to prevent the same sensor-group to be used in more than 1 subscription.
CSCwf74305	Headline: Logging 2.0: reduce number of instances that autcollect run to reduce load on CLI/vsh Symptoms: Slow CLI execution in EXEC mode or CONF mode for period of the time when high rate of severity 0,1,2 unique syslog messages are logged by nxos (high rate is more than 100 messages per second) Workarounds: Disable autocollect feature:event manager applet syslog_trigger override __syslog_trigger_default action 1.0 collect disable. NOTE: Do not add any other action in the EEM applet as that will prevent autocollect to be disabled. Only single action can be present in this applet

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf75437	Headline: N9K-C9504 after system switchover, static route(configured static route BFD) will disappear Symptoms: Before C9504 switchover: N9K-C9504-2(config-if)# sh module Mod Ports Module-Type Model Status --- ----- ------------------------------------ ---------- --------- 1 52 48x10/25G + 4x40/100G Ethernet Module N9K-X97160YC-EX ok 2 52 48x10/25G + 4x40/100G Ethernet Module N9K-X97160YC-EX ok 22 0 4-slot Fabric Module N9K-C9504-FM-E ok 23 0 4-slot Fabric Module N9K-C9504-FM-E ok 24 0 4-slot Fabric Module N9K-C9504-FM-E ok 26 0 4-slot Fabric Module N9K-C9504-FM-E ok 27 0 Supervisor Module N9K-SUP-B+ active * 28 0 Supervisor Module N9K-SUP-B+ ha-standby 29 0 System Controller N9K-SC-A standby 30 0 System Controller N9K-SC-A active N9K-C9504-2(config-if)# sh ip route 172.168.1.0 IP Route Table for VRF "default" '*' denotes best ucast next-hop '**' denotes best mcast next-hop '[x/y]' denotes [preference/metric] '%<string>' in via output denotes VRF <string> 172.168.1.0/24, ubest/mbest: 1/0 *via 192.168.2.2, Eth1/22, [1/0], 00:19:37, static N9K-C9504-2# sh bfd neighbors dest-ip 192.168.2.2 OurAddr NeighAddr LD/RD RH/RS Holdown(mult) State Int Vrf Type 192.168.2.1 192.168.2.2 1090519046/1090519042 Up 5504(3) Up Eth1/22 default SH The information of the engine and Routing table BFD before the switchover that looks like all normal. After C9504 switchover: N9K-C9504-2# sh module Mod Ports Module-Type Model Status --- ----- ------------------------------------ ---------- --------- 1 52 48x10/25G + 4x40/100G Ethernet Module N9K-X97160YC-EX ok 2 52 48x10/25G + 4x40/100G Ethernet Module N9K-X97160YC-EX ok 22 0 4-slot Fabric Module N9K-C9504-FM-E ok 23 0 4-slot Fabric Module N9K-C9504-FM-E ok 24 0 4-slot Fabric Module N9K-C9504-FM-E ok 26 0 4-slot Fabric Module N9K-C9504-FM-E ok 27 0 Supervisor Module N9K-SUP-B+ ha-standby 28 0 Supervisor Module N9K-SUP-B+ active * 29 0 System Controller N9K-SC-A standby 30 0 System Controller N9K-SC-A active N9K-C9504-2# sh ip route 172.168.1.0 IP Route Table for VRF "default" '*' denotes best ucast next-hop '**' denotes best mcast next-hop '[x/y]' denotes [preference/metric] '%<string>' in via output denotes VRF <string> Route not found <<<<<<<<< N9K-C9504-2# sh bfd neighbors dest-ip 192.168.2.2 OurAddr NeighAddr LD/RD RH/RS Holdown(mult) State Int Vrf Type 192.168.2.1 192.168.2.2 1090519046/1090519042 Up 5504(3) Up Eth1/22 default SH

Bug ID	Description
CSCvf15936	Headline: m9700::non-sysmgr cores are seen after topology configuration Symptoms: non-sysmgr crashed on executing event manager action-log cli on 9700 platform2017 Jul 5 02:25:36 sw1-gd122 %SYSMGR-2-LAST_CORE_BASIC_TRACE: : PID 1198 with message non-sysmgr (non-sysmgr) crashed, core will be saved . Workarounds: Action-logs can be regularly cleared from the box to avoid the crash.

CSCwd63941	Headline: License information unavailable and showing error after ISSU
	
	Symptoms: License information unavailable and showing error after moving to 9.3(2) build.
	
	<pre>
	switch01# show license usage
	Could not get license usage: License server is busy. Please retry in a few seconds.
	switch01# show license authorizations
	Could not get data authorizations info: err=144
	
	switch01# show license status
	Could not get utility info: err=1
	
	show utility info failed
	Smart Licensing using Policy:
	  Status: ENABLED
	
	Could not get data privacy info: err=1
	
	</pre>
	
	Workarounds: Contact TAC for workaround.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCvk54147	Headline: LC module goes to failure state while collecting show tech binary after ISSU
	
	Symptoms: While collecting show tech binary, one or more LC modules in the system may reset due to a fln_que crash. A core file can be seen in `show cores` and an exceptionlog in `show module internal exceptionlog`:
	
	`show cores`
	VDC Module Instance Process-name PID Date(Year-Month-Day Time)
	--- ------ -------- --------------- -------- -------------------------
	1 10 1 fln_que 1201 2018-10-22 17:25:12
	1 1 1 fln que 1202 2018-10-22 17:25:26
	
	`show module internal exceptionlog`
	
	********* Exception info for module 10 ********
	
	exception information --- exception instance 1 ----
	Module Slot Number: 10
	Device Id : 134
	Device Name : System Manager
	Device Errorcode : 0x00000426
	Device ID : 00 (0x00)
	Device Instance : 00 (0x00)
	Dev Type (HW/SW) : 04 (0x04)
	ErrNum (devInfo) : 38 (0x26)
	System Errorcode : 0x401e008a Service on linecard had a hap-reset
	Error Type : FATAL error
	PhyPortLayer : 0x0
	Port(s) Affected :
	Error Description : fln_que hap reset
	DSAP : 0 (0x0)
	UUID : 1 (0x1)
	Time : Mon Oct 22 17:25:26 2018
	
	(Ticks: 5BCDBAAE jiffies)
	
	Workarounds: Avoid running show-tech binary and use regular show tech if possible.

CSCwd68297	Headline: SNMPd Crashes when Configuring 'event snmp-notification' EEM Script Symptoms: SNMPd Crashes a couple of times following an EEM config addition. Workarounds: Unknown, aside from avoiding the configuration of the script in the first place.
CSCwd01610	Headline: BGP AS not updated properly in Netflow flow cache Symptoms: - The routes are showing as Src As and Dst AS 0 when they should be populating at least of one of them with a BGP AS- The BGP AS values are shown as non-zero even though the current routes are non-BGP (like static or ospf) Workarounds: Reload the VDC or chassis
CSCwe02602	Headline: PIM-Process Crash Symptoms: %SYSMGR-3-HEARTBEAT_FAILURE: Service "pim" sent SIGABRT for not setting heartbeat for last 7 periods. Last heartbeat 210.94 secs ago. Workarounds: None. PIM will restart after the crash.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwe94284	Headline: OSPF Process is increasing memory utilization Symptoms: Switches are reporting the following memory errors continuously: %OSPF-1-NOMEM: ospf-XXXX [16983] Malloc (65571 bytes) failed for: OSPF MEM CMI SNAPSHOT %OSPF-4-SYSLOG_SL_MSG_WARNING: OSPF-1-NOMEM: message repeated 5 times in last 46375039 sec > "MemUsed" is very close to the "MemLimit" for the OSPF process # show processes memory sort PID MemAlloc MemLimit MemUsed StackBase/Ptr Process ----- -------- ---------- ---------- ----------------- ---------------- 9035 1560645632 2066057164 2066014208 ffe3af40/ffe3a760 ospf >>> 9026 126746624 1451523033 719093760 fff470f0/fff46fd0 mrib 8816 153280512 1588493555 717537280 ffd696b0/ffd69180 igmp 8235 87703552 1419916377 700219392 ffa2e8e0/ffa2e540 arp 8861 208924672 1059769830 699904000 ffe15c20/ffe1409c monitor Workarounds: The leaks can be avoided by avoiding an SNMP walk for the OID 1.3.6.1.2.1.14.17 or any of it's parents OIDs.
CSCwf18783	Headline: Deprecated OSPF Network Commands seen in the running-config after an upgrade. Symptoms: After upgrading from a previous version where OSPF network commands are supported, the network commands cease to function and OSPF neighborships associated with those networks are brought down until an equivalent interface configuration is applied on the switch. Even though it no longer works, the network command is still present in the running config and cannot be removed as all the commands for the OSPF network commands (including the negation of the command) has been removed. Workarounds: Delete and re-apply either the OSPF process or the feature itself. If the feature is removed, all OSPF configuration must be re-applied through interface configuration and any attempt to include the network statements under the routing process again will be denied.
CSCwf38091	Headline: EIGRP distribution-list out allows route that should be denied after SSO Symptoms: In the working state, the switch will block routes with tag 5 from being advertised out e2/1 as expected with this distribution-list. But if I do a SSO of the switch, we see the routes that should be denied are advertised to the EIGRP neighbor. Workarounds: Deny routes going out the interface and set the tag ingress on the other device.
CSCwe55960	Headline: Port LED is not working correctly for N9K-C9364D-GX2A Symptoms: Port LED is not working correctly for some interfaces on N9K-C9364D-GX2AFor example,if we bring up four ports in the same row (E1/19, E1/20,E1/51,E1/52)The LED shows green for port E1/19 & E1/20 only. if we shut down the E1/51, the LED for E1/20 will become amber. Workarounds: N/A
CSCvg06451	Headline: Remove CLI to configure Fill Pattern as only IDLE is supported Symptoms: switchport fill-pattern is not a supported configuration on N9k, as only IDLE patterns are supported. Hence configuration CLI is removed Workarounds: None

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwa54414	Headline: Static MACs conf on int NVE deleted from vPC secondary after int NVE shut/no shut on vPC primary
	
	Symptoms: Static MACs configured on interface NVE may be deleted from vPC secondary device after interface NVE shutdown /no shutdown was executed on vPC primary.
	
	Workarounds: Delete/re-create the VLAN/VNI where the static MAC was originally configured and add back the static MAC on interface NVE.
	
	`no vlan <vlan_id>`
	`vlan <vlan id>`
	`vn-segment <segment-id>`
	
	`mac address-table static <mac address> vni <segment-id> interface nve 1 peer-ip <IP_address>`

CSCwa99186	Headline: IF-MIB : IFHIGHSPEED object returns wrong interface speed values for internal interfaces (ii x/y/z)
	
	Symptoms: Only internal interfaces looks to be affected (example ii1/1/1).
	
	`snmpwalk -v 2c -c test <ip_address> iso.3.6.1.2.1.31.1.1.1.15.1241513984`
	`IF-MIB::ifHighSpeed.1241513984 = Gauge32: 42000 < ---------  expecting 100000`
	
	`` `show hardware internal tah all-ieth-ports` is reporting 100G for both ieth and fp ports. ``
	
	`This problem is not seen if you query front-port interface like Eth1/1:`
	
	`snmpwalk -v 2c -c test <ip_address> iso.3.6.1.2.1.31.1.1.1.15.436207616`
	
	`IF-MIB::ifHighSpeed.436207616 = Gauge32: 100000`
	
	Workarounds: None

CSCwb62058	Headline: N_DEV-125 ipfib core is seen
	
	Symptoms: IPFIP service crash.
	
	`%SYSMGR-SLOT7-2-SERVICE CRASHED: Service "ipfib" (PID 14754) hasn't caught signal 11 (core will be saved).`
	`%SYSMGR-SLOT7-2-HAP_FAILURE_SUP_RESET: Service "ipfib" in vdc 1 has had a hap failure`
	
	This might cause a linecard unexpected reboot.
	
	`` `show system reset-reason module x` ``
	`Error code   : 1075708042`
	`Reset reason : ipfib hap reset => [Failures < MAX] : powercycle`
	`Service name : System manager`
	`Time stamp   : At 12926 usecs after Tue Jun 13 15:16:04 2023`
	
	This is a case where routine used to delete routes has failed in HAL. When Mfib receives the response from HAL, it tries to delete the route again, causing a loop. In this case, trigger seems to be a route delete.
	
	Workarounds: N/A

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwb90953	Headline: POAP not listing all the breakout options supported Symptoms: During the POAP process, the supported breakout options are not listed hence POAP process stops. Workarounds: None
CSCwc06052	Headline: N9K-C93180YC-FX3: GNSS LED is red when GPS/GNSS is not connected and SyncE not configured Symptoms: The GNSS chassis LED of Nexus 93180YC-FX3 switches running NX-OS software releases 10.2(3), 10.2(4), or 10.2(5) may be either lit red, or rapidly oscillating between green and red even though GPS/GNSS is not connected to the switch and the SyncE feature is not configured on the switch. This issue is cosmetic in nature - only the status of the LED is affected by this issue. Workarounds: There is no known workaround for this issue. As this issue is cosmetic in nature, the color of the GNSS LED can be safely ignored if the SyncE feature is not configured and not GPS/GNSS is connected to the switch.
CSCwc72568	Headline: Generation of non-sysmgr core while collecting pss dump of a corrupted db Symptoms: When any application crashes, sysmgr triggers core collection for the application. The pss db dump for the application is also collected as part of core collection. During the process of core collection, if the pss db of the application is already corrupted (due to some memory corruption), then the collection utility (pss2dump utility) crashes. This leads to generation of a non-sysmgr core along with the application core. Workarounds: None

CSCwd11996	Headline: Unable to unconfigure extcommunity-list , if we use the same in 2 route-map at a time
	
	Symptoms: If an extended community list is being called by 2 or more route maps in the current configuration, deleting the list will not reflect properly in running config. Executing "show running-config" will show the removed extcommunity-list still present in the configuration.
	
	`513E-A-17-N93180YC-FX-1(config)# no ip extcommunity-list standard list3 seq 10 permit 4byteas-generic transitive 100:200`
	`513E-A-17-N93180YC-FX-1(config)# sh running-config rpm`
	
	`!Command: show running-config rpm`
	`!Running configuration last done at: Fri Mar 17 02:13:18 2000`
	`!Time: Fri Mar 17 02:13:22 2000`
	
	`version 10.2(5) Bios:version 05.47`
	`ip extcommunity-list standard list3 seq 10 permit 4byteas-generic transitive 100:200`
	`route-map rmap5 permit 10`
	`  match extcommunity list3`
	`route-map rmap6 permit 10`
	
	`  match extcommunity list3`
	
	Workarounds: Unconfigure set command first.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwd23976	Headline: Nexus C9364C-GX - Golden EPLD upgrade times out and resets with Fatal Module Error
	
	Symptoms: A Cisco Nexus9000 C9364C-GX Chassis may experience a reload following a Golden EPLD upgrade. If the EPLD has already been upgraded and is not necessary, performing the unnecessary golden upgrade results in the device timing out and reloading with a Fatal Module Error.
	
	`N9K-C9364C-GX# install epld bootflash:[epld-image] module all golden`
	`...`
	`The switch will be reloaded at the end of the upgrade`
	`Do you want to continue (y/n) ?  [n] y`
	
	`Proceeding to upgrade Modules.`
	
	`Starting Module 1 EPLD Upgrade`
	
	`Module 1 : MI FPGA [Programming] : 100.00% (    # of     # sectors)`
	
	`EPLD process seems to have exited unexpectedly 2700`
	
	`Unable to communicate with the EPLD Process`
	
	`Reloading Supervisor Module 1`
	`...`
	
	`Last reset at [USECS] usecs after [DATE]`
	`Reason: Reset Requested due to Fatal Module Error`
	`System version: [IMAGE]`
	
	`Service:`
	
	Workarounds: None

CSCwd31194	Headline: Nexus 9000 doesn't configure contact and location for snmp.
	
	Symptoms: Devices with release version 10.x. See output:
	
	`switch# conf t`
	`Enter configuration commands, one per line. End with CNTL/Z.`
	`switch(config)# snmp-server location HEREDIA`
	`wrong length of value being set`
	`switch(config)#`
	`switch(config)# snmp-server location Alain`
	
	`wrong length of value being set`
	
	Workarounds: N/A

CSCwd47632	Headline: Memory leak on acllog "acllog_net_l2_pkt_handle" Symptoms: Continuous memory leak observed with ACLLOG process when detailed logging is enabled (logging ip access-list detail), which eventually leads to ACLLOG process crash. Workarounds: Remove "logging ip access-list detailed" from global configuration.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwd63552	Headline: There is mismatch between Scalability Doc and the threshold configured on the switch Symptoms: As per the Cisco Nexus 9000 Series NX-OS Verified Scalability Guide, Release 9.3(10)https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/93x/scalability/guide-9310/cisco-nexus-9000-series-nx-os-verified-scalability-guide-9310.html the switch supports 48000 STP instances however the devices triggers a log message when more than 14000 are in use %STP-2-VLAN_PORT_LIMIT_EXCEEDED: The number of vlan-port instances exceeded [Rapid-PVST mode] recommended limit of 14000 Workarounds: None
CSCwd81099	Headline: CLI CR failing with default radius server configs Symptoms: 1) CLI Config-Replace fails with default 'radius' server group 2) Default 'radius' entry is shown in running-config even after defaulting/removing its children config Workarounds: Use non-default AAA group.
CSCwd82487	Headline: N9K VxLAN: MAC Mobility Sequence is not getting incremented for mac only bgp update Symptoms: • After moving host from one leaf to the other, mac mobility sequence is not getting incremented for mac only bgp update. • For Mac-ip MAC Mobility Sequence gets incremented. Workarounds: Use clear mac address-table dynamic address <address> command.
CSCwd85017	Headline: Rx Pause enabled in h/w even when flow control is disabled in s/w causing pause frames to be honoured Symptoms: • Rx Pause enabled in hardware even when flow-control is disabled in Software causing pause frames to be honoured. • As these pause frames are honoured in the hardware, might cause packet drops/performance issues for both Unicast and Multicast Traffic. Workarounds: None
CSCwd87170	Headline: snmpbulkget to ciscoEntityFRUControlMIB create invalid unicode in sh snmp internal event-his pktdump Symptoms: Unicode Symbols in snmp packet buffer dump when doing FRUget snmpbulkget Workarounds: Use Snmpbulkget with two instances or avoid using "show system internal snmp event-history pktdump" command.
CSCwe06759	Headline: Memory leak at FEX unit Symptoms: FEX device will reload unexpectedly. FEX devices will appear to be offline on the parent switches and will become online again. There will be loss of traffic ingressing on the fex when it reloads. Workarounds: None

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwe09300	Headline: Internal BGP routes are getting installed as External routes with an AD of 20 in the Routing Table
	
	Symptoms: Nexus installs ECMP route with iBGP and eBGP next-hops, even though command "maximum-paths eibgp" is not configured. Here is an example of an unexpected behavior:
	
	`show bgp vrf all all`
	`Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best`
	`Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected`
	`Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2`
	
	`*>e10.0.0.0/24 192.168.0.1 0 64500 64501 ?`
	`* i 192.168.1.1 100 0 64500 64501 ?`
	`*|i 192.168.1.2 100 0 64500 64501 ?`
	`* i 192.168.1.3 100 0 64500 64501 ?`
	
	We can see that eBGP route with next hop 192.168.0.1 was selected as best (>) and iBGP route with next hop 192.168.1.2 as multipath (|).
	
	Workarounds: Workaround is to re-initiate route installation in RIB. Examples of possible workarounds: - re-announce affected network - remove multipath command - change best path selection algorithm "bestpath med non-deterministic" - restart BGP process

CSCwe17461	Headline: Maximum routes X warning-only starts dropping routes under VRF Symptoms: Under a vrf context, the command maximum routes <N> warning-only should give a syslog whenever there are N routes in the VRF. It should not prevent routes to be added if the limit is reached, it only gives a warning log. What we see is the routes are rejected and thus not added to the database when the limit is reached. Workarounds: No workaround other than not using the configuration.
CSCwe18776	Headline: N9k | PBR Traffic may egress out via incorrect interfaces under certain conditions Symptoms: PBR Traffic leaves incorrect egress interfaces. One example of a failed traffic flow would be that the PBR Next hop is reachable via an orphan port but the actual PBR Traffic leaves via an unrelated Port-channel interface Workarounds: Bounce the SVI which has the PBR configured.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwe25343	Headline: Nexus 9000 VxLAN virtual peer-link tunnel recovery failure.
	
	Symptoms: After a specific failure of the virtual peer-link the tunnel used to traverse CFS traffic may fail to recover after the virtual peer-link is brought back up. The virtual peer-link shows "UP" but no VLANS forwarding over the peer-link and CFS traffic does not make it through the tunnel.
	
	<pre>
	N9k-SW01# show vpc brief
	<snip>
	Legend:
	(*) - local vPC is down, forwarding via vPC peer-link
	
	vPC domain id : 1
	Peer status : peer adjacency formed ok
	vPC keep-alive status : peer is alive
	Configuration consistency status : success
	Per-vlan consistency status : success
	Type-2 consistency status : success
	vPC role : secondary
	Number of vPCs configured : 10
	Peer Gateway : Enabled
	Dual-active excluded VLANs : -
	Graceful Consistency Check : Enabled
	Auto-recovery status : Enabled, timer is off.(timeout = 360s)
	Delay-restore status : Timer is off.(timeout = 150s)
	Delay-restore SVI status : Timer is off.(timeout = 10s)
	Operational Layer3 Peer-router : Disabled
	Virtual-peerlink mode : Enabled
	
	vPC Peer-link status
	--- 	
	id Port Status Active vlans
	-- ---- ------ --
	1 Po500 up -
	</pre>
	- The switch shows ICMP connectivity and a route installed to the remote vMCT loopback IP
	- The switch shows vPC vlans in an error state due to "vPC peer is not reachable over cfs"
	<pre>
	N9k-SW01# show interface status err-vlans
	
	-- 	
	Port Name Err-Vlans Status
	-- 	
	Po500 "vpc-peer-link" 1,1401-1404,1406-1407,1411- vPC peer is not
	1412,1421,1424,1429-1430,14 reachable over
	37-1440,3201-3207 cfs
	
	</pre>
	
	Workarounds: Flapping the uplinks that are used to build the virtual peer-link will force the virtual peer-link to re-initialize and should recover from this state. If flapping the uplinks does not recover the peer-link, a reload is required.

CSCwe29161	Headline: N9300-GX: One or more interfaces are configured with an all zero MAC Symptoms: On a N9K-C93600CD-GX switch running impacted code, one or more ports may have an all zeroes mac, like belowEthernet1/22 is upadmin state is up, Dedicated Interface belongs to Po13 Hardware: 40000/50000/100000 Ethernet, address: 6c31.0e2d.e147 (bia 0000.0000.0000) MTU 9216 bytes, BW 100000000 Kbit , DLY 10 usec. This will result in issues with control plane protocols which try to use the port reporting issues with the MAC or failing to function as expected. Remotely connected devices may report errors about this MAC under CDP or LLDP like follows:2023 Feb 4 11:26:35 Remote-N9K %LLDP-2-INVALID_LLDP_RECEIVED: Received an invalid LLDP on Eth3/1 Reason: Invalid LLDP SRC MAC Workarounds: None

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwe31550	Headline: LPM entry gets stuck after multiple bgp withdraws
	
	Symptoms: In an ECMP environment where the BGP prefixes have constant withdraws, it is possible for the "single entries" (no ECMP) to get stuck in LPM table. As a consequence, the ECMP entry ends up installed at the end and deferred, because we honour the first one (older entry).
	
	`show hardware internal tah l3 v4lpm prefix 20.20.48.0/24 table 1"`
	
	`**EPE label`
	`*Flags:`
	`CC=Copy To CPU, SR=SA Sup Redirect,`
	`DR=DA Sup Redirect, TD=Bypass TTL Dec,`
	`DC=SA Direct Connect,DE=Route Default Entry,`
	`LI=Route Learn Info,HR=Host as Route`
	`Idx | vrf | ip/len | mpath | nump | base/l2ptr`
	`|cc|sr|dr|td|dc|de|li|hr|`
	`----------|---------|-------------------------------|-------|-------|------------|--|--|--`
	`|--|--|--|--|--|`
	`13/711 | 1 | x.x.x./24 | 0 | 0 | 0x50005 | | | | |`
	`| | | | <<<<<<<< "No ECMP" path entry not cleared`
	`7/799 | 1 | x.x.x.x/24 | 1 | 2 | 0x400004e1 | | | | |`
	`| | | |. <<<<< new ECPM entry`
	
	Workarounds:
	1. Clear the route works as a temporary workaround, as this forces nxos to put the ECMP entry at the top. However, if another withdraw occurs we end up hitting the same condition.
	
	`clear ip route x.x.x.x/x`
	
	2. Clear the LPM entry works as a better workaround. However, if the constant withdraws continue it might end up hitting the same condition
	
	`Attach module 1 ; clear hardware internal tah l3 ipv4 tcam-bank z tcam-index y`
	`Idx | vrf | ip/len | mpath | nump | base/l2ptr`
	`|cc|sr|dr|td|dc|de|li|hr|`
	`----------|---------|-------------------------------|-------|-------|------------|--|--|--`
	`|--|--|--|--|--|`
	
	`Z/Y | 1 | x.x.x.x/x | 1 | 2 | 0x------------ | |`
	`| | | | | | |`

CSCwe34915	Headline: 92348GC consuming -XF license instead of -GF license Symptoms: Cisco Nexus 92348GC platform usually consumes -GF license; instead, it is consuming -XF license, when respective features are enabled, for example, feature bgp is enabled, and license consumption is checked Workarounds: None

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwe35981	Headline: Unicast route create failed and N9K tahusd crashed
	
	Symptoms: Unicast route create failed logs due to Trie Hw write failed.
	After 1 second, service "tahusd" crashed.
	
	`2023 Feb 13 10:04:41.230 %$ %IPFIB-SLOT1-2-UFIB ROUTE CREATE: Unicast route create failed for INS unit 0, VRF: 1, xxx.xxx.xx.0/24, flags:0x0, intf:0x40000087, Error: Trie Hw write failed(220)`
	`2023 Feb 13 10:04:41.230 %$ %IPFIB-SLOT1-2-UFIB_ROUTE_CREATE: Unicast route create failed for INS unit 0, VRF: 1, xxx.xxx.xx.0/24, flags:0x0, intf:0x40000087, Error: Trie Hw write failed(220)`
	`2023 Feb 13 10:04:41.230 %$ %IPFIB-SLOT1-2-UFIB ROUTE CREATE: Unicast route create failed for INS unit 0, VRF: 1, xxx.xxx.xx.0/24, flags:0x0, intf:0x40000087, Error: Trie Hw write failed(220)`
	`2023 Feb 13 10:04:42.245 %$ %SYSMGR-SLOT1-2-SERVICE CRASHED: Service "tahusd" (PID 25057) hasn't caught signal 11 (core will be saved).`
	`2023 Feb 13 10:04:42.248 %$ %SYSMGR-SLOT1-2-HAP_FAILURE_SUP_RESET: Service "tahusd" in vdc 1 has had a hap failure`
	`2023 Feb 13 10:04:42.249 %$ %SYSMGR-SLOT1-2-LAST CORE BASIC TRACE:`
	`fsm action become offline: PID 10840 with message Could not turn off console logging on vdc 1 error: mts req-response with syslogd in vdc 1 failed (0xFFFFFFFF) .`
	`2023 Feb 13 10:05:21.343 %$ %KERN-2-SYSTEM_MSG: [ 4542.138123] usd process 25057, uuid 1356 (0x54c) tahusd failed to send heartbeat - kernel`
	
	show system reset-reason
	`1) At 647356 usecs after Mon Feb 13 10:07:08 2023`
	`Reason: Reset Requested due to Fatal Module Error`
	`Service: tahusd hap reset`
	`Version: 9.3(9)`
	
	`* You may also observe`
	
	`UFIB MULTIPATH TABLE EXHAUSTION: Unicast ecmp table exhausted.`
	
	Workarounds: None

CSCwe38874	Headline: Certain interfaces shown input/output rate as 0pps on N9K GX platform Symptoms: Certain interfaces shown input/output rate as 0pps on N9K-GX platform without any operations. Workarounds: 1.The issue can be cleared after flapped one of those ports. 2.Reactivate the DMA state with internal commands. 3.Connecting a new terminal to the port to which the DMA Group belongs will reactivate the DMA Group.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwe41298	Headline: xbar multicast show command causes 300MB of memory to be allocated
	
	Symptoms: When "show system internal xbar multicast_id all" is run, there is no output from the command. However, you see that 300MB of memory is allocated to "XBM_MEM_xbm_mcast_group_info_t
	
	`switch# show processes memory | i  "All processes"`
	`All processes: MemAlloc = 20795484`
	
	`switch# show system internal xbar mem-stats detail | i XBM MEM xbm mcast group info t`
	`switch# show system internal xbar multicast id all`
	`switch# show system internal xbar mem-stats detail | i XBM_MEM_xbm_mcast_group_info_t`
	`  18 XBM_MEM_xbm_mcast_group_info_t             512    512  307761152  307761152`
	
	`switch# show processes memory | i  "All processes"`
	
	`All processes: MemAlloc = 21096364`
	
	Workarounds: Clear the memory allocation by reloading the switch. Then avoid running the show system internal xbar multicast_id all command, and also commands such as show tech-support details.

CSCwe41327	Headline: SYSMGR-3-CFGWRITE_FAILED: Configuration copy failed due to 100% usage of startup-cfg
	
	Symptoms: Nexus switches may fail to save the config to startup config due to frequent changes to the ACL config with below errors:
	`2023 Feb 10 13:17:11.746 N3k %PSS-0-PSS WRITE DATA FAILURE: aclmgr: failed to write data to /var/sysmgr/startup-cfg/bin/aclmgr start cfg user block 8785: partial write`
	`2023 Feb 10 13:17:11.746 N3k%PSS-0-PSS WRITE DATA FAILURE: aclmgr: failed to write data to /var/sysmgr/startup-cfg/bin/aclmgr_start_cfg_user block 8785: partial write (message repeated 1 time)`
	`2023 Feb 10 13:17:11.746 N3k %PSS-0-PSS WRITE FAILURE: aclmgr: failed to write data: Operation not permitted`
	`2023 Feb 10 13:17:11.760 N3k %SYSMGR-3-CFGWRITE SRVFAILED: Service "aclmgr" failed to store its configuration (error-id 0x4048000C).`
	`2023 Feb 10 13:17:11.927 N3k %SYSMGR-2-CFGWRITE_ABORTED: Configuration copy aborted.`
	`2023 Feb 10 13:17:14.017 N3k %SYSMGR-3-CFGWRITE_FAILED: Configuration copy failed (error-id 0x401E0000).`
	
	`switch# show system internal flash | in Mount|startup`
	`Mount-on  1K-blocks  Used  Available  Use%  Filesystem`
	`/var/sysmgr/startup-cfg  102400  90632  11768  89  none`
	
	`switch#`
	
	Workarounds: To avoid growing usage of startup-cfg by making config changes without config session/commitOnce startup-cfg usage hits 100%, the config save to startup will fail. Recovery is by switch reload. The running config can be saved to bootflash if additional config can't be saved to the bootflash.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwe42043	Headline: BGP External-Fallover not working when TTL-Security is enabled eBGP non multihop sessions)
	
	Symptoms: When TTL-Security is enabled, BGP External-Fallover does not bring down the eBGP session as soon as the physical interface (Eth1/47) is down. Instead, BGP is waiting for the hold timer to expire (which in below example was configured to be 20 secs)
	
	`2023 Feb 20 09:50:47 N9K-2 %ETHPORT-5-IF_DOWN_PARENT_DOWN: Interface Ethernet1/47.3182 is down (Parent interface is down)`
	`2023 Feb 20 09:50:47 N9K-2 %ETHPORT-5-IF DOWN ADMIN DOWN: Interface Ethernet1/47 is down (Administratively down)`
	`2023 Feb 20 09:51:03 N9K-2 %BGP-5-ADJCHANGE:  bgp- [22583] (VPN IC) neighbor 10.83.1.25 Down - sent:  holdtimer expired error`
	
	Workarounds: TTL-Security on eBGP sessions which are non eBGP mulithop sessions should be disabled.

CSCwe43450	Headline: Kernel panic due to "Fatal Module Error" after ND ISSU on N9K-C9348GC-FXP Symptoms: After ND-ISSU, an unexpected reload due to kernel panic is noticed in N9K-C9348GC-FXP switches. This symptom can be verified by running either the `show logging onboard internal reset-reason` command or the `show system reset-reason` command. The output shows Reset Requested due to Fatal Module Error. Workarounds: Use disruptive/normal upgrade procedure.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwe46769	Headline: N9K CoPP logging drop threshold dropped count changes to a large value after upgrade to 64-bit image
	
	Symptoms: A Nexus platform configured with a custom CoPP policy and using the 'logging drop threshold' option may encounter the following issue when upgrading from a 32-bit to a 64-bit NXOS image.
	
	On a 32-bit NXOS image, ex: nxos.7.0.3.I7.9.bin:
	
	`policy-map type control-plane copp-policy-strict-custom`
	`  class copp-class-l3uc-data-custom`
	`  logging drop threshold 7500 level 1`
	
	After upgrading to 64-bit image, nxos64-cs.10.3.2.F.bin:
	
	`policy-map type control-plane copp-policy-strict-custom`
	`  class copp-class-l3uc-data-custom`
	`  set cos 1`
	`  logging drop threshold 32212254720000 level 1`
	
	`C93180YC-FX(config-pmap-c)# show policy-map int control-plane`
	`Control Plane`
	
	`  Service-policy  input: copp-policy-strict-custom`
	
	`  class-map copp-class-l3uc-data-custom (match-any)`
	`  match exception glean`
	`  threshold: 32212254720000, level: 1`
	
	The logging drop drop-count increased to a large value after upgrading.
	
	The new drop-count value cannot be overwritten or removed:
	
	`C93180YC-FX(config-pmap-c)# logging drop threshold 7500 level 1`
	`C93180YC-FX(config-pmap-c)# do sh run | beg "class copp-class-l3uc-data-custom"`
	`  class copp-class-l3uc-data-custom`
	`  logging drop threshold 32212254720000 level 1`
	
	`C93180YC-FX(config)# policy-map type control-plane copp-policy-strict-custom`
	`C93180YC-FX(config-pmap)#  class copp-class-l3uc-data-custom`
	`C93180YC-FX(config-pmap-c)# no  logging drop threshold 32212254720000 level 1`
	`  ^`
	`% Invalid number, range is (1:80000000000) at '^' marker.`
	
	Workarounds: Workaround before upgrading:
	• Remove logging drop threshold command from the configuration
	• Save the configuration Upgrade to the 64-bit image
	• Add logging drop threshold command to the configuration
	• Save the configuration
	
	The following workaround is disruptive:
	• Enable one of the CoPP built in templates
	• Remove the custom CoPP policy configuration
	• Recopy the CoPP profile
	
	Configure a new logging drop threshold for each class

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwe48938	Headline: N9K: "show spanning-tree internal info global" command output truncated "SWOVER Timeout" Symptoms: - `show spanning-tree internal info global` does not show the full output- The output will stop at this line: "SWOVER Timeout (sec)."- This will happen every time when the command is executed once we hit the issue for the first time. Workarounds: None.

CSCwe50502	Headline: Nexus 9K Unexpected Reload due to Watchdog with High "ktah_nl_asic_isr" Interrupts
	
	Symptoms: A Nexus 9K running NX-OS 9.3(9) may reload unexpectedly with a reason of "Watchdog Timeout" due to a high amount of "ktah_nl_asic_isr" hardware interrupt events seen in a kernel panic log. Eg:
	
	<pre>SWITCH# show system reset-reason
	----- reset reason for module 1 (from Supervisor in slot 1) ---
	1) At 431246 usecs after Sun Jan 1 00:00:00 2023
	Reason: Watchdog Timeout
	Service:
	Version: 9.3(9)
	
	SWITCH# show logging onboard stack
	(snip)
	
	**********
	STACK TRACE GENERATED AT Sun Jan 1 00:00:00 2023 UTC
	**********
	(snip)
	
	<0>[XXXXXXXX.XXXXXX] NMI due to BROADWELL FPGA WDT GPIO LEVEL MASK error
	(snip)
	
	<6>[XXXXXXXX.XXXXXX] Dumping interrupt statistics
	<6>[XXXXXXXX.XXXXXX] CPU0 CPU1 CPU2 CPU3 CPU4 CPU5
	CPU6 CPU7 intrs/last_sec max_intrs/se
	(snip)
	
	<6>[XXXXXXXX.XXXXXX] 50: 0 92735896 313 0 0 0
	0 0 655 2871 PCI-MSI-edge ktah nl asic isr</pre>
	
	Workarounds: No workaround. The "watchdog" reset is a recovery mechanism to prevent the switch from becoming unresponsive. Monitor the device for stability. In rare instances, if the issue happens repeatedly, RMA may resolve the issue, but this may not always work as "ktah_nl_asic_isr" is a very generic interrupt counter for any event on the Tahoe ASIC, so at this time, it cannot be assumed this is necessarily a hardware issue. In addition, it is generally expected for that counter to always be non-zero as part of normal uptime of the device.

CSCwe51271	Headline: show process memory should be available for network-operator role Symptoms: show process memory is not available for network-operator role:93240YC-FX2-L1-S1# show processes memory% Permission denied for the role Workarounds: N/A
CSCwe52736	Headline: N9300 NBM related syslog does not comply with standard nxos syslog format Symptoms: NBM error message logged too frequently. Workarounds: N/A

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwe52879	Headline: Syslog %LICMGR-6-LOG_LICAPP_SMART_GLOBAL_CHANGE removal Symptoms: A Nexus device that is enabled with Smart Licensing, but not yet connected to CSLU/CSSM, will log the following message every hour:%LICMGR-6-LOG_LICAPP_SMART_GLOBAL_CHANGE: A global notification change of type 11 is sent out for Application(s)when Smart licensing is used. Workarounds: Change the logging level, message is cosmetic only in nature
CSCwe53655	Headline: Revert reserved MAC blocking behavior for VRRP macs on SVIs Symptoms: User is not able to configure VRRP VMAC on SVI interfaces. Workarounds: None

CSCwe56514	Headline: N9K-C9364C Interface counters is stuck due to EDMA Channel Stuck
	
	Symptoms: Interface counter is stuck and does not increase. Console log would not show any special messages, but "slot 1 show hardware internal tah event-history error" may show the following message:
	
	`2022-12-14T19:16:54.663049000+09:00 [M 1] [tahusd] E_STRING`
	`[tahusd edma handle dma timer event (2267)] inst = 0 EDMA Reserve Channel 6 Stuck`
	
	`2022-12-14T19:16:54.663049000+09:00 [M 1] [tahusd] E STRING`
	`[tahusd edma handle dma timer event (2266)] inst = 0 EDMA Reserve Channel 6 Stuck`
	
	Workarounds: None

CSCwe60434	Headline: n9k: urib core seen leaving only local routes on switch Symptoms: N9k running impacted code keeps running but starts to see the urib process crashing: ``` %KERN-3-SYSTEM_MSG: [6867726.665799] (1:273:1)urib[3871]: rwsem_flush process/thread still exist (1:273:1009)urib:urib-dme-t[3894] - kernel %KERN-3-SYSTEM MSG: [6867726.665814] (1:273:1)urib[3871]: rwsem flush 1 threads still exist - kernel %SYSMGR-2-SERVICE_CRASHED: Service "urib" (PID 3871) hasn't caught signal 6 (core will be saved). ``` The switch remains up but since this process operated the Unicast RIB table, routing information is lost for a short moment. The tables will be repopulated and normal operation should resume. Workarounds: Restart the BGP process to get out of the problem state. Instead of using the non-cached ip route commands, use following CLI through Netconf:show ip route summary cached vrf allshow ipv6 route summary cached vrf all
CSCwe60547	Headline: Memory leak under PIM_MEM_ECMP_REDIR_Q_ENTRY on Nexus9000 Symptoms: - PIM process discloses an unusual amount of memory allocated under PIM_MEM_ECMP_REDIR_Q_ENTRY. Workarounds: 1. Identify any device and flow responsible for hundreds of thousands of ECMP Redir Queue entries using the CLI command "show ip pim internal ecmp-redir queue". 2. Block the problem IP(s) using an ACL caused the PIM_ECMP_REDIR_Q_ENTRY to stop growing.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwe61944	Headline: C93600CD-GX: VPC BGW peer reload might cause up to 20s of traffic Blackhole Symptoms: On a multisite setup reload of VPC BGW peer can cause up to 20s of packet loss. Workarounds: Some setups improve when advertise-pip advertise virtual-rmac is configured.
CSCwe65091	Headline: gnmic reply missing key-value pairs due to PropertyName=<> is_set=0 and is_default_defined=0 Symptoms: the gnmic client call receives a reply but the paylod is missing some key:value pairs that are reported by "show system internal dme running-config all" Workarounds:

CSCwe67205	Headline: Credit Loss Recovery not triggered for FC interface with no transmit credits
	
	Symptoms: A Fibre Channel interface that stays at 0 transmit credits is not recovered by the Credit Loss Recovery agent.
	
	The show interface output shows the affected interface is up but with 0 transmit credits available. Some frames may have been transmitted.
	
	<pre>
	fc1/1 is up
	    ...
	    Port mode is F, FCID is 0x123456
	    ...
	    Operating Speed is 8 Gbps
	    ...
	      3107 frames output,186756 bytes
	        0 discards,0 errors
	    ...
	      0 transmit B2B credit remaining
	      0 low priority transmit B2B credit remaining
	</pre>
	
	The output of slot 1 show logging onboard credit-loss does not contain any recovery events for the interface.
	
	Workarounds: If the interface has switchport ignore bit-errors configured then remove it with the no switchport ignore bit-errors interface configuration command.

CSCwe67953	Headline: N9K: Error or incorrect result when computing multiple file hashes simultaneously Symptoms: The switch may throw an error when multiple file hashes are being computed at the same time by multiple SSH sessions: "ck.out Empty / cksum or md5sum or sha256sum or sha512 calculation Error". Alternately, the wrong hash may be reported for some of the files computed. Workarounds: The only known way to avoid this issue is to wait for File A to hash before starting the hash for File B.
CSCwe67996	Headline: Routes with BGP backup route will not get advertised to a BGP peer Symptoms: BGP route which is not marked best-route in unicast routing table may not get advertised to peer in rare scenario. Workarounds: Clear the IP route,

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwe72406	Headline: Show commands collected via NX-API have missing character with 8k fastcgi buffers. Symptoms: Nexus 93240YC-FX2 using 10.2(2a) is using 8k fastgci buffers by default. In customer switch,when large outputs are requested from Rest, errors are consistently seen where the 8076th characters is missing. Workarounds: There is no current workaround at this time.
CSCwe72834	Headline: N9K BGP peer session stuck in Closing with AF LU prefix-priority high Symptoms: Nexus 9000 with BGP peer session may get stuck in Closing state. Workarounds: Restart the BGP Process or reload the device.
CSCwe74517	Headline: eBGP-Removing template peer-policy under l2vpn EVPN address-family deletes prefixes Symptoms: N9k or N3k 10.2(4)Using eBGP for l2VPN EVPN with rewrite-asn using template peer-policy for l2VPN EVPN neighbor. After removing the template peer-policy under the L2VPN neighbor configuration, the nexus will delete the rewrite-asn from BGP neighbor even though it is hardcoded under the neighbor config. Workarounds: Reconfigure affected neighbor or restart BGP. Note: Restarting the BGP process is disruptive, all the BGP neighbors/BGP routes will bounce, suggestion is to reconfigure the affected BGP neighbor.
CSCwe79884	Headline: Stuck MDIO access with BV ports leads to missed HBs and USD kill. Symptoms: Tahusd process crashes resulting in box reload. Workarounds: NA
CSCwe90801	Headline: Storm-Control CC fail even though Policer/Burst values are correct Symptoms: Storm-Control CC fail even though Policer/Burst values are correct Workarounds: Manually dump commands and verify the values

CSCwe92797	Headline: Nexus 9000 pruning VLANs even when VTP Pruning is not enabled
	
	Symptoms: VLANs in the normal range (2-1001) are pruned off the trunk interfaces after reload
	`--------------------------------------------------------------------------------`
	`Port Vlans in spanning tree forwarding state and not pruned`
	`--------------------------------------------------------------------------------`
	`Eth1/1 none`
	`Eth1/45 none`
	`Eth1/53 none`
	`Eth1/54 none`
	`Po1 2000-2001,3600`
	`Po2 2000-2001,3600`
	
	`Po500 1,4-6,30-32,34-35,74,200-207,209-212,214,216,219-220,224-228,232-240,242-244,246-249,253,271-272,291,298,514,531,533,535,543,583,593,730,732-734,950,2000-2001,3600`
	
	Workarounds: Re-enable feature VTP
	`no feature vtp`
	`feature vtp`
	`vtp domain cisco`

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwe93779	Headline: Object Tracking si not being created when using ip options
	
	Symptoms: When creating a tracking object to associate it with an existing IP SLA is generating an error message.
	
	`test#  config t`
	`Enter configuration commands, one per line. End with CNTL/Z.`
	`test(config)# track 1 ip sla 1 reachability`
	`ERROR: % Object does not exist`
	`test(config)#`
	
	Accounting log shows:
	
	`Tue Apr  4 22:43:30 2023:type=update:id=10.21.66.63@pts/0:user=admin:cmd=configure terminal ; track 1 ip sla 1 reachability (FAILURE)`
	
	`Also, when creating an ip sla and exiting it, an invalid mode is shown.`
	`test# config t`
	`Enter configuration commands, one per line. End with CNTL/Z.`
	`test(config)# ip sla 2`
	`test(config-ip-sla)# exit`
	`## Invalid mode ## #`
	`## Invalid mode ## #`
	`## Invalid mode ## # show ip int brief`
	
	`IP Interface Status for VRF "default"(1)`
	`Interface  IP Address  Interface Status`
	`Vlan100  1.1.1.1  protocol-up/link-up/admin-up`
	`Vlan200  2.2.2.1  protocol-up/link-up/admin-up`
	`## Invalid mode ## #`
	
	Workarounds: Downgrade to 10.1.1, 9.3.11, 9.3.9. Use tracking for interface line-protocol.

CSCwe95715	Headline: N9K: Install epld <image name> module all does not upgrade the system controller in slot 30 Symptoms: Executing the command "install epld bootflash:<image name> module all" does not upgrade the EPLD firmware version of the system controller in slot 30. Workarounds: Running the install command once again specifically for the non-upgraded module resolves the issue:install epld bootflash:<image name> module 30.
CSCwf01120	Headline: N9K FX3 VTEP does not perform VXLAN encapsulation when transmitting GRE packets Symptoms: N9K VTEP does not perform VXLAN encapsulation when transmitting GRE packets Workarounds: Reload N9K.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwf03457	Headline: Auto-complete for VRF name can cause unexpected config changes
	
	Symptoms: When the first letter of a VRF is typed and enter key is pressed, if there is only one VRF starting with that letter, the switch may o may not try to auto-complete the VRF name this depending on the command that is used.
	
	For commands such as "show run vrf" or "vrf member" the auto-complete feature doesn't kick in but for the command "no vrf context" the auto-complete feature kicks in, it means that the switch will delete the VRF starting with that letter instead of showing a message that the VRF doesn't exist, this behavior can lead customer to delete a VRF by mistake causing network disruptions.
	
	* For the following tests, type the first letter of the VRF's name and hit enter. Don't use TAB key.
	
	`------------------`
	`1. Starting with the following VRFs created:`
	`------------------`
	
	`SWITCH(config-if)# show vrf`
	`VRF-Name VRF-ID State Reason`
	`ANOTHERTEST 6 Up -- <<<`
	`default 1 Up --`
	`management 2 Up --`
	
	`------------------`
	`2. Checking the behavior with show run`
	`------------------`
	
	`SWITCH# show run vrf A`
	
	`!Command: show running-config vrf A <<< Auto-complete didn't kick in`
	
	`------------------`
	`3. Assigning a VRF to an interface`
	`------------------`
	
	`SWITCH(config)# inter e1/1`
	`SWITCH(config-if)# vrf member A`
	`Warning: Deleted all L3 config on interface Ethernet1/1`
	`VRF A does not exist. Create vrf to make interface Ethernet1/1 operational <<< Auto-complete feature is not triggered the switch logs a message`
	
	`----------------`
	`4. Delete none-existing VRF "A"`
	`----------------`
	
	`SWITCH(config-if)# no vrf context A <<< "A" vrf doesn't exist but the switch will auto-complete and will delete vrf "ANOTHERTEST"`
	
	`SWITCH(config)# show vrf`
	`VRF-Name VRF-ID State Reason`
	`default 1 Up --`
	`management 2 Up --`
	
	No message was shown for the not existing VRF and due to auto-complete feature, the VRF starting with the letter A was deleted.
	
	Workarounds:
	• Make sure to write all the VRF name while deleting a VRF.
	
	Use TAB key to auto-complete, however this options is not always possible for example when using scripts for automation.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf08533	Headline: Netflow traffic getting dropped under Custom CoPP class-default Symptoms: A N9K-C93108TC-FX3P running 10.3.2 NX-OS, with a custom CoPP configuration may experience a condition where all Netflow packets stop being classified by the hardware rate-limiter and become classified under the custom CoPP class-default class. Most of the Netflow packets are dropped by the strict policer. Workarounds: Remove and reapply the custom CoPP policy:control-plane no service-policy input copp-policy-strict-custom service-policy input copp-policy-strict-custom

CSCwf08661	Headline: vsh.bin service crashed in Nexus 9k
	
	Symptoms: vsh.bin service crashed:
	
	`2023 Apr 11 08:14:50.828 netb1d1-brtr040b %$ VDC-1 %$ %SYSMGR-2-LAST_CORE_BASIC_TRACE: : PID 8735 with message vsh.bin(non-sysmgr) crashed, core will be saved .`
	
	`show cores`
	`VDC  Module  Instance  Process-name  PID  Date(Year-Month-Day Time)`
	`---  ------  --------  ---------------  --------  -------------------------`
	
	`1  1  1  vsh.bin  8735  2023-04-11 08:15:01`
	
	Workarounds: No workaround known so far

CSCwf10110	Headline: Seeing error logs on upgrade: invalid group file entry delete line 'aaa-db-operator:508:'? No Symptoms: Error log during bootupinvalid group file entrydelete line 'aaa-db-operator:508:'? No Workarounds: None needed.
CSCwf13179	Headline: VLAN Filter Allows an ACL with the 'log' keyword to be used when applied to a non-existent VLAN Symptoms: VLANs are suspended on interfaces after a VLAN Filter is applied.- At first the only VLANs affected are the ones identified by the 'vlan-list' component of the VLAN Filter. If the affected interfaces are flapped, the suspension will spread to all VLANs currently hosted on the interfaces that were flapped. The switch will generate logs similar to:- %ETHPORT-3-IF_ERROR_VLANS_SUSPENDED: VLANs 2707 on Interface Ethernet1/2 are being suspended. (Reason: ACL Logging is not supported in egress direction.)- %ETHPORT-5-IF_SEQ_ERROR: Error ("ACL Logging is not supported in egress direction.") communicating with MTS_SAP_SPM for opcode MTS_OPC_ETHPM_PORT_LOGICAL_BRINGUP (RID_PORT: Ethernet1/2) Workarounds: - Remove the VLAN Filter configuration or remove the 'log' keyword from the ACL used by the VACL. After either of the two steps above are done, it is necessary to flap all of the affected interfaces by doing 'shutdown' and then 'no shutdown'.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf17839	Headline: BGP core file generated after using "show l2vpn vpn route x.x.x.x" command Symptoms: There are two issues that the DDTS resolves: 1) BGP core file generated on N9K after using the command: "show l2vpn vpn route x.x.x.x". 2) Asserts seen when displaying BGP routes It was first documented in N9K running 10.2(4) but it can happen in non-patched version. The BGP core is generated with show l2vpn vpn route 0.0.0.0 but it might happen with other routes. It can happen with the following logs when the show command is executed: `2023 Apr 21 19:11:05 sp4.tmp1 %BGP-3-ASSERT: bgp- [16130] ../routing-sw/routing/bgp/bgp path.c:84: Assertion `bgp path base has ext path(path base)' failed.` `2023 Apr 21 19:11:05 sp4.tmp1 %BGP-3-ASSERT: bgp- [16130] -Traceback:XXXXXXXXXXXXX` `2023 Apr 21 19:11:05 sp4.tmp1 %KERN-3-SYSTEM_MSG: [1102.196122] (1:283:1003)bgp:bgp-mts[26675]: rwsem_flush process/thread still exist (1:283:1049)bgp:bgp-cli-tra[32084] - kernel` `2023 Apr 21 19:11:05 sp4.tmp1 %KERN-3-SYSTEM MSG: [1102.196126] (1:283:1003)bgp:bgp-mts[26675]: rwsem flush 1 threads still exist - kernel` `2023 Apr 21 19:11:05 sp4.tmp1 %SYSMGR-2-SERVICE CRASHED: Service "bgp" (PID 16130) hasn't caught signal 11 (core will be saved).` `2023 Apr 21 19:11:11 sp4.tmp1 %BGP-5-ADJCHANGE: bgp- [32090] (default) neighbor fe80::XXXX:XXXX:XXXX:XXXX Up` `<snip>` `2023 Apr 21 19:11:18 sp4.tmp1 %BGP-5-ADJCHANGE: bgp- [32090] (default) neighbor fe80::af3:fbff:fe6a:33fe Up` `2023 Apr 21 19:11:36 sp4.tmp1 %BGP-3-ASSERT: bgp- [32090] ../routing-sw/routing/bgp/bgp_path.c:84: Assertion `bgp_path_base_has_ext_path(path_base)' failed.` `2023 Apr 21 19:11:36 sp4.tmp1 %BGP-3-ASSERT: bgp- [32090] -Traceback: XXXXXXXXXXXXX` `2023 Apr 21 19:11:37 sp4.tmp1 %SYSMGR-2-SERVICE CRASHED: Service "bgp" (PID 32090) hasn't caught signal 11 (core will be saved).` `2023 Apr 21 19:11:43 sp4.tmp1 %BGP-5-ADJCHANGE: bgp- [1314] (default) neighbor fe80::XXXX:XXXX:XXXX:XXXX Up` Workarounds: Solve the BGP assert error logs before using "show l2vpn vpn route x.x.x.x" command.

CSCwf19968	Headline: vsh.bin fails after setting a SPAN capture with thousands of source vlans on a single session
	
	Symptoms: When configuring thousands of source vlans the device got unresponsive, the switch hangs and closes the ssh session, then the vsh.bin process reloads unexpectedly:
	
	`switch(config-monitor)# source vlan 31-3999`
	`^`
	`reserved vlans 3968-3999 at '^' marker.`
	
	`switch(config-monitor)# source vlan 31-3967`
	`Warning: Tx Vlan Span is not supported`
	
	`2023 Mar 8 10:10:07 switch %SYSMGR-2-LAST CORE BASIC TRACE: : PID 19845 with message vsh.bin(non-sysmgr) crashed, core will be saved .`
	
	`switch# show cores`
	`VDC Module Instance Process-name PID Date(Year-Month-Day Time)`
	`--- ------ -------- --------------- -------- -------------------------`
	
	`1 1 1 vsh.bin 19845 2023-03-08 10:10:07`
	
	Workarounds: Avoid setting a higher amount of source vlans as the supported, from the bug CSCwf16220, it is documented the SPAN has a limit for the total source vlans supported on a single session which is 32 vlans.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwf20782	Headline: Nexus 9000 - N2K HIF interface configurations erased while upgrading to NX-OS 10.2(4)
	
	Symptoms: During upgrading Nexus 9000 disruptively from NX-OS 9.3(11) --> 10.2(4), Configurations on all N2K HIF interface configurations partially erased resulting in loss of traffic.
	
	- Configurations on all N2K HIF interface configurations partially erased.
	- Interfaces were in inactive status, As 'switchport access vlan x' was missing from HIF interfaces.
	
	Before the Upgrade:
	`----------`
	`interface Ethernet102/1/17`
	`  description xxxxx`
	`  switchport access vlan x`
	`  spanning-tree port type edge`
	`  no shutdown`
	
	After the Upgrade:
	`----------`
	`interface Ethernet102/1/17`
	`  description xxxxxx`
	`  switchport`
	`  no shutdown`
	
	- Added missing configuration to problematic interfaces and the CLI was accepting missing configurations but
	configurations were not reflected in running-configuration.
	
	Workarounds: Temporary workaround is to default the affected interfaces and reconfiguring it.

CSCwf21483	Headline: Block `evpn multisite dci-tracking` and `evpn multisite fabric-tracking` on BGW SVI
	
	Symptoms: The CLI `evpn multisite dci-tracking` and `evpn multisite fabric-tracking` are able to configure on the SVI interface. The only supported uplinks in VXLAN are on physical interface and this CLI should only exist on uplink ports. Fabric tracking should be on uplink ports in the site-internal network and the dci-tracking CLI should be on the uplink ports between sites (site external). The ability to add this command to an SVI even when the ports may be disconnected or admin down, can lead the switch to presume that the ports are back to the site internal or site external or are active and working. The Border-Gateway's isolation mechanism functions by recognizing when the tracking ports are down and by bringing down the loopback for the NVE. With this behavior the NVE remains up.
	
	`Site2-BGW1# show nve multisite dci-links`
	`Interface  State`
	`---------  -----`
	`Vlan144  Up  <<<< SVI added to DCI tracking makes BGW think DCI connectivity is still present`
	
	`Ethernet1/2  Down`
	
	Workarounds: Do not configure BGW SVIs with DCI-tracking or fabric tracking configuration

CSCwf21554	Headline: NXOS: Special character ">" causing issues with clear-text key-string in key chain Symptoms: Using a key-string with ">" character may cause the parser to not capture the string, resulting in an empty string and missing configuration. Workarounds: Avoid using the ">" character in a clear-text (pre encrypted) key-string configuration.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf32021	Headline: Debug message of PTP cause crash Symptoms: PTP process crashes with a backtrace. Workarounds: Disable PTP using no feature ptp command.
CSCwf32330	Headline: Enabling "feature intersight" on 10.2(4) or 10.2(5) leaves TCP Port 9002 open to telnet and http Symptoms: Following an upgrade any release prior to 10.2(4) or 10.2(5) which is not the 10.3(x) train, one can see that a new port 9002 is open and listening. It may be detected via security scan, and one can telnet to that port or start an HTTP connection to it. When checking for any changes made on the switch following the upgrade to see why the port might be open, you will find that feature intersight has been enabled by default. Workarounds: Disable feature intersight via [no feature intersight]- Upgrade to 10.3(x)
CSCwf34746	Headline: Configuring track in role-interface mode can cause vsh crash Symptoms: Configuring track in role-interface mode as below can cause vsh crash.N9K# conf tN9K(config)# role name eem-roleN9K(config-role)# interface policy denyN9K(config-role-interface)# permit interface Ethernet1/8N9K(config-role-interface)# track 8 interface Ethernet1/8 line-protocol->ssh crashed Workarounds: Configure track in global configuration mode.

CSCwf36533	Headline: N9K - VPC ports stuck into STP BLK state on secondary VPC after vPC Fabric Peering sync is recovered
	
	Symptoms: When Spine drops the vPC peer communication, Nexus vPC peer-link detected BPDU receive timeout, and will block VPC port-channel. After vPC peer communication is restored, vPC port-channel is stuck into STP BLK state on secondary vPC.
	
	EVIDENCE:
	
	- On secondary Nexus VPC, output from "show logging logfile" shows that vPC port-channel was unblock:
	
	`switch# show logging logfile`
	`<snip>`
	`%STP-2-VPC_PEER_LINK_INCONSIST_UNBLOCK: vPC peer-link inconsistency cleared unblocking port-channel100 MST0001.`
	`%STP-2-VPC PEER LINK INCONSIST UNBLOCK: vPC peer-link inconsistency cleared unblocking port-channel1 MST0001.`
	
	- Output from "show spanning-tree" show that vPC peer-link port-channel is in Forwarding state, but the vPC port-channel still in Blocking state:
	
	`switch# show spanning-tree mst`
	`<snip>`
	`Interface Role Sts Cost Prio.Nbr Type`
	`---------- ---- --- --------- -------- --------------------------------`
	`Po1 Root FWD 20000 128.4096 (vPC peer-link) Network P2p <<<<<<<< Peer-link in Forwarding (FWD) state`
	
	`Po100 Desg BLK 200 128.4195 (vPC) Edge P2p`
	`<<<<<<<<<<<<<<<<<<<< VPC port in Blocking (BLK) state`
	
	Workarounds: To recover, need to shut/no shut the port that is on BLK state.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf37901	Headline: VXLAN VLANs suspended by vpc consistency checker due to different replication mode and VNI type Symptoms: VXLAN vlans are getting suspended due to different replication mode, but configuration shows that replication mode is the same on both vPC switches. Workarounds: Remove and add back affected VNI.
CSCwf37914	Headline: Reduce severity of syslog when mc-drop command is enabled Symptoms: After enabling "hardware qos pf mc-drop" we start to see a lot of sev2 messages that are strictly informational flood the logs. Workarounds: Change global logging level (not recommended)Logging discriminator
CSCwf39373	Headline: OSPFv3 is adding a link-local forwarding address for NSSA type 7 originated routes Symptoms: While a user redistribute a static route into the device in the Not-so-stubby Area (NSSA) [NSSA is an OSPF Stub Area, which can carry routes learned by other protocols such as BGP or RIP and here Static in this case] and making the current OSPF router as a border router. User is observing that in the database that the forwarding address is link local and not any of the available global addresses on ospfv3 enabled interfaces in the area. This seems inconsistent with the RFC which states that we should not use link local address in the NSSA LSA forwarding address. Workarounds: Configure a loopback with a global ipv6 address in the NSSA area to pick a global address from.
CSCwf42887	Headline: N9K -FX3: VXLAN storm-control policer fabric bandwidth does not update after fabric link flap Symptoms: N9K -FX3: VXLAN storm-control policer fabric bandwidth does not update after fabric link flap Workarounds: Shut/No-shut multi-site loopback- Will reset policer but not stop the behaviorDisable evpn storm-control
CSCwf47425	Headline: Ports of C93600CD-GX with QSA (10G) don't forward traffic Symptoms: 10G port with QSA adapter in up state, but doesn't forward traffic Workarounds:
CSCwf48266	Headline: Debug hardware internal command on N9K-C93360YC-FX2 causes tahusd crash Symptoms: N9K-C93360YC crashes after entering debug commands while in the module shell, generating a tahusd core Workarounds: For any debugging in the module shell, please contact TAC
CSCwf48692	Headline: N9300-FX3S/FX3 may randomly timeout syncE peer Symptoms: Random syncE peer timeouts. Workarounds: N/A

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwf50018	Headline: Fibre Channel snmp trap configuration causing CLI to hang
	
	Symptoms: CLI hangs after entering some snmp traps related to Fibre Channel. If command is aborted with ctrl + c, message "Failed to collect returned stimulus" is observed. After this, other configuration changes not necessarily related to FC experience the same issue.
	`switch(config)# snmp-server enable traps fcdomain dmDomainIdNotAssignedNotify`
	`^CFailed to collect returned stimulus`
	`switch(config)#`
	
	`switch(config)# feature interface-vlan`
	`^CFailed to collect returned stimulus`
	`switch(config)#`
	
	`switch(config)# vlan 10`
	`^V^CFailed to collect returned stimulus`
	
	`switch(config)#`
	
	Workarounds: Reloading the switch clears the issue; however, if snmp command is re-entered, the issue will come up again.

CSCwf50388	Headline: tahusd crash due to InPhi retimer quad port dead lock Symptoms: Tahusd reloaded due to software mutex lock. Last reset at 22466 usecs after Sat May 20 11:29:29 2023 Reason: Reset Requested due to Fatal Module Error System version: 9.3(9) Service: tahusd hap reset Workarounds: N/A, the chassis will silently reload and recovers post reload.As part of the proactive action to avoid similar issue, to ensure users do not see this issue (we can only reduce probability and not eliminate it), is to keep all links that unused in admin shut state.
CSCwf54392	Headline: N9K: "radius-server key <>" missing after ND ISSU Symptoms: When upgrading a Nexus 9000 switch from a code where LXC boot mode is not the default mode, to one where it is, using non-disruptive ISSU, the "radius-server key 7 <string>" configuration can go missing. This is due to to the change in default boot mode. Workarounds: Remove and reapply the missing CLI string to reconfigure and ensure it is consistent between configuration and DME.
CSCwf55892	Headline: N9500-G - Total Power Allocated (budget) is incorrect calculation Symptoms: The sum of individual module power in "show environment power" isn't equal to total power allocated (budget) value. Workarounds: None
CSCwf56954	Headline: N9K Reset-Reason is not showing correctly after module power cycle Symptoms: A Nexus 9000 platform may display the reset reason incorrectly after has been power cycled. Workarounds: None

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf57648	Headline: Nexus 9500 -R modules incorrect outer DMAC after initialization Symptoms: Nexus 9500 with linecards to perform a MPLS to VxLAN handoff. The Nexus 9500 acting as the PE device is adding an incorrect DMAC to the outer VxLAN header, which is causing the downstream VxLAN leaf to drop the packet. The incorrect MAC observed is "00:D0:00:00:00:88". Workarounds: If the switch is found to be in this state a "reload ascii" will fix the mis-programming. We have also seen flapping the BGP neighbors correct the issue, as this is related to the advertising and programming of the next-hop VTEPs RMAC. Another workaround is to ?poweroff module <>? and ?no poweroff module <>? of the problematic Linecards to recover from the issue.
CSCwf58507	Headline: FEX 2348UPQ brings hosts link too fast after powercycle causing traffic blackholing for around 1min. Symptoms: when FEX 2348TQ power-cycles all hosts using (GLC-T, GLC-T-C,SFP-H10GB-CU3M,SFP-H10GB-CU3M) SFPs will have link up while the FEX is down. after ~5 seconds links go up for around 1 minute, despite the fact that FEX from N9k perspective and its FIs are down. Workarounds: Use LACP or any other SFP or upgrade the version.
CSCwf60819	Headline: n9k/NXOS - QOS policer/CIR - traffic rate policed 1% below configured CIR Symptoms: On N9K -GX switches, the effective policing rate for NBM flows is observed to be about 1% lower than the configured value. This might lead to traffic drops. Workarounds: None
CSCwf61588	Headline: %NFM-1-RTP_FLOW_ERROR_DETECTED - CLI execution slowness seen Symptoms: General slowness in command execution in CLI may be seen if there are multiple "%NFM-1-RTP_FLOW_ERROR_DETECTED:" syslogs are generated. This is because if large unique flows present, for every flow that become from lossless to lossy, a new syslog is logged. Same is true for lossy to lossless. Workarounds: No workarounds.
CSCwf61602	Headline: duplicate bia seen on N9K - GX Symptoms: On 93600CD-GX switch, BIA block is reused after port 1/25. So there will duplicate bia mac. See output below. ports 1-24 gets unique BIA mac, then the same set of BIA macs are assigned for the next set of ports 25-48. LEAF35-PMN-SOLN-WOLFRIDGE(config)# sh system internal ethpm info all | inc "Backplane MAC address in GLDB" | sort | uniq -c | excl '1 Backp' 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:68 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:70 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:78 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:80 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:88 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:90 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:98 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:a0 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:a8 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:b0 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:b8 2 Backplane MAC address in GLDB: 10:b3:d6:bc:54:c0 LEAF35-PMN-SOLN-WOLFRIDGE(config)# Workarounds: None

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwf63078	Headline: 'ip dhcp relay' commands are not retained on SVI after changing VRF membership
	
	Symptoms: 'ip dhcp relay address' commands are no longer present in the SVI after changing the VRF membership, even though the 'system vrf-member-change retain-l3-config' command is globally configured.
	
	N9K-1(config-if)# show run int vlan 10
	
	!Command: show running-config interface Vlan10
	!Running configuration last done at: Wed Mar 29 20:01:02 2023
	!Time: Wed Mar 29 20:01:08 2023
	
	version 10.3(3) Bios:version 08.39
	
	interface Vlan10
	  no shutdown
	  vrf member A
	  ip address 10.1.1.1/24
	  ip pim sparse-mode
	  ip pim border
	  ip pim hello-interval 300000
	  ip dhcp relay address 129.132.128.135
	  ip dhcp relay address 129.132.128.199
	
	N9K-1(config-if)#
	
	N9K-1(config-if)# vrf member B
	Warning: Please ensure igmp snooping is enabled on the corresponding vlan, using the command
	'show ip igmp snooping vlan <x>'. IGMP snooping MUST be enabled for correct forwarding
	functionality.
	
	N9K-1(config-if)# end
	
	[+] After changing the VRF membership for Vlan10 the 'ip dhcp relay' commands are gone
	
	N9K-1# show run int vlan 10
	
	!Command: show running-config interface Vlan10
	!Running configuration last done at: Wed Mar 29 20:01:42 2023
	!Time: Wed Mar 29 20:01:48 2023
	
	version 10.3(3) Bios:version 08.39
	
	interface Vlan10
	  no shutdown
	  vrf member B
	  ip address 10.1.1.1/24
	  ip pim sparse-mode
	  ip pim border
	  ip pim hello-interval 300000
	
	Workarounds: Apply back to the affected SVI the missing 'ip dhcp relay address' commands.
	
	N9K-1# conf t
	Enter configuration commands, one per line. End with CNTL/Z.
	N9K-1(config)# interface vlan 10
	N9K-1(config-if)# ip dhcp relay address 129.132.128.135
	N9K-1(config-if)# ip dhcp relay address 129.132.128.199
	N9K-1(config-if)# end
	N9K-1# show run int vlan 10
	
	!Command: show running-config interface Vlan10
	!Running configuration last done at: Wed Mar 29 20:02:06 2023
	!Time: Wed Mar 29 20:02:32 2023
	
	version 10.3(3) Bios:version 08.39
	
	interface Vlan10
	  no shutdown
	  vrf member B
	  ip address 10.1.1.1/24
	  ip pim sparse-mode
	  ip pim border

Bug ID	Description
CSCwf64467	Headline: Incorrect config-profile configuration cannot remove corresponding entry in vsh config Symptoms: When entering a command with an incorrect network in a config profile and applying the profile to the main configuration, the profile cannot be unapplied. For example, when applying a config profile with a line like this:ip prefix-list PL-1 permit 10.0.0.10/24 => the correct network address would be 10.0.0.0/24in the "show run", the address will configuration will correctly be:ip prefix-list PL-1 permit 10.0.0.0/24 but the config profile cannot be unapplied Workarounds: None

CSCwf64695	Headline: 'ip dhcp relay' commands are not getting applied
	
	Symptoms: 'ip dhcp relay address' commands fail to be applied under the SVI, even though the 'show accounting log' will say the commands were successfully executed. However, DME consistency-checker will fail.
	
	`N9K-2(config-if)#  ip dhcp relay address 10.1.1.100`
	`N9K-2(config-if)#  ip dhcp relay address 10.1.1.200`
	`N9K-2(config-if)# end`
	`N9K-2# show run int vlan 10`
	
	`!Command: show running-config interface Vlan10`
	`!Running configuration last done at: Tue Mar 21 19:03:51 2023`
	`!Time: Tue Mar 21 19:03:55 2023`
	
	`version 10.2(4) Bios:version 08.39`
	
	`interface Vlan10`
	`  no shutdown`
	`  vrf member B`
	`  ip address 10.10.10.2/24`
	`  ip pim sparse-mode`
	`  ip pim border`
	`  ip pim hello-interval 300000`
	
	`N9K-2#`
	
	`Tue Mar 21 19:03:50 2023:type=update:id=console0:user=admin:cmd=configure terminal ;`
	`interface Vlan10 ; ip dhcp relay address 10.1.1.100 (REDIRECT)`
	`Tue Mar 21 19:03:50 2023:type=update:id=console0:user=admin:cmd=configure terminal ;`
	`interface Vlan10 ; ip dhcp relay address 10.1.1.100 (SUCCESS)`
	`Tue Mar 21 19:03:51 2023:type=update:id=console0:user=admin:cmd=configure terminal ;`
	`interface Vlan10 ; ip dhcp relay address 10.1.1.200 (REDIRECT)`
	`Tue Mar 21 19:03:51 2023:type=update:id=console0:user=admin:cmd=configure terminal ;`
	`interface Vlan10 ; ip dhcp relay address 10.1.1.200 (SUCCESS)`
	
	Workarounds: Apply the 'no' form of the missing commands, then reapply manually

CSCwf67122	Headline: MSDP instability if VRF not enabled for BGP Symptoms: MSDP performance issues may be observed if it's running in a VRF that's not enabled in BGP while BGP is configured globally.- Adjacencies may flap due to keepalives expiring, while packet captures indicate no drops- SA propagation may be significantly delayed, or SAs may never be received- The receive queue for MSDP TCP sockets in "show sockets connection" would be full. Workarounds: Enable the MSDP VRF under the BGP process. No additional configuration such as address families or neighbors is required. Router bgp 65500 vrf MSDP-VRF. Alternatively, disable BGP entirely:no router bgp 65500

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf67373	Headline: Configuration update aborted: request was aborted, post 10.4.1 bin to upg ND-ISSU Symptoms: "copy r s" command is aborted after ND-ISSU from older releases to OR1F with fex. The layer for fex ports are inconsistent across DME and backend. Since DME has layer 2 for fex ports it was trying to configure vlans and in backend vlan_mgr was rejecting the config since it was L3 in backend. The correct layer was L3.Since nxapi retries are happening in continuous loop, "copy r s" command is getting aborted. Workarounds: Configuring swithcport/no switchport on all affected interfaces should recover the switch from issue scenario.
CSCwf69556	Headline: Nexus 9000: Interface description including string "%n" or "%ln" crash Service "port-profile" Symptoms: An error occurs when we put the string "%n" or "%ln" as interface description. ----- switch(config-if)# interface Ethernet1/5 switch(config-if)# description % Inter XXXX-1001-XX % 2023 Jun 8 14:44:58 switch %$ VDC-1 %$ %SYSMGR-2-SERVICE_CRxSHED: Service "port-profile" (PID 16609) hasn't caught signal 11 (core will be saved). 2023 Jun 8 14:44:58 switch %$ VDC-1 %$ %USER-2-SYSTEM MSG: ssnmgr app init called on ssnmgr up - aclmgr ----- You could see same issue with following strings as well. description %In description % In description %n description % n After repeating to configure description with these strings several times, NXOS was reloaded. Workarounds: Do not use "%n" or "%ln" in interface description.

CSCwf70004	Headline: Several core files for the process sysmgr-cserver appear in the N9K.
	
	Symptoms: The following logs will appear in the show log which mean that several cores are created.
	
	<pre>
	SYSMGR-2-LAST_CORE_BASIC_TRACE: : PID 18491 with message sysmgr-cserver(non-sysmgr) crashed, core will be saved.
	SYSMGR-2-SERVICE CRASHED: Service "System Manager (core-server)" (PID xxxxx) hasn't caught signal 11 (core will be saved).
	</pre>
	
	And you can verify this with the output of show cores where the sysmgr-cserver will be:
	<pre>
	`show cores`
	VDC Module Instance Process-name PID Date(Year-Month-Day Time)
	--- ------ -------- --------------- -------- ------------ -------------
	1 1 1 sysmgr-cserver xxxx date
	1 1 1 sysmgr-cserver xxxx date
	1 1 1 sysmgr-cserver xxxx date
	
	</pre>
	
	Workarounds: None

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf70688	Headline: N9K: NX-API ospfAdjEp.json cannot retrieve OSPF neighbors correctly after interface flapping Symptoms: NX-API ospfAdjEp.json cannot retrieve OSPF neighbors correctly after interface flapping. For example, when the ospf processes have 10 neighbors and the response from the "show ip ospf neighbors vrf all" command always lists 10 neighbors, but the request via ospfAdjEp.json lists only 9 neighbors. Workarounds: Use "/api/mo/sys/ospf.json?query-target=subtree&target-subtree-class=ospfAdjEp" instead of "/api/class/ospfAdjEp.json"
CSCwf75767	Headline: Syslog to enable layer3 peer-router seen, even with no routing enabled on VPC VLAN's Symptoms: • Below log is being seen in # show logging log, even if no routing is enabled on VPC VLAN's • Suppressing syslog by "no layer3 peer-router syslog" has no impact"ipfib: Routing adjacency over vPC detected without required configuration. Please configure layer3 peer-router under the vPC domain. See https://cisco.com/go/l3pr for more information." Workarounds: • Enable "layer3 peer-router " under VPC configuration on both VPC peers Minor cosmetic thing, can be ignored
CSCwf75862	Headline: Bi-dir traffic received on DF winner interface is dropped towards the RPF (RP) Symptoms: Bi-dir traffic received on DF winner interface is dropped towards the RPF (RP)All consistency checker are cleanElam summary shows forwardElam detail shows RPF failure Workarounds: NA at the time of opening the defect
CSCwf77827	Headline: N9300 inconsistent ACE sequence number check causing NDB switch to be in NOTREADY state Symptoms: NDB nexus switch is in NOTREADY state after NDB controller process restart. Workarounds: Reload nexus NDB switch or delete all NDB acls manually.

CSCwf79132	Headline: After ISSU upgrade performed with maintenance mode SVIs stay down
	
	Symptoms: On Nexus 9000 all SVIs remain down after exiting maintenance mode post ISSU upgrade.
	
	Example VLAN interface in problematic state:
	`nexus9000# show interface vlan 3600`
	`Vlan3600 is down (suspended), line protocol is down, autostate enabled`
	
	`nexus9000# show system internal eltm event-history vxlan-error`
	`2023-06-26T06:27:59.067882000+00:00 [M 1] [eltm] E DEBUG Outer Bd Alloc failed Nve ifx 0x1a005e00, Overlay state = Up v4/v6 tid (0x1 80000001 1 0) en (0 0)`
	
	`2023-06-23T13:39:01.222605000+00:00 [M 1] [eltm] E DEBUG Failed to force alloc outer bd 16 for vxlan infra vlan 3600`
	
	Workarounds: Rebooting the switch restores the SVIs.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwf81820	Headline: EvtLog: Evtlog decoder not working
	
	Symptoms: Decoding of blogger log-snapshot does not work in 10.2(5) releases:
	`agargula@ag-ubuntu-ML1:~/snapshot-bin$ ./NexusEvtLogDecoder.10.2.5.I9.1.py -input ../snapshot.tar`
	`ERROR:root:Parsing Decode Table failed with exception(Expecting object: line 226341 column 1 (char 20287150))`
	
	`ERROR:root:An error ('NoneType' object has no attribute 'keys') occurred during decode operation: /tmp/BLOGGERD0.939064446123/1-230622215047-230622215705-720109-Caaa-U181-M1-V1-I0-0-P13707-messages.gz @ ../snapshot.tar/20230622215659524263 evtlog archive.tar`
	
	Workarounds: Use decoder from 10.4(1) once available

CSCwf82223	Headline: On N9K switch SNMP polling unable to retrieve QSFP DOM values. Symptoms: A third patry DOM supported QSFP will not provide any DOM information when polled via SNMP. Workarounds: Physically reset QSFP28. OIR.
CSCwf84373	Headline: Admin shut the interface "mgmt0", other end port remains up/flapping. Symptoms: Admin shut interface "mgmt0", other end port remains up/flapping. Workarounds: No workaround.
CSCwf85413	Headline: ICAM scale table values inconsistent between IPv4 and IPv6 on 9364C Symptoms: The N9K-C9364C switch configured with the LPM heavy template shows warning log messages like %ICAM-4-SCALE_THRESHOLD_EXCEEDED_WARN: Utilization of 97 percent for feature IPv4 LPM routes is over the warning threshold even when the IPv4 routing table size is relatively small for such template, e.g. 14k routes. Workarounds: Add the following line to the configuration**icam monitor scale unicast-routing routing lpm-route-v4 limit 262000**

CSCwf86821	Headline: ACL not programming in hardware with udp-relay config on Nexus 9000
	
	Symptoms: ACL entries not programming in hardware with udp-relay config on Nexus 9000 when adding under the SVI.
	`slot 1`
	`=======`
	
	`ERROR: no ACL related hardware resources for vdc [1], vlan [210]`
	
	Workarounds: Doing the following commands with same order under the SVI:switch #(config-if) no ip udp relay addrgroup <object-group name>switch #(config-if) ip udp relay addrgroup <object-group name>

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description

CSCwf88604	Headline: EPBR crash @iscm_parse_ipv6_ace upon epbr ipv6 policy configuration with object groups
	
	Symptoms: 'epbr' process crashes are observed upon configuring an ebpr ipv6 policy on an interface on a Nexus 9000.
	The process crash triggers a reload of the Nexus 9000 switch. Following the reload the following information reports the reason of the crash.
	
	`N9K# show system reset-reason`
	`----- reset reason for module 1 (from Supervisor in slot 1) ---`
	`1) At 582169 usecs after Sat Jul 8 13:48:28 2023`
	`Reason: Reset triggered due to HA policy of Reset`
	`Service: epbr hap reset`
	`Version: 10.3(3)`
	
	`N9K# show core`
	`VDC Module Instance Process-name PID Date(Year-Month-Day Time)`
	`--- ------ -------- --------------- -------- ------------------------`
	`1 1 1 epbr 28608 2023-07-08 12:50:22 *`
	`1 1 1 epbr 6054 2023-07-08 12:50:22 *`
	`1 1 1 epbr 5998 2023-07-08 12:53:58`
	
	`* HAP reset core restored on reload from previous boot`
	
	Workarounds: Do not use IPv6 object-group in EPBR policy ACL match condition. Replace the object-group with standard ACE entry in the IPv6 access-list configuration.

CSCwf97134	Headline: Several core files for the process sysmgr-cserver appear in the N9K.
	
	Symptoms: The following logs will appear in the show log which mean that several cores are created.
	
	`<pre>`
	`SYSMGR-2-LAST_CORE_BASIC_TRACE: : PID 18491 with message sysmgr-cserver(non-sysmgr) crashed, core will be saved.`
	`SYSMGR-2-SERVICE CRASHED: Service "System Manager (core-server)" (PID xxxxx) hasn't caught signal 11 (core will be saved).`
	`</pre>`
	
	`And we can verify this with the output of show cores where the sysmgr-cserver will be:`
	`<pre>`
	`` `show cores` ``
	`VDC Module Instance Process-name PID Date(Year-Month-Day Time)`
	`--- ------ -------- --------------- -------- ------------ -------------`
	`1 1 1 sysmgr-cserver xxxx date`
	`1 1 1 sysmgr-cserver xxxx date`
	`1 1 1 sysmgr-cserver xxxx date`
	
	`</pre>`
	
	Workarounds: None

CSCwf97335	Headline: N9k | Host routes transition with Punt index in v4host table after route change Symptoms: Multiple Symptoms may be reported. 1) Traffic to certain IP addresses may end up having high response times or 2) Intermittent packet loss to certain HostsCoPP Drops may be seen depending upon the traffic rate for the class- copp-system-p-class-l3uc-dataData plane traffic to these impacted hosts are software switched/punted and is visible via Ethanalyzer. Workarounds: Introduce back the AM or /32 Host routes

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf98148	Headline: Port-security static binding gets stuck if the interface is converted to L3 before deleting config Symptoms: if the interface is configured with a port-security static binding and then shut down before converting it to L3 using "no switchport" command, the static entry will continue to be shown statically associated to the interface and the only way to remove this static entry is to disable the port-security feature using "no feature port-security". If the static binding is removed before converting the interface to L3, or if the "default interface" command is entered the issue is not seen. Workarounds: • Do not shut down the interface before converting it to L2. • Default the interface before converting it to L3. • Remove the static entry with the no command before converting the interface to L3. If the entry is stuck, disable and reenable the port-security feature.
CSCwf98194	Headline: PBR in combination with a NULL 0 route not working properly , we see drops (UC_PC_CFG_TABLE_DROP) Symptoms: The issue is seen when we have a NULL 0 route present for the destination in combination with PBR which is causing the traffic to be dropped. Workarounds: If the NULL 0 route to the destination is removed no drop is seen
CSCwf98753	Headline: Nexus 9000: Redistributed Routes Are Not Removed from Routing Protocol when deleted from URIB. Symptoms: A prefix that is redistributed into a routing protocol is stuck in its database even though the route is no longer in URIB. Workarounds: - Configure a Null0 route for the affected prefix and then remove it. ip route 0.0.0.0/0 null0no ip route 0.0.0.0/0 null0
CSCwh11140	Headline: Remote Address not sent in TACACS Authorization packet when using NETCONF over SSH Symptoms: Remote Address not sent in TACACS Authorization packet when using NETCONF over SSH Workarounds: N/A
CSCwf25135	Headline: MSDP peer flapping Symptoms: MSDP performance issues may be observed if it's running in a VRF that's not enabled in BGP while BGP is configured globally.- Adjacencies may flap due to keepalives expired, while packet captures indicate no drops- SA propagation may be significantly delayed, or SAs may never be received- The receive queue for MSDP TCP sockets in "show sockets connection" would be full. Workarounds: Enable the MSDP VRF under the BGP process. No additional configuration such as address families or neighbors is required. Router bgp 65500 vrf MSDP-VRF. Alternatively, disable BGP entirely:no router bgp 65500.

© 2023 Cisco and/or its affiliates. All rights reserved.

Bug ID	Description
CSCwf62452	Headline: Mac learning issue after reload with private vlan configuration. Symptoms: MAC learning is not happening after reload for private VLAN Workarounds: 1st Workaround – 1. Delete the Primary and Secondary VLAN 2. Delete the Interface level PVLAN config, in Lab setup, it is Eth1/3 3. Add VLAN configuration back 4. Add Interface configuration Back 5.MAC address will be learned 2nd Workaround – Add static Mac.
CSCwf61686	Headline: Nexus 9500: Configuring "feature nv overlay" breaks non-VXLAN VLAN Multicast across FM-G modules Symptoms: Multicast across Fabric Modules is dropping silently. Inter-VLAN and intra-VLAN multicast across FMs are affected. Only Multicast received on Trunk ports should be affected. Ingress L3 or access ports seems to be unaffected. Multicast flows where the Source and Receiver are connected to the same line card do not appear to be affected since they bypass the Fabric Modules. Workarounds: Disable "feature nv overlay" if possible. The ingress port can be changed to an Access, L3, or L3 Sub interface. Otherwise, install non-FM-G Fabric Modules.
CSCwh04496	Headline: SNMP reloaded unexpectedly without collecting the main thread on the core file Symptoms: The issue was seen for first time on a Nexus N9K-C9508 running in 10.3(2).The failure was reported by SNMP, the device did not reloaded but the process .%SYSMGR-2-SERVICE_CRASHED: Service "snmpd" (PID 3678) hasn't caught signal 6 (core will be saved). Workarounds: Disable SNMP by using the "no snmp-server protocol enable" command.
CSCwh30962	Headline: NXOS - BGP Graceful Restart Helper ignores BFD down event when TCP FIN is received from restarting peer. Symptoms: When reload is issued on ASR1K a TCP FIN is sent out for BGP session termination. BFD also goes down between peers, however this is being ignored by BGP and Graceful Restart is not terminated. Workarounds: Shutdown BGP neighbor prior to the reload.
CSCwh30104	Headline: N9k/icam system monitor - history do not have all entries present Symptoms: The icam monitoring system is missing history values for some processes. show icam system <...> history XXX will not show history for all processes. Workarounds: Manually collect "show icam system" and save on external NMS or monitor processes memory by using the show CLI.
CSCwh32362	Headline: Evora crash when incorrect evora register is given. Symptoms: 9732C-EXM card may reload when doing data collection. Workarounds: Ensure the correct register is issued by using the show command.

© 2023 Cisco and/or its affiliates. All rights reserved.

Known Issues

Bug ID	Description
CSCwi99525	On Cisco Nexus N2K-C2348TQ HIFs fail to utilize redundant Port-Channel links, to NIF, during link failover events.

Device Hardware

The following tables list the Cisco Nexus 9000 Series hardware that Cisco NX-OS Release 10.2(6)M supports. For additional information about the supported hardware, see the Hardware Installation Guide for your Cisco Nexus 9000 Series device.

Table 1. Cisco Nexus 9500 Switches

Product ID	Description
N9K-C9504	7.1-RU modular switch with slots for up to 4 line cards in addition to two supervisors, 2 system controllers, 3 to 6 fabric modules, 3 fan trays, and up to 4 power supplies.
N9K-C9508	13-RU modular switch with slots for up to 8 line cards in addition to two supervisors, 2 system controllers, 3 to 6 fabric modules, 3 fan trays, and up to 8 power supplies.
N9K-C9516	21-RU modular switch with slots for up to 16 line cards in addition to two supervisors, 2 system controllers, 3 to 6 fabric modules, 3 fan trays, and up to 10 power supplies.

Table 2. Cisco Nexus 9500 Cloud Scale Line Cards

Product ID	Description	Maximum Quantity		
		Cisco Nexus 9504	Cisco Nexus 9508	Cisco Nexus 9516
N9K-X9716D-GX	Cisco Nexus 9500 16-port 400-Gigabit Ethernet QSFP line card	4	8	N/A
N9K-X9736C-FX	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet QSFP28 line card	4	8	16
N9K-X9788TC-FX	Cisco Nexus 9500 48-port 1/10-G BASE-T Ethernet and 4-port 40/100 Gigabit Ethernet QSFP28 line card	4	8	16
N9K-X97160YC-EX	Cisco Nexus 9500 48-port 10/25-Gigabit Ethernet SFP28 and 4-port 40/100 Gigabit Ethernet QSFP28 line card	4	8	16
N9K-X9732C-FX	Cisco Nexus 9500 32-port 40/100 Gigabit Ethernet QSFP28 line card	4	8	16
N9K-X9732C-EX	Cisco Nexus 9500 32-port 40/100 Gigabit Ethernet QSFP28 line card	4	8	16
N9K-X9736C-EX	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet	4	8	16

© 2023 Cisco and/or its affiliates. All rights reserved.

Product ID	Description	Maximum Quantity		
		Cisco Nexus 9504	Cisco Nexus 9508	Cisco Nexus 9516
	QSFP28 line card			

Table 3. Cisco Nexus 9500 R-Series Line Cards

Product ID	Description	Maximum Quantity	
		Cisco Nexus 9504	Cisco Nexus 9508
N9K-X9636C-R	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet QSFP28 line card	4	8
N9K-X9636C-RX	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet QSFP28 line card	4	8
N9K-X9636Q-R	Cisco Nexus 9500 36-port 40 Gigabit Ethernet QSFP line card	4	8
N9K-X96136YC-R	Cisco Nexus 9500 16-port 1/10 Gigabit, 32-port 10/25 Gigabit, and 4-port 40/100 Gigabit Ethernet line card	4	8
N9K-X9624D-R2	Cisco Nexus 9500 24-port 400 Gigabit QDD line card	Not supported	8

Table 4. Cisco Nexus 9500 Cloud Scale Fabric Modules

Product ID	Description	Minimum	Maximum
N9K-C9504-FM-E	Cisco Nexus 9504 100-Gigabit cloud scale fabric module	4	5
N9K-C9504-FM-G	Cisco Nexus 9500 4-slot 1.6Tbps cloud scale fabric module	4	5
N9K-C9508-FM-E	Cisco Nexus 9508 100-Gigabit cloud scale fabric module	4	5
N9K-C9508-FM-E2	Cisco Nexus 9508 100-Gigabit cloud scale fabric module	4	5
N9K-C9508-FM-G	Cisco Nexus 9500 8-slot 1.6Tbps cloud-scale fabric module	4	5
N9K-C9516-FM-E2	Cisco Nexus 9516 100-Gigabit cloud scale fabric module	4	5

Table 5. Cisco Nexus 9500 R-Series Fabric Modules

Product ID	Description	Minimum	Maximum
N9K-C9504-FM-R	Cisco Nexus 9504 100-Gigabit R-Series fabric module	4	6

© 2023 Cisco and/or its affiliates. All rights reserved.

Product ID	Description	Minimum	Maximum
N9K-C9508-FM-R	Cisco Nexus 9508 100-Gigabit R-Series fabric module	4	6
N9K-C9508-FM-R2	Cisco Nexus 9508 400-Gigabit R-Series fabric module	4	6

Table 6. Cisco Nexus 9500 Supervisor Modules

Supervisor	Description	Quantity
N9K-SUP-A	1.8-GHz supervisor module with 4 cores, 4 threads, and 16 GB of memory	2
N9K-SUP-A+	1.8-GHz supervisor module with 4 cores, 8 threads, and 16 GB of memory	2
N9K-SUP-B	2.2-GHz supervisor module with 6 cores, 12 threads, and 24 GB of memory	2
N9K-SUP-B+	1.9-GHz supervisor module with 6 cores, 12 threads, and 32 GB of memory	2

Note: N9K-SUP-A and N9K-SUP-A+ are not supported on Cisco Nexus 9504 and 9508 switches with -R line cards.

Table 7. Cisco Nexus 9500 System Controller

Product ID	Description	Quantity
N9K-SC-A	Cisco Nexus 9500 Platform System Controller Module	2

Table 8. Cisco Nexus 9500 Fans and Fan Trays

Product ID	Description	Quantity
N9K-C9504-FAN	Fan tray for 4-slot modular chassis	3
N9K-C9504-FAN2	Fan tray that supports the Cisco N9K-C9504-FM-G fabric module	3
N9K-C9508-FAN	Fan tray for 8-slot modular chassis	3
N9K-C9508-FAN2	Fan tray that supports the Cisco N9K-C9508-FM-G fabric module	3
N9K-C9516-FAN	Fan tray for 16-slot modular chassis	3

Table 9. Cisco Nexus 9500 Fabric Module Blanks with Power Connector

Product ID	Description	Minimum	Maximum
N9K-C9504-FAN-PWR	Nexus 9500 4-slot chassis 400G cloud scale fan tray power connector	1	2
N9K-C9508-FAN-PWR	Nexus 9500 4-slot chassis 400G cloud scale fan tray power connector	1	2

© 2023 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Table 10. Cisco Nexus 9500 Power Supplies

Product ID	Description	Quantity	Cisco Nexus Switches
N9K-PAC-3000W-B	3 KW AC power supply	Up to 4 Up to 8 Up to 10	Cisco Nexus 9504 Cisco Nexus 9508 Cisco Nexus 9516
N9K-PDC-3000W-B	3 KW DC power supply	Up to 4 Up to 8 Up to 10	Cisco Nexus 9504 Cisco Nexus 9508 Cisco Nexus 9516
N9K-PUV-3000W-B	3 KW Universal AC/DC power supply	Up to 4 Up to 8 Up to 10	Cisco Nexus 9504 Cisco Nexus 9508 Cisco Nexus 9516
N9K-PUV2-3000W-B	3.15-KW Dual Input Universal AC/DC Power Supply	Up to 4 Up to 8 Up to 10	Cisco Nexus 9504 Cisco Nexus 9508 Cisco Nexus 9516

Table 11. Cisco Nexus 9200 and 9300 Switches

Cisco Nexus Switch	Description
N9K-C9316D-GX	1-RU switch with 16x400/100/40-Gbps ports.
N9K-C9364C-GX	2-RU fixed-port switch with 64 100-Gigabit SFP28 ports.
N9K-C93600CD-GX	1-RU fixed-port switch with 28 10/40/100-Gigabit QSFP28 ports (ports 1-28), 8 10/40/100/400-Gigabit QSFP-DD ports (ports 29-36)
N9K-C9364C	2-RU Top-of-Rack switch with 64 40-/100-Gigabit QSFP28 ports and 2 1-/10-Gigabit SFP+ ports. • Ports 1 to 64 support 40/100-Gigabit speeds. • Ports 49 to 64 support MACsec encryption. Ports 65 and 66 support 1/10 Gigabit speeds.
N9K-C9332C	1-RU fixed switch with 32 40/100-Gigabit QSFP28 ports and 2 fixed 1/10-Gigabit SFP+ ports.
N9K-C9332D-GX2B	1-Rack-unit (1RU) spine switch with 32p 400/100-Gbps QSFP-DD ports and 2p 1/10 SFP+ ports.
N9k-9348D-GX2A	48p 40/100/400-Gigabit QSFP-DD ports and 2p 1/10G/10G SFP+ ports
N9k-9364D-GX2A	64p 400/100-Gigabit QSFP-DD ports and 2p 1/10 SFP+ ports
N9K-C93180YC-FX3	48 1/10/25 Gigabit Ethernet SFP28 ports (ports 1-48) 6 10/25/40/50/100-Gigabit QSFP28 ports (ports 49-54)
N9K-C93180YC-FX3S	48 1/10/25 Gigabit Ethernet SFP28 ports (ports 1-48)

© 2023 Cisco and/or its affiliates. All rights reserved.

Cisco Nexus Switch	Description
	6 10/25/40/50/100-Gigabit QSFP28 ports (ports 49-54)
N9K-C93360YC-FX2	2-RU switch with 96 10-/25-Gigabit SFP28 ports and 12 40/100-Gigabit QSFP28 ports
N9K-C93240YC-FX2	1.2-RU Top-of-Rack switch with 48 10-/25-Gigabit SFP28 fiber ports and 12 40-/100-Gigabit Ethernet QSFP28 ports.
N9K-C93216TC-FX2	2-RU switch with 96 100M/1G/10G RJ45 ports, 12 40/100-Gigabit QSFP28 ports, 2 management ports (one RJ-45 and one SFP port), 1 console, port, and 1 USB port.
N9K-C93180YC-FX	1-RU Top-of-Rack switch with 10-/25-/32-Gigabit Ethernet/FC ports and 6 40-/100-Gigabit QSFP28 ports. You can configure the 48 ports as 1/10/25-Gigabit Ethernet ports or as FCoE ports or as 8-/16-/32-Gigabit Fibre Channel ports.
N9K-C93180YC-FX-24	1-RU 24 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.
N9K-C93108TC-FX	1-RU Top-of-Rack switch with 48 100M/1/10GBASE-T (copper) ports and 6 40-/100-Gigabit QSFP28 ports
N9K-C93108TC-FX-24	1-RU 24 1/10GBASE-T (copper) front panel ports and 6 fixed 40/100-Gigabit Ethernet QSFP28 spine-facing ports.
N9K-C93108TC-FX3P	1-RU fixed-port switch with 48 100M/1/2.5/5/10GBASE-T ports and 6 40-/100-Gigabit QSFP28 ports
N9K-C9348GC-FXP *	Nexus 9300 with 48p 100M/1 G, 4p 10/25 G SFP+ and 2p 100 G QSFP
N9K-C92348GC-X	The Cisco Nexus 92348GC-X switch (N9K-C92348GC-X) is a 1RU switch that supports 696 Gbps of bandwidth and over 250 mbps. The 1GBASE-T downlink ports on the 92348GC-X can be configured to work as 100-Mbps, 1-Gbps ports. The 4 ports of SFP28 can be configured as 1/10/25-Gbps and the 2 ports of QSFP28 can be configured as 40- and 100-Gbps ports. The Cisco Nexus 92348GC-X is ideal for big data customers that require a Gigabit Ethernet ToR switch with local switching.
N9K-C93180YC-EX	1-RU Top-of-Rack switch with 48 10-/25-Gigabit SFP28 fiber ports and 6 40-/100-Gigabit QSFP28 ports
N9K-C93180YC-EX-24	1-RU 24 1/10/25-Gigabit front panel ports and 6-port 40/100 Gigabit QSFP28 spine-facing ports
N9K-C93108TC-EX	1-RU Top-of-Rack switch with 48 10GBASE-T (copper) ports and 6 40-/100-Gigabit QSFP28 ports
N9K-C93108TC-EX-24	1-RU 24 1/10GBASE-T (copper) front panel ports and 6 40/100-Gigabit QSFP28 spine facing ports.
N9K-C9336C-FX2	1-RU switch with 36 40-/100-Gb Ethernet QSFP28 ports
N9K-C9336C-FX2-E	1- RU switch with 36 40-/100-Gb QSFP28 ports

***Note:** For N9K-C9348GC-FXP 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".

Table 12. Cisco Nexus 9200 and 9300 Fans and Fan Trays

© 2023 Cisco and/or its affiliates. All rights reserved.

Product ID	Description	Quantity	Cisco Nexus Switches
NXA-FAN-160CFM-PE	Fan module with port-side exhaust airflow (blue coloring)	3	9364C [1] 93360YC-FX2
NXA-FAN-160CFM-PI	Fan module with port-side intake airflow (burgundy coloring)	3	9364C [1] 93360YC-FX2
NXA-FAN-160CFM2-PE	Fan module with port-side exhaust airflow (blue coloring)	4	9364C-GX
NXA-FAN-160CFM2-PI	Fan module with port-side intake airflow (burgundy coloring)	4	9364C-GX
NXA-FAN-30CFM-B	Fan module with port-side intake airflow (burgundy coloring)	3	93108TC-EX 93108TC-FX [1] 93180YC-EX 93180YC-FX [1] 9348GC-FXP [1]
NXA-FAN-30CFM-F	Fan module with port-side exhaust airflow (blue coloring)	3	93108TC-EX 93108TC-FX [1] 93180YC-EX 93180YC-FX [1] 9348GC-FXP
NXA-FAN-35CFM-PE	Fan module with port-side exhaust airflow (blue coloring)	4 6	92300YC [1] 9332C [1] 93180YC-FX3S [2] 93180YC-FX3 93108TC-FX3P 9336C-FX2-E 9316D-GX 93600CD-GX
NXA-FAN-35CFM-PI	Fan module with port-side intake airflow (burgundy coloring) Fan module with port-side exhaust airflow (blue	4 6 6	92300YC [1] 9332C [1] 93180YC-FX3S [2] 93180YC-FX3 93108TC-FX3P 9316D-GX 93600CD-GX 9336C-FX2-E

[1] For specific fan speeds see the Overview section of the Hardware Installation Guide.

[2] This switch runs with +1 redundancy mode so that if one fan fails, the switch can sustain operation. But if a second fan fails, this switch is not designed to sustain operation. Hence before waiting for the major threshold temperature to be hit, the switch will power down due to entering the fan policy trigger command.

© 2023 Cisco and/or its affiliates. All rights reserved.

Product ID	Description	Quantity	Cisco Nexus Switches
	coloring)		
NXA-FAN-65CFM-PE	Fan module with port-side exhaust airflow (blue coloring)	3	93240YC-FX2 [1] 9336C-FX2 [1]
NXA-FAN-65CFM-PI	Fan module with port-side exhaust airflow (burgundy coloring)	3	93240YC-FX2 9336C-FX2 [1]

Table 13. Cisco Nexus 9200 and 9300 Power Supplies

Product ID	Description	Quantity	Cisco Nexus Switches
NXA-PAC-500W-PE	500-W AC power supply with port-side exhaust airflow (blue coloring)	2	93108TC-EX 93180YC-EX 93180YC-FX
NXA-PAC-500W-PI	500-W AC power supply with port-side intake airflow (burgundy coloring)	2	93108TC-EX 93180YC-EX 93180YC-FX
NXA-PAC-650W-PE	650-W power supply with port-side exhaust (blue coloring)	2	92300YC 93180YC-FX3S 93108TC-EX 93180YC-EX 93180YC-FX3
NXA-PAC-650W-PI	650-W power supply with port-side intake (burgundy coloring)	2	92300YC 93180YC-FX3S 93108TC-EX 93180YC-EX 93180YC-FX3
NXA-PAC-750W-PE	750-W AC power supply with port-side exhaust airflow (blue coloring) 1	2	9336C-FX2 9336C-FX2-E 9332C 93240YC-FX2
NXA-PAC-750W-PI	750-W AC power supply with port-side intake airflow (burgundy coloring) 1	2	9336C-FX2 9336C-FX2-E 9332C 93240YC-FX2
NXA-PAC-1100W-PE2	1100-W AC power supply with port-side exhaust airflow (blue coloring)	2	93240YC-FX2 9332C 9316D-GX 9336C-FX2 9336C-FX2-E 93600CD-GX
NXA-PAC-1100W-PI2	1100-W AC power supply with port-side intake airflow (burgundy coloring)	2	93240YC-FX2 9332C 9316D-GX 9336C-FX2 9336C-FX2-E 93600CD-GX
NXA-PAC-1100W-PI	Cisco Nexus 9000 PoE 1100W AC PS, port-side intake	2	93108TC-FX3P

© 2023 Cisco and/or its affiliates. All rights reserved.

Product ID	Description	Quantity	Cisco Nexus Switches
NXA-PAC-1100W-PE	Cisco Nexus 9000 PoE 1100W AC PS, port-side exhaust	2	93108TC-FX3P
NXA-PAC-1900W-PI	Cisco Nexus 9000 PoE 1900W AC PS, port-side intake	2	93108TC-FX3P
NXA-PAC-1200W-PE	1200-W AC power supply with port-side exhaust airflow (blue coloring)	2	93360YC-FX2 9364C
NXA-PAC-1200W-PI	1200-W AC power supply with port-side intake airflow (burgundy coloring)	2	93360YC-FX2 9364C
N9K-PUV-1200W	1200-W Universal AC/DC power supply with bidirectional airflow (white coloring)	2	92300YC 93108TC-EX 93108TC-FX 93360YC-FX2 93180YC-FX3S 93180YC-EX 93180YC-FX 9364C
NXA-PDC-930W-PE	930-W DC power supply with port-side exhaust airflow (blue coloring)	2	93108TC-EX 93180YC-EX 93360YC-FX2 93180YC-FX3S 93180YC-FX 9364C
NXA-PDC-930W-PI	930-W DC power supply with port-side intake airflow (burgundy coloring)	2	93108TC-EX 93180YC-EX 93360YC-FX2 93180YC-FX3S 93180YC-FX 9364C
NXA-PDC-1100W-PE	1100-W DC power supply with port-side exhaust airflow (blue coloring)	2	93240YC-FX2 93600CD-GX 9316D-GX 9332C 9336C-FX2 9336C-FX2-E
NXA-PDC-1100W-PI	1100-W DC power supply with port-side intake airflow (burgundy coloring)	2	93240YC-FX2 93600CD-GX 9316D-GX 9332C 9336C-FX2 9336C-FX2-E
UCSC-PSU-930WDC	930-W DC power supply with port-side intake (green coloring)	2	93108TC-EX 93180YC-EX
UCS-PSU-6332-DC	930-W DC power supply with port-side exhaust (gray coloring)	2	93108TC-EX 93180YC-EX
NXA-PHV-1100W-PE	1100-W AC power supply with port-side exhaust airflow (blue coloring)	2	93240YC-FX2 9336C-FX2
NXA-PHV-1100W-PI	1100-W AC power supply with port-side intake airflow	2	93240YC-FX2

© 2023 Cisco and/or its affiliates. All rights reserved.

Product ID	Description	Quantity	Cisco Nexus Switches
	(burgundy coloring)		9336C-FX2
NXA-PAC-2KW-PE	2000-W AC power supply with port-side exhaust airflow (blue coloring)	2	9364C-GX
NXA-PAC-2KW-PI	2000-W AC power supply with port-side intake airflow (burgundy coloring)	2	9364C-GX
NXA-PDC-2KW-PE	2000-W DC power supply with port-side exhaust airflow (blue coloring	2	9364C-GX
NXA-PDC-2KW-PI	2000-W DC power supply with port-side intake airflow (burgundy coloring)	2	9364C-GX
N2200-PAC-400W	400-W AC power supply with port-side exhaust airflow (blue coloring)	2	92348GC-X
N2200-PAC-400W-B	400-W AC power supply with port-side intake airflow (burgundy coloring)	2	92348GC-X
N2200-PDC-350W-B	350-W DC power supply with port-side intake airflow	2	92348GC-X
N2200-PDC-400W	400-W DC power supply with port-side exhaust airflow (blue coloring)	2	92348GC-X

Compatibility Information

Fabric Module and Line Card compatibility details are listed below.

Table 14. Cisco Nexus 9500 Cloud Scale Line Cards

Product ID	N9K-C9504-FM-G	N9K-C9508-FM-G	N9K-C9504-FM-E	N9K-C9508-FM-E	N9K-C9508-FM-E2	N9K-C9516-FM-E2
N9K-X9716D-GX	4	4	No	No	No	No
N9K-X9736C-FX	5	5	5	5	5	5
N9K-X97160YC-EX	4	4	4	4	4	4
N9K-X9788TC-FX	4	4	4	4	4	4
N9K-X9732C-EX	4	4	4	4	4	4
N9K-X9736C-EX	4	4	4	4	4	4
N9K-X9732C-FX	4 5 (n+1 redundancy)	4 5 (n+1 redundancy)	4 5 (n+1 redundancy)	4 5 (n+1 redundancy)	4 5 (n+1 redundancy)	4 5 (n+1 redundancy)

Table 15. Cisco Nexus 9500 R-Series Line Cards

© 2023 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential

Product ID	N9K-C9504-FM-R	N9K-C9508-FM-R
N9K-X9636C-RX	6	6
N9K-X9636Q-R	4 6 (n+2 redundancy)	4 6 (n+2 redundancy)
N9K-X9636C-R	5 6 (n+1 redundancy)	5 6 (n+1 redundancy)
N9K-X96136YC-R	6	6

Table 16. Cisco Nexus 9500 R2-Series Line Cards

Product ID	N9K-C9508-FM-R2
N9K-X9624D-R2	6

Optics

To determine which transceivers and cables are supported by a switch, see the Transceiver Module (TMG) Compatibility Matrix. To see the transceiver specifications and installation information, see the Install and Upgrade Guides.

Cisco Nexus Dashboard Insights for Data Center

Cisco NX-OS Release 10.2(6)M supports the Nexus Dashboard Insights on Cisco Nexus 9300-FX, 9300-FX2, 9300-FX3, 9300-GX, 9300-GX2, 9400, and 9800 platform switches and 9500 platform switches with -EX/FX/GX Line Cards. For more information, see the Cisco Nexus Insights documentation.

Upgrade and Downgrade

To perform a software upgrade or downgrade, follow the instructions in the Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide, Release 10.2(x). For information about an In Service Software Upgrade (ISSU), see the Cisco NX-OS ISSU Support Matrix.

Related Content

Cisco Nexus 9000 Series documentation: Cisco Nexus 9000 Series Switches

Cisco NX-OS Software Release and Image-naming Convention: Cisco NX-OS Software Strategy and Lifecycle Guide

Cisco Nexus 9000 and 3000 Series NX-OS Switch License Navigator: Cisco Nexus 9000 and 3000 Series NX-OS Switch License Navigator

Cisco Nexus 9000 Series Software Upgrade and Downgrade Guide: Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide, Release 10.2(x)

© 2023 Cisco and/or its affiliates. All rights reserved.

Cisco Nexus 9000 Series FPGA/EPLD Upgrade Release Notes: Cisco Nexus 9000 Series FPGA/EPLD Upgrade Release Notes, Release 10.2(6).

Cisco Nexus 3000 and 9000 Series NX-API REST SDK User Guide and API Reference: Cisco Nexus NX-API Reference

Cisco NX-OS Supported MIBs: ftp://ftp.cisco.com/pub/mibs/supportlists/nexus9000/Nexus9000MIBSupportList.html

Supported FEX modules: Cisco Nexus 9000 Series Switch FEX Support Matrix

Licensing Information: Cisco NX-OS Licensing Guide and Cisco Nexus Smart Licensing Using Policy User Guide

When you downgrade from Cisco NX-OS Release 10.2(6)M to an earlier release, the features that use the ACI+NX-OS Essentials, Advantage, and add-on licenses or the Hardware Streaming Telemetry license continue to work in honor mode in the downgraded version. In addition, the output of the show license usage command continues to include entries for these unsupported licenses.

For more information, see the Cisco NX-OS Licensing Guide.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus9k-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: https://www.cisco.com/c/en/us/about/legal/trademarks.html. 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. (1721R)

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.

© 2023 Cisco Systems, Inc. All rights reserved.