cisco

Cisco Nexus 9000 Series NX-OS Release Notes, Release 9.3(7)

This document describes the features, issues, and exceptions of Cisco NX-OS Release 9.3(7) software for use on Cisco Nexus 9000 Series switches.

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.

Table 1. Changes to this Document

Date	Description
April 25, 2024	Added CSCwh50989 and CSCwe53655 to Open Issues.
July 18, 2023	Added flex link to the Enhanced Features section.
October 5, 2021	Added details about Thousand Eyes (TE) Integration feature in the New and Enhanced Software Features section.
March 11, 2021	Cisco NX-OS Release 9.3(7) became available.

New and Enhanced Software Features

New Features	
Feature	**Description**
RPC to get generate diff between candidate and running datastores	Introduced Cisco NETCONF RPC - get_diff that enables a NETCONF server to return the difference between configuration databases. This new option instructs the switch to calculate on-box configuration difference in candidate or running datastore. For more information, see the [Cisco Nexus 9000 Series NX-OS Programmability Guide, Release 9.3(x)](#).
Thousand Eyes (TE) Integration	Introduced Thousand eyes integration support with Cisco Nexus 9000 Series switches. For product overview look at: https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/at-a-glance-c45-2431016.html It is a must to install the following general SMU when TE integration is performed: nxos.CSCvz52812-n9k_ALL-1.0.0-9.3.7.lib32_n9000.tar For SMU installation please refer to the following guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/93x/system-management/b-cisco-nexus-9000-series-nx-os-system-management-configuration-guide-93x/b-cisco-nexus-9000-series-nx-os-system-management-configuration-guide-93x_chapter_010111.html

The enhanced features listed below are existing features introduced in earlier releases but enhanced with new support in Cisco NX-OS Release 9.3(7).

Enhanced Features	
Feature	**Description**
SNMPv3 encryption with SHA256	Introduced HMAC-SHA-256 authentication protocol support for SNMPv3. For more information, see the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide, Release 9.3(x)](#) .
MACsec with QSA	Added MACsec support when QSA is enabled in Cisco Nexus N9K-C9364C and N9K-C9336C-FX2 platform switches. Cisco Nexus N9k-C9364C and N9K-C9336C-FX2 supports MACSEC with QSA and 10G link

Enhanced Features	
Feature	**Description**
	combination from Cisco NX-OS Releasee 9.3(7). For more information, see the [Cisco Nexus 9000 Series Security Configuration Guide, Release 9.3(x)](#).
Flex Link	Added support for flex link on Cisco N9K-C93180YC-FX3 platform switch. For more information, see [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide, Release 9.3(x)](#).
FEX	Added support for FEX on Cisco Nexus 93180YC-FX3S platform switch. This switch can be used in either FEX mode or TOR mode. For more information on converting the switch from one mode to another. **Note:** **The Cisco Nexus 93180YC-FX3S switch in software FEX mode supports 25-G FEX connectivity to the host for single-point-of-management.** For more information, see the [Cisco Nexus 2000 Series NX-OS Fabric Extender Configuration Guide for Cisco Nexus 9000 Series Switches, Release 9.3(x).](#)

New Hardware Features

The following new hardware are introduced in Cisco NX-OS Release 9.3(7):

- NXA-SFAN-65CFM-PI Fan for N9K-C9336C-FX2
- NXA-SFAN-65CFM-PE Fan for N9K-C9336C-FX2

Note: To enable or disable displaying the serial number of the NXA-SFAN-65CFM-PI or NXA-SFAN-65CFM-PE fan, enter the **hardware fan-sprom** command.

Open Issues

Bug ID	Description
[CSCvx53013](#)	Headline: Unsupported TCAM region config added when ISSU from 9.2.4 to 9.3.4 Symptoms: Following lines exist in running-config and cannot be removed: hardware access-list tcam region ing-racl-lite 534756472 double-widehardware access-list tcam region ing-short 189402400 double-widehardware access-list tcam region ing-ifacl-ipv4-lite 1hardware access-list tcam region ing-ifacl-ipv6-lite 32784hardware access-list tcam region ing-cntacl 1hardware access-list tcam region ing-mvpn 370590546 double-widehardware access-list tcam region ing-l2-l3-qos 28hardware access-list tcam region hw-telemetry 1 Workaround: Perform write-erase followed by reload.

Bug ID	Description
CSCvx60047	Headline: "no negotiate auto" is removed from RJ-45 member interface and not from PO

Symptoms: If a RJ-45 port is member of a port-channel, with "no negotiate auto" configuration on the port-channel and member interfaces, that configuration gets removed from the member interface, but not from the port-channel after ISSU from a pre-9.3(4) or pre-7.0(3)I7(8) image to later images.

"show running interface e<x/y>" would not have "no negotiate auto" config

"show running interface po<>" would have "no negotiate auto" config

"no channel-group <>" config on the member interface will not take effect

Workaround:
Workaround for Pre-ISSU:

  - Remove "no negotiate auto" config from port-channel interface with RJ45 ports

 Workaround for Post-ISSU (if the pre-ISSU workaround is not applied):

  - Remove "no negotiate auto" config from port-channel interface with RJ45 ports

  - Run "channel-group " config on all members again |
| CSCvx56768 | Headline: 1G GLC-TE port was not coming up on N9K-C93360YC-FX2 after unplug and replug the SFP

Symptoms: Connect C2960X to Nexus 9000 with GLC-TE. Keep the connection up first and then unplug and replug the SFP on Nexus 9000 side. The port will not come up untill you shut /no shut the port on Nexus 9000 side.

Workaround: Shut /no shut the port on Nexus 9000 side. |
| CSCvx58626 | Headline: SNMP Crash in Nexus 9000 after ISSU Upgrade

Symptoms: Nexus9000 C9504, supervisor "N9K-SUP-A" crashed in during no disruptive upgrade from 7.0(3)I4(7) to 7.0(3)I7(9) aborting the ISSU.

Workaround: None |
| CSCvx60909 | Headline: Installing multiple SMUs do not remain committed after reload.

Symptoms: When installing multiple SMUs they do not remain committed after reload.

Workaround: Upgrade to a version that contains the fixes natively. |
| CSCvx61330 | Headline: Nexus 9000 aclqos cores - ERSPAN w/source VLAN mapped to VNI on certain ports

Symptoms: Gen1 Nexus 9000 models may see a core file from the aclqos process when trying to do an ERSPAN on a VLAN |

Bug ID	Description
	mapped to a VNI, using ports from speicifc ASICs. Workaround: None
CSCvx39125	Headline: Module ejector interrupt storm causing plfm mgr crash Symptoms: Marginal seating of IO module can lead to ejector button driven interrupt storm which causes platform manager to crash. Workaround: Reseat module to verify good connection, verify chassis grounding, etc.
CSCvx41778	Headline: BGP flaps (with holdtimer expired) every time a new leaf is added or reloaded Symptoms: VXLAN spine and leaf mode on infrastructure of Nexus 93180YC works fine without any problem. The issue arises when reloading leaf X or introducing a new leaf, this will make that a random leaf of the network connected to same spines loses BGP connection. This happens randomly, affecting different leafs.The evpn architecture implemented is based on a bgp ipv4 underlay, making it possible to announce the loopbacks of the Nexus serving as termination for the VTEPs. Workaround: None
CSCvx59644	Headline: HW Multipath table is not programming some links. Symptoms: Available paths seen on show ip route for equal cost paths are not being fully utilized. Workaround: Bundle uplinks in L3 port-channel
CSCvx60023	Headline: MAC Mobility Seq is not updated correctly upon MAC moves Symptoms: MAC Mobility sequence number is not correctly updated on EVPN routes when a host moves from one site to another.Missing the correct Sequence number might lead BGP to select non-optimum paths. Workaround: None
CSCvx60758	Headline: Bringing up SPAN session silently fails when sFlow data sources are configured Symptom: A Nexus 9000 series switch configured with sFlow data sources is not able to administratively bring SPAN sessions online. This is expected behavior and is a documented limitation. If at least one sFlow data source is configured, the SPAN sessions cannot be brought up. However, no error message or feedback is presented if one attempts to bring a SPAN session up while an sFlow data source is configured. Workaround: There is no known workaround for this issue. This defect introduces an error message to the CLI of the switch when one attempts to bring a SPAN session up while sFlow data sources are configured on the switch.
CSCvx60778	Headline: SPAN session administrative status becomes stuck when brought up alongside sFlow data sources

Bug ID	Description
	Symptom: A Nexus 9000 series switch configured with sFlow data sources is not able to administratively bring SPAN sessions online. This is expected behavior and is a documented limitation. If at least one sFlow data source is configured, the SPAN sessions cannot be brought up. However, if one removes all sFlow data sources from the switch's configuration and attempts to administratively bring the SPAN session up, the SPAN session will remain in a down state. Workaround: Administratively shutting down the SPAN session (even though it's technically already administratively shut down), then attempt to bring it back up.
CSCvx61244	Headline: `udld disable` on breakout interfaces is removed from running configuration after upgrade Symptom: UDLD operational on a 4x10G interface that was previously configured to have udld disabled. Workaround:None
CSCvx61532	Headline: CFS HAP reset and core file during system switchover Symptom: A Nexus 9500 switch on 7.0(3)I7(x) version may undergo a HAP reset and generate a core file on performing a system switchover. Workaround: None
CSCvw67472	Headline: Nexus9300 unexpect options (port 1 216) are set on show run all Symptom: When configuring "hardware rate-limiter access-list-log XXX" the port 1 216 are set on show run all like below. ---hardware rate-limiter layer-3 glean 100 module 1 port 1 216hardware rate-limiter layer-3 multicast local-groups 3000 module 1 port 1 216hardware rate-limiter access-list-log 100 module 1 port 1 216hardware rate-limiter bfd 10000 module 1 port 1 216hardware rate-limiter fex 12000 module 1 port 1 216hardware rate-limiter span 50 module 1 port 1 216hardware rate-limiter sflow 40000 module 1 port 1 216hardware rate-limiter vxlan-oam 1000 module 1 port 1 216hardware rate-limiter 100M-Ethports 10000 module 1 port 1 216hardware rate-limiter DOT1X 3000 module 1 port 1 216hardware rate-limiter mpls-oam 300 module 1 port 1 216hardware rate-limiter netflow 120000 module 1 port 1 216---But it is not possible to even configure "... port 1 216" on the device.---switch# conf tEnter configuration commands, one per line. End with CNTL/Z.switch(config)# hardware rate-limiter layer-3 glean 100 module 1 port 1 216 ^% Invalid number, range is (1:48) at '^' marker.switch(config)# --- Workaround: Write erase and then reload the device.
CSCvx57409	Headline: BGP remote-as route-map will be delete after default afi Symptom: The BGP neighbor remote-as route-map will be deleted after the default afi show running config neighbor 10.1.1.0/24 remote-as route-map CLI is run. Workaround: This is not an issue. The route-map is actually in effect.
CSCvx60094	Headline: Random link flaps on N9K-C93108TC-EX with QSFP-4x10G-AOC Symptom: Random L1 link flaps on N9K-C93108TC-EX with QSFP-4x10G-AOC

Bug ID	Description
	Workaround:Disable/blacklist/admin-down all not-connected/unused breakout ports to reduce the frequency of link flap.
CSCvx61314	Headline: N9k VXLAN - Non-Zero ESI value can result in blackholing.
	Symptom: Example output:
	N9K-VTEP# show mac address-table vlan 2431 | i beefC 2431 dead.beef.aaaa dynamic 0 F F nve1(192.168.50.10 192.168.50.11) <- - - MAC is associated with two NVE peers.
	N9K-VTEP# show l2route evpn mac evi 2431 mac dead.beef.aaaa detailFlags -(Rmac):Router MAC (Stt):Static (L):Local (R):Remote (V):vPC link(Dup):Duplicate (Spl):Split (Rcv):Recv (AD):Auto-Delete (D):Del Pending(S):Stale (C):Clear, (Ps):Peer Sync (O):Re-Originated (Nho):NH-Override(Pf):Permanently-Frozen, (Orp): OrphanTopology Mac Address Prod Flags Seq No Next-Hops---------- -------------- ------ ------------ --------- ----------------2431 dead.beef.aaaa BGP Rcv 0 192.168.50.10 Route Resolution Type: ESI Forwarding State: Resolved (PL) Resultant PL: 192.168.50.10, 192.168.50.11 <- - - Two Resultant PL entries. Sent To: L2FM ESI : 03aa.aaaa.aaaa.aa00.0001 <- - - ESI is present.
	N9K-VTEP# show hardware mac address-table 1 vlan 2431 address dead.beef.aaaaFE | VLAN | MAC | Dynamic | Port |Location Index| | | | | | |---+------+--------------+---------+---------+--------------+0 2431 dead.beef.aaaa dynamic <- - - Egress port does not exist.Remove one of the remote VTEP and it works:
	N9K-VTEP# show l2route evpn mac evi 2431 mac dead.beef.aaaa detailFlags -(Rmac):Router MAC (Stt):Static (L):Local (R):Remote (V):vPC link(Dup):Duplicate (Spl):Split (Rcv):Recv (AD):Auto-Delete (D):Del Pending(S):Stale (C):Clear, (Ps):Peer Sync (O):Re-Originated (Nho):NH-Override(Pf):Permanently-Frozen, (Orp): OrphanTopology Mac Address Prod Flags Seq No Next-Hops---------- -------------- ------ ------------ --------- ----------------2431 dead.beef.aaaa BGP Rcv 0 192.168.50.10 Route Resolution Type: ESI Forwarding State: Resolved (PL) Resultant PL: 192.168.50.10 <- - - One Resultant PL entries. Sent To: L2FM ESI : 03aa.aaaa.aaaa.aa00.0001 <- - - ESI is present.
	N9K-VTEP# show hardware mac address-table 1 vlan 2431 address dead.beef.aaaaFE | VLAN | MAC | Dynamic | Port |Location Index| | | | | | |---+------+--------------+---------+---------+--------------+0 2431 dead.beef.aaaa dynamic nve1 <- - - Egress port is now populated.
	Workaround: Do not use ESI. ESI is not currently supported on Cloudscale platforms. For Nexus 9000 switches vPC can be used instead of ESI for dual homing and redundancy.
CSCvx50717	Headline: Nexus 9500 EX/FX SVI unicast counters do not work
	Symptom: Unicast SVI counters remain 0 in `show interface` and `show interface` and the `show vlan id X counters`.
	Workaround: None
CSCvy19448	Headline: SSH Connection Rejected with FIPS enabled using any ssh key
	Symptom: SSH connections will be rejected if FIPS feature is enabled on version 9.3.7. This issue is similar to CSCvu10721 but ssh is also rejected for ecdsa key.
	Workaround: Downgrade to version 9.3.6 or earlier, or upgrade to 10.1.1. Issue will be resolved in next NXOS release, 9.3.8.
	There is a general available SMU to address this issue on 9.3(7):
	https://software.cisco.com/download/home/286314783/type/286278856/release/9.3(7)

Bug ID	Description
	SMU install instructions: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/93x/system-management/b-cisco-nexus-9000-series-nx-os-system-management-configuration-guide-93x/b-cisco-nexus-9000-series-nx-os-system-management-configuration-guide-93x_chapter_010111.html
CSCvy24198	Headline: l2fm process crash after l2fm_mcec_get_mac_handler Symptom: l2fm process crashed after vPC came online: %$ VDC-1 %$ %ASCII-CFG-2-CONF_CONTROL: System ready %$ VDC-1 %$ %VPC-2-PEER_KEEP_ALIVE_RECV_FAIL: In domain 100, VPC peer keep-alive receive has failed %$ VDC-1 %$ %VPC-2-PEER_KEEP_ALIVE_RECV_FAIL: In domain 100, VPC peer keep-alive receive has failed (message repeated 1 time) %$ VDC-1 %$ %VPC-2-PEER_KEEP_ALIVE_RECV_FAIL: In domain 100, VPC peer keep-alive receive has failed (message repeated 1 time) ... %$ VDC-1 %$ %SYSMGR-2-SERVICE_CRASHED: Service "l2fm" (PID 7824) hasn't caught signal 11 (core will be saved). Workaround: Disconnect vPC peer link and upgrade both peers separately.
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.
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.

Resolved Issues

Bug ID	Description
CSCvx21260	Headline: Nexus 9000/3000 NXOS : Micron_M500IT Bootflash in read only mode Symptoms: Nexus 9000/3000 switch bootflash goes into read-only mode with Micron_M500IT SSD drive after 28,224 power-on-hours (POH) for the first time. The bootflash will stop responding causing failure of operations such as config changes/save, read/write operations etc.

Bug ID	Description
	Workarounds: Reload the switch. However, this failure will reappear after 1008 hours of operation.
CSCvt19116	Headline: Port-channel load-balancing for MPLS tagged traffic on L2 transit N9k FX/EX switches
	Symptoms: Device does not perform 5-tuple hashing for MPLS packets passing through. Even if switch does not perform label switching and treats packets as simple L2 packets it still does not check inner headers. As result switch does not perform optimal load-balancing over port-channel. Here only src-dst mac of outer L2 header are used only.
	Workarounds:
	• Enable feature mpls segment-routing (without further label switching use).
	• MPLS SR and VPC features are mutually exclusive. So device must be no VPC enabled.
CSCvx25283	Headline: msdp owned (s,g) mroute does not inherit pim oif from (*,g)
	Symptoms: Receivers are not getting multicast stream.
	MSDP owned (S,G) mroute sync'd by MSDP peer/RP/Catalyst to RP/Nexus, does not inherit OIF from (*,G). Issue is specific to few groups, some groups are working fine too.
	Workarounds: Add static OIF (S,G) for outgoing interface.
	interface Ethernet x/y
	ip igmp static-oif multicast_group source multicast_sender_ip
CSCvx24194	Headline: 100G FR: Transceiver not coming up between C93600CD-GX and Mellanox switch
	Symptoms: Mellanox side is using a breakout cable breaking out from 400G into 4x100G.
	From Mellanox side one leg is connected to another vendor switch (port is up), the other 100G connected to N9k is down.
	Workarounds: NA.
CSCvk45018	Headline: BFD is blocked over Unnumbered Interfaces
	Symptoms: BFD would not come up over L3 interfaces if configured as unnumbered.
	Workarounds: None
CSCvt19116	Headline: Non-MPLS device does not perform 5-tuple hashing for MPLS packets passing through. (FX,EX)
	Symptoms: Device does not perform 5-tuple hashing for MPLS packets passing through. Even if switch does not perform label switching and treats packets as simple L2 packets it still does not check inner headers. As result switch does not perform optimal load-balancing over port-channel. Here only src-dst mac of outer L2 header are used only.
	Workarounds: Enable feature mpls segment-routing (without further label switching use). MPLS SR and VPC features are mutually exclusive. So device must be no VPC enabled.
CSCvv00706	Headline: Need to add a cli to view SFP details on N9K Mgmt0 port
	Symptoms: No user friendly cli available to view if SFP is inserted or not on N9K mgmt port
	Workarounds: Need to go into bash and get the info, reachout to TAC
CSCvv21797	Headline: no vpc domain 20 removes previously configure vpc domain 10 without any warning

Bug ID	Description
	Symptoms: On a Nexus 9000 switch one can remove an existing vpc domain with the wrong vpc id number.I.e examplevpc domain 10 is configured but the executed command is:no vpc domain 20this will remove domain 10 without any further warning, despite that we tried to remove domain 20 which does not exist in this example. Workarounds: None
CSCvv75490	Headline: NFM process crashes leads to hap-reset Symptoms: A Nexus device is seeing multiple nfm (netflow) process crashes. This results in the system getting reloaded due to a hap-reset:2020 Sep 16 19:51:15 N9K %$ VDC-1 %$ %SYSMGR-2-SERVICE_CRASHED: Service "nfm" (PID 31961) hasn't caught signal 6 (core will be saved).2020 Sep 16 19:52:16 N9K %$ VDC-1 %$ %SYSMGR-2-SERVICE_CRASHED: Service "nfm" (PID 3667) hasn't caught signal 6 (core will be saved).2020 Sep 16 19:53:17 N9K %$ VDC-1 %$ %SYSMGR-2-SERVICE_CRASHED: Service "nfm" (PID 4754) hasn't caught signal 6 (core will be saved).2020 Sep 16 19:54:18 N9K %$ VDC-1 %$ %SYSMGR-2-SERVICE_CRASHED: Service "nfm" (PID 5703) hasn't caught signal 6 (core will be saved).2020 Sep 16 19:54:18 N9K %$ VDC-1 %$ %SYSMGR-2-HAP_FAILURE_SUP_RESET: Service "nfm" in vdc 1 has had a hap failure Workarounds: Remove netflow if possible. Downgrading to a version bellow 9.2(4) or 9.3(1) should also stop the crash
CSCvv98385	Headline: Nexus 9000 "Fabric Port Class X Output Drop Packets" counter constantly incrementing Symptoms: A Nexus 9000 Series switch equpped with a Cloud Scale ASIC may have one or more large non-zero "Fabric Port Class X Output Drop Packets" counters in the output of `show hardware internal errors all`. An example of this output is shown below.switch# show hardware internal errors all <snip>ID Name Value Ports-- ---- ----- -----6029314 Fabric Port Class 6 Output Drop Packets 0000000000000016 1,9,17 - 6029316 Fabric Port Class 6 Output Drop Packets 0000000000000004 1,17 - 6029327 Fabric Port Class 6 Output Drop Packets 0000000000229751 1,17,61 - 6029328 Fabric Port Class 6 Output Drop Packets 0000000000307963 1,17 - 6160386 Fabric Port Class 7 Output Drop Packets 0000000000000016 1,9,17 - 6160388 Fabric Port Class 7 Output Drop Packets 0000000000000004 1,17 - 6160399 Fabric Port Class 7 Output Drop Packets 0000000000228748 1,17,61 - 6160400 Fabric Port Class 7 Output Drop Packets 0000000000307834 1,17 - 6422530 Fabric Port Class 1 Output Drop Packets 0000000000000001 1,9,17 - 6422532 Fabric Port Class 1 Output Drop Packets 0000000000000004 1,17 - 6422543 Fabric Port Class 1 Output Drop Packets 0000000000000081 1,17,61 - 6422544 Fabric Port Class 1 Output Drop Packets 0000000000000269 1,17 - 6684672 Fabric Port Class 3 Output Drop Packets 0000000048748371 1,17 - Workarounds: There is no known workaround to this issue.
CSCvw24004	Headline: SVI goes down as soon as access interfaces goes down and only flex link stays up Symptoms: SVI interface can go down when corresponding VLAN is active (forwarded by) on FlexLink only. Workarounds: In order to restore SVI state, you will need to shutdown/no shutdown it.
CSCvw30171	Headline: N9k EX - Dhcp snooping binding not refresh Symptoms: If have approximate 2000 dhcp snooping entries in N9K after reloading the dhcp snooping binding will not refresh. Workarounds: None
CSCvw34566	Headline: NXOS rfc1583compatibility not consistent with IOS/XE implementation Symptoms: NXOS by default supports rfc2328. To support inter-op with the devices which supports rfc1583 NXOS provides configuration option "rfc1583compatibility" . This makes NXOS to follow rfc1583. When NXOS devices operating in rfc1583compatibility mode it may choose a different

Bug ID	Description
	path to ASBR compare to IOS/XE.There is no concern if "rfc1583compatibility" is not configured. Workarounds: Use RFC2328 throughout the network instead
CSCvw43139	Headline: after changing VPC HIF ports to orphan, Flood traffic crossing MCT does not egress HIF Symptoms: + HIF ports are configured as VPCs members.+ after changing the config in a certain way to be orphan ports.+ broadcast traffic are not egressing hif ports and causing disconnection.+ any broadcast traffic ingressing peer-link will not forwarded to orphan ports which were members of vpcs before Workarounds: Reconfigure the ports.
CSCvw43442	Headline: VXLAN: MAC address learned on Fabric Layer3 interfaces Symptoms: MAC address learnt on a L3 Fabric interface of VXLAN EVPN Border Gateway (BGW) switch.As result, MAC address may move between the L3 interface and NVE interface and eventually gets `Permanently-Frozen` on one of the leaf. Workarounds: None. Excessive MAC move will cause software to permanently keep the MAC in frozen state. On L3 interface MAC learning shouldn't happen and MAC frozen state doesn't impact L3 forwarding
CSCvw48498	Headline: N9K - NXAPI : %SAFE_STR CONSTRAINT: strncpy_s: slen is zero, error code=401 Symptoms: Viewing following error message : ?%SAFE_STR CONSTRAINT: strncpy_s: slen is zero, error code=401? into /var/sysmgr_nxapi/logs/error.log file Workarounds: From bash mode, clear the content of the error.log manually by echo-ing empty string inside the file.(To get into bash mode - have the bash feature enabled and run bash to get inside)bash-4.3$ echo "" > /var/sysmgr_nxapi/logs/error.log
CSCvw49456	Headline: Nexus switch "ipfib" crash Symptoms: A Nexus switch might experience an IPFIB crash Reset Reason for this card:Image Version : 9.3(3)Reset Reason (LCM): Unknown (0) at time Mon Nov 16 12:07:20 2020Reset Reason (SW): Reset Requested due to Fatal Module Error (4) at time Mon Nov 16 12:04:27 2020 Service (Additional Info): System managerReset Reason (HW): Module PowerCycled (112) at time Mon Nov 16 12:07:20 2020Last log in OBFL was written at time Mon Nov 16 11:23:17 2020%SYSMGR-SLOT1-2-SERVICE_CRASHED: Service "ipfib" (PID 1797) hasn't caught signal 11 (core will be saved).%SYSMGR-SLOT1-2-HAP_FAILURE_SUP_RESET: Service "ipfib" in vdc 1 has had a hap failure%SYSMGR-SLOT1-2-LAST_CORE_BASIC_TRACE: fsm_action_become_offline: PID 19347 with message Could not turn off console logging on vdc 1 error: mts req-response with syslogd in vdc 1 failed (0xFFFFFFFF) . %SYSMGR-SLOT1-2-LAST_CORE_BASIC_TRACE: save_core: PID 2428 with message cd /var/sysmgr/tmp_logs ; /isan/bin/sysmgr_logmgr 0x102_ipfib 1797 1 1>> /mnt/pss/core_handling.log 2>> /mnt/pss/core_handling.log failed with ret val 768 .%MODULE-2-MOD_DIAG_FAIL: Module 1 (Serial number: FDO241109A5) reported failure due to Service on linecard had a hap-reset in device DEV_SYSMGR (device error 0x17a)VDC Module Instance Process-name PID --- --- --- -------- --------------- --------1 1 1 ipfib 1801 Workarounds: No known workaround
CSCvw50234	Headline: NX93180LC - GLC-TEs not working after upgrade to 9.3.5 (notconnected) on even ports only Symptoms: GLC-TE transceiver stays in not connected state after upgrade to 9.3.5 on even ports (i.e. eth1/20). Odd ports are working fine (i.e. eth1/19)513E-C.06-N9K-C93180LC-EX# sh int eth1/20Ethernet1/20 is down (Link not connected)513E-C.06-N9K-C93180LC-EX# sh int eth1/20 transceiver Ethernet1/20 transceiver is present type is 1000base-T name is CISCO part number is SBCU-5740ARZ-CS1 revision is G3.1 serial number is AVC204822X4 nominal bitrate is 1300 MBit/sec Link length supported for copper is 100 m cisco id is 3 cisco extended id number is 4 cisco part number is 30-1475-01 cisco product id is GLC-TE cisco

Bug ID	Description
	version id is V01 Workarounds: Downgrade to nxos.7.0.3.I7.X
CSCvw51632	Headline: Observing PTP port flap which is connected to STU(GM) Symptoms: In scale environment, customer is facing port flap issue and due to this their cell site operation got impacted. Workarounds: Reduce the no.of master ports below 25 and do fine tuning of policer operation, so that policer wont discard ptp pkts.
CSCvw53323	Headline: CAP_FEATURE_MULTIHOP_BFD capability is not getting removed from the system after disabling BFD Symptoms: Switch gives an error during downgrade if BFD feature is directly disabled, instead of removing the multihop BFD configuration. CAP_FEATURE_MULTIHOP_BFD capability is not getting removed from the system and hence causing error.Ideally all the BFD related capabilities should have been deleted when feature is disabled. Workarounds: Following workaround can be used by customer to overcome this1.Enable BFD "feature bfd" 2.Add Multihop BFD configuration "bfd multihop interval 500 min_rx 500 multiplier 5" 3.*Delete Multihop BFD configuration "no bfd multihop interval 500 min_rx 500 multiplier 5" 4.Disable BFD "no feature bfd"
CSCvw53976	Headline: Incorrect Forwarding on Hardware. VXLAN Flood & Learn. IPFIB process stuck and MTS queue build up. Symptoms: We see traffic forwarded incorrectly to other interfaces.Traffic that should be punted to CPU is getting VXLAN encapsulatedshow forwarding commands do not workshwo tech l3 forwarding generates no output.show ip mroute for underlay mcast shows pending routesWe see messages stuck on queue for ipfib.SPINE1# show system internal mts buffers summary* recv_q: not received yet (slow receiver)* pers_q/npers_q/log_q: received not dropped (leak)node sapno recv_q pers_q npers_q log_q app/sap_descriptionlc 207 142 0 0 0 ipfib/ipfib SAP for lp msgslc 205 118 0 0 0 ipfib/ipfib SAP for mcast statsmsgslc 201 30 0 0 0 ipfib/ipfib SAP for mcast msgslc 200 13 1 0 0 ipfib/ipfib SAP for ucast msgslc 199 2 0 0 0 ipfib/ipfib SAP for sp msgsup 214 0 40 0 0 mfdm/MFDM SAP for L3 update msgssup 212 0 0 45 0 mfdm/MFDM SAP for sys msgssup 196 0 12 0 0 ufdm/ufdm SAP for hp msgssup 284 0 6 0 0 netstack/TCPUDP process client MTS queuesup 252 0 0 4 0 am/AM LOW PRI Q MTS SAP Workarounds: Passive-interface for SVIs shared between spines to avoid creating an IGP adjacency betweenSPINES over VXLAN and reload if mts suck in queue.
CSCvw56696	Headline: NXOS: VLAN Manager crashes with VLAN mapping configuration change Symptoms: The VLAN Manager (vlan_mgr) service reports a crash and a core file is saved: %SYSMGR-2-SERVICE_CRASHED: Service "vlan_mgr" (PID ###) hasn't caught signal ## (core will be saved). To see the core files run 'show cores': `show cores`VDC Module Instance Process-name PID Date(Year-Month-Day Time)--- ------ -------- --------------- --------- ------------------------# # # vlan_mgr # YYYY-MM-dd HH:mm:ss The core triggers after a configuration change in VLAN mapping. Workarounds: None
CSCvw57406	Headline: N9k IPv6 Neighbor Discovery COPP looks at mac & ipv6 addr that it does not own Symptoms: In a vxlan fabric, currently all IPV6 ND packets are sent to the supervisior regardless wheather or not the switch owns the ipv6 address and or the destination mac address. In a routing loop, one packet is sent in the loop until the TTL expires. Each time the packet hit the switch it is punted to the CPU and hits the COPP class. This over runs the ND class in COPP and devices can not communicate with their default gateway.

Bug ID	Description
	Workarounds: None
CSCvw59799	Headline: Hairpin of L2 multicast NLB KA frames on the same interface it was received
	Symptoms: NLB KA packets, received on a VPC Po are hair-pining and exiting the same interface they were received on.
	Workarounds: Disable IGMP snooping
CSCvw60409	Headline: HSRP vmac is not cleared and remains as static entry after shutting down SVI.
	Symptoms: On N9k switch running HSRP, when SVI is shut down on active HSRP switch, HSRP vmac is not cleared andremains as static entry. This may cause traffic disruption.
	Workarounds: Remove HSRP configurations from SVI then shut down SVI.switch(config-if)# no hsrp 1 ipv4switch(config-if)# no hsrp 1 ipv6
CSCvw60736	Headline: N9K-C9348GC link up delay after reloaded
	Symptoms: After vPC peer reloaded, vPC member port linkup delay than the other end,which lead to packet loss when using channel group mode on.
	Workarounds: None
CSCvw65224	Headline: N9K-PAC-650W for N9k reporting Fail/Shut status
	Symptoms: The below Syslog will be seen on the switch: %KERN-3-SYSTEM_MSG: [6960403.222450] cctrlib_tor2_get_psu_env_info.557: PSU 1 failed to read CCTRL_PSU_READ_VIN (3 =>2.58.88.2) - kernelSyslog can be seen also for PSU 0 slot.The output of show environment will be as follows with either one or both of the PSUs showing the fail/shut state:# sh env powerPower Supply:Voltage: 12 VoltsPower Actual Actual TotalSupply Model Output Input Capacity Status (Watts) (Watts) (Watts)------- ---------- --------------- ------ ---------- ---------------- ----1 N9K-PAC-650W 0 W 0 W 0 W Fail/Shut2 N9K-PAC-650W 0 W 0 W 0 W Fail/Shut
	Workarounds: None
CSCvw66557	Headline: EOR - default v6 route not correctly programmed in hw with template-service-provider
	Symptoms: Traffic that hitting default v6 route is forwarded in software. This can cause packet drop for traffic using v6 default route due to sw switching and CoPP
	Workarounds: If using IPv6 static route:* Add and remove static route after device reload.If using IPv6 dynamic route:* Add a more specific IPv6 route for the affected prefix(es)
CSCvw68897	Headline: segmentation fault on call home service turn on after enabling smart license
	Symptoms: + Below logs2020 Dec 2 19:42:14 R3-USFWT-05-LF1 %LICMGR-5-LOG_SMART_LIC_EVAL_START: (pid=2294) Entering evaluation period2020 Dec 2 19:42:16 R3-USFWT-05-LF1 last message repeated 1 time2020 Dec 2 19:42:16 R3-USFWT-05-LF1 %SYSMGR-2-SERVICE_CRASHED: Service " licmgr" (PID 2294) hasn't caught signal 11 (core will be saved).2020 Dec 2 19:42:16 R3-USFWT-05-LF1 %CALLHOME-2-EVENT: SW_CRASH2020 Dec 2 19:42:16 R3-USFWT-05-LF1 %LICMGR-5-LOG_SMART_LIC_EVAL_START: (pid=26014) Entering evaluation period2020 Dec 2 19:42:17 R3-USFWT-05-LF1 last message repeated 1 time2020 Dec 2 19:42:17 R3-USFWT-05-LF1 %LICMGR-5-LOG_SMART_LIC_COMM_RESTORED: (pid=26014) Communications with the Cisco Smart Software Manager or satellite restored2020 Dec 2 19:43:36 R3-USFWT-05-LF1 %SYSMGR-2-SERVICE_CRASHED: Service " licmgr" (PID 26014) hasn't caught signal 6 (core will be saved).2020 Dec 2 19:43:37 R3-USFWT-05-LF1 %LICMGR-5-LOG_SMART_LIC_EVAL_START: (pid=6479) Entering evaluation period2020 Dec 2 19:43:38 R3-USFWT-05-LF1 last message repeated 1 time2020 Dec 2 19:43:38 R3-USFWT-05-LF1 %LICMGR-5-LOG_SMART_LIC_COMM_RESTORED: (pid=6479) Communications with the Cisco

Bug ID	Description
	Smart Software Manager or satellite restored2020 Dec 2 19:43:44 R3-USFWT-05-LF1 %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty by <user> on <ip>@pts/42020 Dec 2 19:44:47 R3-USFWT-05-LF1 %UFDM-3-FIB_IPv4_ADJ_CONSISTENCY_CHECKER_PASS: FIB IPv4 adjacency consistency checker PASSED on slot 12020 Dec 2 19:44:47 R3-USFWT-05-LF1 %UFDM-3-FIB_IPv4_ROUTE_CONSISTENCY_CHECKER_PASS: FIB IPv4 route consistency checker PASSED on slot 12020 Dec 2 19:45:16 R3-USFWT-05-LF1 %CALLHOME-2-EVENT: SW_CRASH licmgr in slot 1 crashed with crash type : stateful crash2020 Dec 2 20:24:48 R3-USFWT-05-LF1 %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty by A1264370-3 on <ip>@pts/42020 Dec 2 20:28:32 R3-USFWT-05-LF1 %LICMGR-3-LOG_SMART_LIC_AGENT_REG_FAILED: (pid=6479) Smart Agent for Licensing Registration with the Cisco Smart Software Manager or satellite failed: Response error: The product '<id>' and sudi { udi_pid: nil+licmgr cores generatedshow coreVDC Module Instance Process-name PID Date(Year-Month-Day Time)--- ------ -------- --------------- -------- ------------------------1 1 1 licmgr 2294 2020-12-02 19:42:431 1 1 licmgr 26014 2020-12-02 19:43:54 Workarounds:
CSCvw69648	Headline: 'vpc orphan-port suspend' configuration can't be removed on interfaces associated to vPC PO Symptoms: 'vpc orphan-port suspend' configuration can't be removed on interfaces associated to vPC port-channel. Workarounds: 1) the physical interface needs to be removed from the vPC PO2) the "vpc orphan-port suspend" command can now be removed from the physical interface3) Re-add the physical interface to the vPC PO
CSCvw70948	Headline: vPC BGW : multisite bgp-if is up for 30s after peer-link failure on vpc secondary Symptoms: After peer-link failure on VPC secondary:* the NVE source interface is immediately brought down* NVE is kept UP for additional 30s * the "multisite bgp-if" (dedicated loopback for multisite) is kept up for 30s (not tunable "Source Interface hold-up-time") This causes the traffic from DCI side to still be attracted possibly causing some blackholing Workarounds: None at this point, connectivity will be re-established after 30s
CSCvw73129	Headline: Unknown Unicast brief report doesn't indicate L2 miss Symptoms: ELAM outgoing interface may print an index giving an impression of unicast fowarding while traffic is L2 flood Workarounds: Refer to ELAM full report for FLOOD|MISS bit set:report detail	egrep "FLOOD|MISS"
CSCvw73676	Headline: Nexus 9K - Unable To Boot NXOS Version 9.x From USB Symptoms: + On Nexus 9K, when attempting to boot NXOS version 9.x from USB via the loader prompt you may see the output below, followed by getting kicked back to the loader prompt.===Bootable Disk is detected. Device Name: Micron_1100_MTFDDAV256TBNVersion 2.18.1260. Copyright (C) 2020 American Megatrends, Inc. FPGA SPI Flash Micron(Numonyx) N25Q128Board type 4IOFPGA @ 0xd8000000SLOT_ID @ 0xfSet fan speed to 60% Filesystem type is ext2fs, partition type 0x83ACI chassisTrying to read config file /boot/grub/menu.lst.local from (hd0,0) Filesystem type is fat, partition type 0xcTrying to read config file /boot/grub/menu.lst.local from (hd1,4) Filesystem type is ext2fs, partition type 0x83Trying to read config file /boot/grub/menu.lst.local from (hd1,5) Filesystem type is ext2fs, partition type 0x83Auto boot configuration file is absent.Autoboot image boot failed. Trying recovery imageTrying to read config file /boot/grub/menu.lst.recovery from (hd1,4) Filesystem type is ext2fs, partition type 0x83Trying to read config file /boot/grub/menu.lst.recovery from (hd1,5) Filesystem type is ext2fs, partition type 0x83Auto boot configuration file is absent.No autoboot or failed autoboot. falling to loader===

Bug ID	Description
	Workarounds: 1. Have 7.x and 9.x NXOS versions on USB2. In loader, boot 7.x version from USB3. Once successfully booted in 7.x, copy the 9.x image file from USB to bootflash via the "copy" command4. Modify the boot statements from the 7.x image to the 9.x image using the "boot" command5. Reload the switch to boot to the 9.x image
CSCvw76165	Headline: N9500:other end port of mgmt port with shutdown force is up

Symptoms: Other end port of mgmt port on N9500 SUP is up even though shutdown force is configured on mgmt port.

Workarounds: remove "shutdwon force" and configure "shutdown force" again on mgmt 0 |
| CSCvw76327 | Headline: Nexus N9K-X9536PQ T2-EOR-LC card crashing and causing bcm-crash core

Symptoms: Nexus N9K-X9536PQ line card crashing and causing bcm-crash core.The following may be reported in the log:2020 Nov 21 09:44:20 %SYSMGR-SLOT3-2-SERVICE_CRASHED: Service "bcm_usd" (PID 8338) hasn't caught signal 6 (core will be saved).2020 Nov 21 09:44:23 %SYSMGR-SLOT3-3-BASIC_TRACE: generate_kernel_traces_file: PID 6553 with message failed to write kernel trace to /var/sysmgr/tmp_logs/0x302_bcm_usd_kernel-trace.8338. return value -1 . 2020 Nov 21 09:44:23 %SYSMGR-SLOT3-2-HAP_FAILURE_SUP_RESET: Service "bcm_usd" in vdc 1 has had a hap failure2020 Nov 21 09:45:37 %MODULE-2-MOD_DIAG_FAIL: Module 3 (Serial number: SAL2034ULTR) reported failure due to Service on linecard had a hap-reset in device DEV_SYSMGR (device error 0x30b)

Workarounds: None |
| CSCvw76953 | Headline: ip routing multicast holddown 0 does not propagate in running config

Symptoms: When trying to configure "ip routing multicast holddown 0" command, it does not propagate to running config.

Workarounds: It looks like problem is happening only for fresh configuraion. Remove/Re-apply of the 'ip routing multicast holddown 0' should solve the problem. |
| CSCvw78632 | Headline: Elysian_TOR: Wrong delivery order and incorrect Timestamp(T1) carried by Follow-up pkt

Symptoms: TOR switch(BC) port that the T1 timestamps carried in Follow-up pkts towards O-DU are behind in time w.r.t the Delay Response packets which are received before Follow-up pkt. Because of this the end offset computed by ptp4l servo goes high and this is screwing up the time and impacting the cell operations/sectors on O-DU.

Workarounds: No workaround, Dev is debugging the issue. |
| CSCvw78982 | Headline: High memory usage by "MRIB" process on 9.3(3)

Symptoms: amount of MRIB-3-MALLOC_FAILED logs on 9.3(3).2020 Dec 7 13:01:04 switch %MRIB-3-MALLOC_FAILED: mrib [517] mrib_smalloc() failed for mrib_mpib_oif_datatype2020 Dec 7 13:01:14 switch %MRIB-4-SYSLOG_SL_MSG_WARNING: MRIB-3-MALLOC_FAILED: message repeated 11017 times in last 3172422 sec

Workarounds: N/A |
| CSCvw80210 | Headline: On N9K-C9332C - 100G macsec links stop forwarding traffic

Symptoms: On N9K-C9332C & N9K-C9364C - 100G macsec links stop forwarding traffic after some time (depending on the traffic rate), if no sak-expiry-time is configured. This will result in connectivity loss and LACP suspension

Workarounds: Workaround 1) Configure XPN Cipher-suite instead of non-xpnWorkaound 2) Configure "sak-rekey-time 300" under macsec security policy |

Bug ID	Description
CSCvw81106	Headline: STOMP generated on heavenly when using 1G xcvr Symptoms: In a VxLAN fabric with –FX2 devices as leaf, packets over 740 bytes are not correctly forwarded over the fabric. Workarounds: Use 10G xcvr in between Leaf and Spine
CSCvw83503	Headline: Nexus 9000 forwards ARP traffic received on suspended interface, causing Layer 2 loop Symptoms: A Nexus 9000 switch running NX-OS software release 9.3(6) may begin forwarding ARP traffic received on a physical interface that is configured to be a member of an LACP port-channel, even though the physical interface is suspended from the port-channel for not receiving LACPDUs.switch# show port-channel summary interface port-channel 10Flags: D - Down P - Up in port-channel (members) I - Individual H - Hot-standby (LACP only) s - Suspended r - Module-removed b - BFD Session Wait S - Switched R - Routed U - Up (port-channel) p - Up in delay-lacp mode (member) M - Not in use. Min-links not met--Group Port-Type Protocol Member Ports Channel---18 Po10(SU) Eth LACP Eth1/49(s) Eth1/50(P)switch# show interface Ethernet1/49-50 counters <snip>---Port InOctets InUcastPkts---Eth1/49 54152856037 795819621 <<<Eth1/50 76792969 470720---Port OutOctets OutUcastPkts---Eth1/49 34045 0Eth1/50 54149810013 795971744 <<<An ELAM confirms that an ARP packet received on Ethernet1/49 is forwarded out of Ethernet1/50:switch(TAH-elam-insel6)# reportSUNDOWN1 ELAM REPORT SUMMARYslot - 1, asic - 0, slice - 0===========================Incoming Interface: Eth1/49Src Idx : 0x602, Src BD : 10Outgoing Interface Info: dmod 1, dpid 20Dst Idx : 0x602, Dst BD : 10Packet Type: ARPDst MAC address: A8:0C:0D:9B:60:BFSrc MAC address: A8:0C:0D:9B:68:3F.1q Tag0 VLAN: 10, cos = 0x6Target Hardware address: A8:0C:0D:9B:60:BFSender Hardware address: A8:0C:0D:9B:68:3FTarget Protocol address: 192.0.2.10Sender Protocol address: 192.0.2.20ARP opcode: 2Drop Info:----------LUA:LUB:LUC:LUD:Final Drops:vntag:vntag_valid : 0vntag_vir : 0vntag_svif : 0switch(TAH-elam-insel6)# show system internal ethpm info all | include dmod=1,dpid=20 IF_STATIC_INFO: port_name=Ethernet1/50,if_index:0x1a006200,ltl=5948,slot=0, nxos_port=196,dmod=1,dpid=20,unit=0,queue=65535,xbar_unitbmp=0x0,ns_pid=255,slice_num=0 ,port_on_slice=20,src_id=40 Workarounds: This issue can be resolved through one of two actions:1. Administratively shut down the suspended port-channel member.2. Reconfigure the remote network device such that the remote network device's interface sends LACPDUs as expected. This will bring the suspended port-channel member out of a suspended state.
CSCvw84051	Headline: ISIS routes are not removed from routing table when the interface goes down. Symptoms: When ISIS configured interface goes down, the routes are still present in the routing table and causes connectivity issue. Workarounds: Check the table-map configuration and remove it if not required. If table-map is required, make sure corresponding route-map is also configured.ORclear ip route x.x.x.x
CSCvw84453	Headline: src_mac is 00:00:00:00:00:00 afer PBR routing after rebooting of one Nexus in VPC pair Symptoms: src_mac is 00:00:00:00:00:00 afer PBR routing after rebooting of one Nexus in VPC pair Workarounds: 1. Link shutdown/no shutdown2.Delete/create static MAC
CSCvw86078	Headline: N9k:–FX and Other Platforms - DHCP Binding lease will not refresh Symptoms: If have approximate 2000 dhcp snooping entries in N9K after reloading the dhcp

Bug ID	Description
	snooping binding will not refresh. Workarounds: Remove DHCP binding and re-configure
CSCvw92365	Headline: cosmetic: Next Hop v6 Filter does not work properly on the show ip route output. Symptoms: show ip route next-hop-v6 X:X:X:X not only shows ipv6 next-hop routes. From the command line we can see it?s only display the v6 next-hop route, but from the test result, it?s also shows the v4 routes. When we use ?show ipv6 route next-hop X:X:X:X? it?s only display the v6 routes. Below is the test output on 9.3(6) N9K-C93180YC-EX-1# show ip route ? next-hop-v6 Display routes with this V6 next-hop only R4# show ip route next-hop-v6 2001::1IP 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> 1.1.1.2/32, ubest/mbest: 1/0 *via 2001::1%default, Eth1/1, [20/0], 00:01:12, bgp-65444, external, tag 652221.1.1.3/32, ubest/mbest: 1/0 *via 10.1.34.10, [20/0], 01:15:13, bgp-65444, external, tag 653331.1.1.4/32, ubest/mbest: 2/0, attached *via 1.1.1.4, Lo0, [0/0], 01:25:42, local, tag 100 *via 1.1.1.4, Lo0, [0/0], 01:25:42, direct, tag 10010.1.34.10/31, ubest/mbest: 1/0, attached *via 10.1.34.11, Eth1/2, [0/0], 01:24:19, direct10.1.34.10/32, ubest/mbest: 1/0, attached *via 10.1.34.10, Eth1/2, [250/0], 01:24:17, am10.1.34.11/32, ubest/mbest: 1/0, attached *via 10.1.34.11, Eth1/2, [0/0], 01:24:19, local20.1.1.0/30, ubest/mbest: 1/0 *via 2001::1%default, Eth1/1, [20/0], 00:01:12, bgp-65444, external, tag 6522230.1.1.0/30, ubest/mbest: 1/0 *via 10.1.34.10, [20/0], 01:15:13, bgp-65444, external, tag 65333 R4# show ipv6 route next-hop 2001::1IPv6 Routing Table for VRF "default" '*' denotes best ucast next-hop'**' denotes best mcast next-hop'[x/y]' denotes [preference/metric] 2001::1/128, ubest/mbest: 1/0, attached*via 2001::1, Eth1/1, [250/0], 00:01:56, am Workarounds: NA
CSCvw92732	Headline: Vfc interface mode can't be set to E Symptoms: VFC interface port mode config is not configurable as E.This is the expected behaviour. Currently, there is no support for PORT mode E for vfc links. Workarounds: Please do not configure switchport mode as E for vfc links.
CSCvw94313	Headline: N9k: port-security does not effect on the VPC port-channel Symptoms: N9k: Port-security does not effect on the VPC port-channel. Violated traffic can pass though Workarounds: NA
CSCvx02717	Headline: After upgrade 9.x N9k cannot be downgraded via install all Symptoms: Upgrade from 9.3(3) to 9.3(6) then when testing downgrade back via "install all" there are incompatibilities that prevent the downgrade:N9K-C9336-FX2-Z-PI# show incompatibility-all nxos nxos.9.3.3.binChecking incompatible configuration(s) for vdc 'N9K-C9336-FX2-Z-PI':--The following configurations on active are incompatible with the system image 1) Service : rpm , Capability : CAP_FEATURE_RPM_TUNNELENCRYPT_KEYSTR_NOSHOWDescription : RPM tunnel-encryption key chain command "key-chain tunnelencrypt-psk no-show" is configuredCapability requirement : STRICTEnable/Disable command : Use "no key-chain tunnelencrypt-psk no-show" to remove it2) Service : rpm , Capability : CAP_FEATURE_RPM_TUNNELENCRYPT_KEYCHAINDescription : RPM tunnel-encryption key chain command "key chain <keychain> tunnel-encryption" is configuredCapability requirement : STRICTEnable/Disable command : Use "no key chain <keychain> tunnel-encryption" to remove it3) Service : rpm , Capability : CAP_FEATURE_IP_NEXTHOP_VERIFY_FORCE_ORDERDescription : RPM-PBR "set ip next-hop verify-availability <ip1> ... force-order" is configuredCapability requirement : STRICTEnable/Disable command : Please Use "no set ip next-hop verify-availability force-order" to remove the configuration4) Service : rpm , Capability : CAP_FEATURE_IP_NEXTHOP_VERIFY_DROP_ON_FAILDescription : RPM-PBR "set ip next-hop verify-availability <ip1> ... drop-on-fail" is configuredCapability requirement : STRICTEnable/Disable command : Please Use "no set ip next-hop verify-availability drop-on-fail"

Bug ID	Description
	to remove the configuration5) Service : rpm , Capability : CAP_FEATURE_IPV6_NEXTHOP_VERIFY_FORCE_ORDERDescription : RPM-PBR "set ipv6 next-hop verify-availability <ip1> ... force-order" is configuredCapability requirement : STRICTEnable/Disable command : Please Use "no set ipv6 next-hop verify-availability force-order" to remove the configuration6) Service : rpm , Capability : CAP_FEATURE_IPV6_NEXTHOP_VERIFY_DROP_ON_FAILDescription : RPM-PBR "set ipv6 next-hop verify-availability <ip1> ... drop-on-fail" is configuredCapability requirement : STRICTEnable/Disable command : Please Use "no set ipv6 next-hop verify-availability drop-on-fail" to remove the configuration7) Service : rpm , Capability : CAP_FEATURE_RPM_MACSEC_TYPE_6_KEYCHAINDescription : RPM type-6-encryption key chain command "key-octet-string <> <> cryptographic-algorithm <>" is configuredCapability requirement : STRICTEnable/Disable command : Use "no key-octet-string" to remove it8) Service : rpm , Capability : CAP_FEATURE_RPM_EVPN_GWIP_USE_NEXTHOPDescription : The following EVPN CLIs are configured on route-map"set evpn gateway-ip use-nexthop" Capability requirement : STRICTEnable/Disable command : Please remove this configuration from all route-map9) Service : rpm , Capability : CAP_FEATURE_RPM_COMM_LIST_GRACEFUL_SHUTDOWNDescription : Community list standard "graceful-shutdown community" is configuredCapability requirement : STRICTEnable/Disable command : please remove the graceful-shutdown configuration and retryChecking dynamic incompatibilities:-------------------------------No incompatible configurationsN9K-C9336-FX2-Z-PI#Configurations listed above are not present in the show run or show run all. Attempting to remove per the instructions above does not fix. Write erase, reload, then downgrade of the N9k via "install all" is required. Workarounds: Write erase, reload, then downgrade of the N9k via "install all"
CSCvx03005	Headline: Nexus 9000 floods ARP Request with unicast destination MAC to VLAN Symptoms: A Nexus 9000 switch may occasionally generate or receive an ARP Request that refreshes an existing ARP entry. This ARP Request packet is a unicast packet with a destination MAC address of another device on the network. However, this ARP Request packet is flooded out of all interfaces carrying the relevant VLAN as if it were a normal broadcast ARP Request packet. Workarounds: There is no known workaround for this issue at this time.
CSCvx04061	Headline: Incorrect NIV programming after changing from VPC HIF to orphan port Symptoms: Traffic traversing vPC peer-link in communication with orphan host connected behind single homed FEX is dropped.ELAM report will show the following drop reason: UC_DF_CHECK_FAIILURE Workarounds: None
CSCvx04916	Headline: segmentation fault for Tunnel Manager (TM) Daemon in CLI for Tunnel mgr process Symptoms: +Crash of Tunnel Manager (TM) Daemon+No noticeable flapping of tunnels, they were all up; normally, it would come up fast if it goes down (from customer perspective). "show logging logfile" and "show logging nvram" do not show any tunnel flaps. Workarounds: tbd
CSCvx07013	Headline: No negotiate auto command missing after upgrade to 7.0(3)I7(9) Symptoms: when upgrading from 7.0(3)I7(6) to 7.0(3)I7(9), the command "no negotiate auto" is missing from the interfaces:version 7.0(3)I7(6) Bios:version 08.36 interface Ethernet1/2 switchport speed 1000 no negotiate auto no shutdownBRU-N9K5-7# show run int eth2/1!Command: show running-config interface Ethernet2/1!No configuration change since last restart!Time: Mon Nov 30 00:46:00 2020version 7.0(3)I7(9) Bios:version 08.36 interface Ethernet2/1 switchport speed 1000 no shutdownThis caused the interfaces in customer setup to go down and we have to manually configure "no negotiate auto" under the interface configuration to recover. Workarounds: we have to manually configure "no negotiate auto" under the interface configuration to recover.

Bug ID	Description
CSCvx07403	Headline: N9K - Some 'redistribute direct' prefixes not advertised to eBGP peer. Symptoms: + Prefixes are not advertised to eBGP neighbors.+ Redistribution into BGP is via network statement and 'redistribute direct' Workarounds: 'clear ip bgp * soft' will restore all prefixes being advertised.
CSCvx09137	Headline: Nexus 9K Linecard Memory Leak in /tmp/logs/l2mcast_lib.log Symptoms: A Nexus 9k switch running NX-OS 9.3(6) may begin printing the below logs indicating one of the linecards has high memory usage in a tmp directory: %SYSMGR-SLOTX-2-TMP_DIR_FULL: System temporary directory usage is unexpectedly high at 100%.In addition, if you attach into the impacted linecard and run "show system internal flash", the "aufs" filesystem will be fully used: SWITCH# attach module x module-x# show system internal flash <snip> aufs 2097152 2097152 0 100% /lcFinally, when checking the filesystem more thoroughly via the bash shell, a file at "/tmp/logs/l2mcast_lib.log" will be growing abnormally large (replace "x" in "login lcx" with the slot number of the impacted linecard): SWITCH# conf t SWITCH(config)# feature bash SWITCH(config)# exit SWITCH# run bash bash-4.2$ sudo bash bash-4.2# rlogin lcx root@lcx:/# ls -la /tmp/logs total 15400 drwxr-xr-x 2 root floppy 60 Jan 1 12:22 . drwxrwxrwx 3 root floppy 600 Jan 1 15:22 .. -rw-rw-rw- 1 root root 15766000 Jan 1 15:44 l2mcast_lib.log <===== Workarounds: The offending file can be periodically deleted, eg: every 24 hours, via a very basic shell script (replace the "x" in "lcx" with the slot number of the impacted linecard):conf t feature bashexitrun bash rlogin lcx while [True]; do echo "Deleting l2mcast_lib.log..."; rm /tmp/logs/l2mcast_lib.log; sleep 86400; done & exitexit
CSCvx18993	Headline: ON portfast enabled interface STP goes through BLK-LRN-FWD convergence after VPC sec turns Op. Prim Symptoms: On portfast enabled interface, STP goes through BLK-LRN-FWD convergence after VPC seceondary turns to operational primary. Workarounds: Shut/ no shut interface bring it back without convergence wait.
CSCvx24711	Headline: 'lcnd: Vxlan-InnerIp size xx is less than expected' is shown in syslog Symptoms: The following message may be shown in syslog;%KERN-3-SYSTEM_MSG: lcnd: Vxlan-InnerIp size 46 is less than expected – kernel Workarounds: No workaround.But the message is harmless and can be ignored.
CSCvx26057	Headline: Unicast traffic punted to CPU due to HW ADJ pointing to global glean adj Symptoms: Unicast traffic punted to CPU and dropped by COPP.HW adj is pointed to global glean adj after route flap or change. Workarounds: None
CSCvt97607	Headline: 9736C-FX imposing two FCS headers on RX frames after link flap Symptoms: A Cisco Nexus 9500 Series Switch with 9736C-FX line cards might experience a condition in which all ingress traffic on a port incorrectly has two FCS's appended to the end of the frame instead of one. Both the outer and inner FCS calculations will be valid, so traffic will not be dropped unless the four extra bytes of FCS causes traffic to exceed MTU, or if a given protocol performs length validity checks on frames. Workarounds: Configure "diagnostic bootup level minimal" and reload the device.

Known Issues

Bug ID	Description
CSCvw60588	10G QSA link in PCS mode is not working even without Macsec configuration.
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 9.3(7) 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

Bug 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-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-EX	Cisco Nexus 9500 32-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-X9736C-EX	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet QSFP28 line card	4	8	16

Product ID	Description	Maximum Quantity		
		Cisco Nexus 9504	Cisco Nexus 9508	Cisco Nexus 9516
N9K-X9736C-FX	Cisco Nexus 9500 36-port 40/100 Gigabit Ethernet 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

Table 3.   Cisco Nexus 9500 R-Series Line Cards

Product ID	Description	Maximum Quantity	
		Cisco Nexus 9504	Cisco Nexus9508
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

Table 4.   Cisco Nexus 9500 Classic Line Cards

Product ID	Description	Maximum Quantity		
		Cisco Nexus 9504	Cisco Nexus 9508	Cisco Nexus 9516
N9K-X9408C-CFP2	Line card with 8 100 Gigabit CFP2 ports	4	8	16
N9K-X9432C-S	Cisco Nexus 9500 32-port 40/100 Gigabit Ethernet QSFP28 line card	4	8	N/A
N9K-X9432PQ	Cisco Nexus 9500 32-port 40 Gigabit Ethernet QSFP+ line card	4	8	16
N9K-X9636PQ	Cisco Nexus 9500 36-port 40 Gigabit Ethernet QSFP+ line card	4	8	N/A
N9K-X9464PX	Cisco Nexus 9500 48 1/10-Gigabit SFP+ and 4-port 40-Gigabit Ethernet QSFP+ line card	4	8	16
N9K-X9464TX	Cisco Nexus 9500 48 port 1/10-Gigabit BASE-T Ethernet and 4-port 40-Gigabit Ethernet QSFP+ line	4	8	16

Product ID	Description	Maximum Quantity		
		Cisco Nexus 9504	**Cisco Nexus 9508**	**Cisco Nexus 9516**
	card			
N9K-X9464TX2	Cisco Nexus 9500 48 port 1/10-Gigabit BASE-T Ethernet and 4-port 40-Gigabit Ethernet QSFP+ line card	4	8	16
N9K-X9536PQ	Cisco Nexus 9500 36-port 40 Gigabit Ethernet QSFP+ line card	4	8	16
N9K-X9564PX	Cisco Nexus 9500 48 1/10-Gigabit SFP+ and 4 port 40-Gigabit Ethernet QSFP+ line card	4	8	16
N9K-X9564TX	Cisco Nexus 9500 48 port 1/10-Gigabit BASE-T Ethernet and 4 port 40-Gigabit Ethernet QSFP+ line card	4	8	16

Table 5.   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-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-C9516-FM-E	Cisco Nexus 9516 50-Gigabit cloud scale fabric module	4	5
N9K-C9516-FM-E2	Cisco Nexus 9516 100-Gigabit cloud scale fabric module	4	5

Table 6.   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
N9K-C9508-FM-R	Cisco Nexus 9508 100-Gigabit R-Series fabric module	4	6

Table 7.   Cisco Nexus 9500 Fabric Modules

Product ID	Description	Minimum	Maximum
N9K-C9504-FM	Cisco Nexus 9504 40-Gigabit fabric module	3	6

Product ID	Description	Minimum	Maximum
N9K-C9508-FM	Cisco Nexus 9508 40-Gigabit fabric module	3	6
N9K-C9516-FM	Cisco Nexus 9516 40-Gigabit fabric module	3	6
N9K-C9504-FM-S	Cisco Nexus 9504 100-Gigabit fabric module	4	4
N9K-C9508-FM-S	Cisco Nexus 9508 100-Gigabit fabric module	4	4

Table 8.   Cisco Nexus 9500 Fabric Module Blanks with Power Connector

Product ID	Description	Minimum	Maximum
N9K-C9508-FM-Z	Cisco Nexus 9508 Fabric blank with Fan Tray Power Connector module	N/A	2
N9K-C9516-FM-Z	Cisco Nexus 9516 Fabric blank with Fan Tray Power Connector module	N/A	2

Table 9.   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 10.   Cisco Nexus 9500 System Controller

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

Table 11.   Cisco Nexus 9500 Fans and Fan Trays

Product ID	Description	Quantity
N9K-C9504-FAN	Fan tray for 4-slot modular chassis	3
N9K-C9508-FAN	Fan tray for 8-slot modular chassis	3
N9K-C9516-FAN	Fan tray for 16-slot modular chassis	3

Table 12. 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 13. Cisco Nexus 9200 and 9300 Fans and Fan Trays

Product ID	Description	Quantity	Cisco Nexus Switches
N9K-C9300-FAN1	Fan 1 module with port-side intake airflow (burgundy coloring)	3	9396PX (early versions)
N9K-C9300-FAN1-B	Fan 1 module with port-side exhaust airflow (blue coloring)	3	9396PX (early versions)
N9K-C9300-FAN2	Fan 2 module with port-side intake airflow (burgundy coloring)	3	93128TX 9396PX 9396TX
N9K-C9300-FAN2-B	Fan 2 module with port-side exhaust airflow (blue coloring)	3	93128TX 9396PX 9396TX
N9K-C9300-FAN3	Fan 3 module with port-side intake airflow (burgundy coloring)	3	92304QC 9272Q[a] 93120TX
N9K-C9300-FAN3-B	Fan 3 module with port-side exhaust airflow (blue coloring)	3	92304QC 9272Q[a] 93120TX
NXA-FAN-160CFM-PE	Fan module with port-side exhaust airflow (blue coloring)	3	9364C[a] 93360YC-FX2
NXA-FAN-160CFM-PI	Fan module with port-side intake airflow (burgundy coloring)	3	9364C[a] 93360YC-FX2
NXA-FAN-160CFM2-PE	Fan module with port-side exhaust airflow (blue coloring)	4	9364C-GX

Product ID	Description	Quantity	Cisco Nexus Switches
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	92160YC-X 9236C[a] 93108TC-EX 93108TC-FX[a] 93180LC-EX[a] 93180YC-EX 93180YC-FX[a] 9332PQ 9372PX 9372PX-E 9372TX 9372TX-E 9348GC-FXP[a]
NXA-FAN-30CFM-F	Fan module with port-side exhaust airflow (blue coloring)	3	92160YC-X 9236C[a] 93108TC-EX 93108TC-FX[a] 93180LC-EX[a] 93180YC-EX 93180YC-FX[a] 9332PQ 9372PX 9372PX-E 9372TX 9372TX-E 9348GC-FXP
NXA-FAN-35CFM-PE	Fan module with port-side exhaust airflow (blue coloring)	4	92300YC[a] 9332C[a] 93108TC-FX3P 93180YC-FX3S[b]
		6	9316D-GX 93600CD-GX
NXA-FAN-35CFM-PI	Fan module with port-side intake airflow (burgundy coloring)	4	92300YC[a] 9332C[a] 93108TC-FX3P 93180YC-FX3S[b]
		6	9316D-GX 93600CD-GX
NXA-FAN-65CFM-PE	Fan module with port-side exhaust airflow (blue coloring)	3	93240YC-FX2[a] 9336C-FX2[a]
NXA-FAN-65CFM-PI	Fan module with port-side exhaust airflow (burgundy coloring)	3	93240YC-FX2[a] 9336C-FX2[a]

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

[b] 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.

Table 14. 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 93180LC–EX 93180YC–EX 93180YC–FX
NXA-PAC-500W-PI	500-W AC power supply with port–side intake airflow (burgundy coloring)	2	93108TC–EX 93180LC–EX 93180YC–EX 93180YC–FX
N9K-PAC-650W	650-W AC power supply with port–side intake (burgundy coloring)	2	9332PQ 9372PX 9372PX–E 9372TX 9372TX–E 9396PX 9396TX
N9K-PAC-650W-B	650-W AC power supply with port–side exhaust (blue coloring)	2	9332PQ 9372PX 9372PX–E 9372TX 9372TX–E 9396PX 9396TX
NXA-PAC-650W-PE	650-W power supply with port–side exhaust (blue coloring)	2	92160YC–X 9236C 92300YC 93180YC–FX3S 92304QC 93108TC–EX 93180YC–EX
NXA-PAC-650W-PI	650-W power supply with port–side intake (burgundy coloring)	2	92160YC–X 9236C 92300YC 93180YC–FX3S 92304QC 93108TC–EX 93180YC–EX
NXA-PAC-750W-PE	750-W AC power supply with port–side exhaust airflow (blue coloring) [1]	2	9336C–FX2 93240YC–FX2 9332C 9336C–FX2
NXA-PAC-750W-PI	750-W AC power supply with port–side exhaust airflow (burgundy coloring) [1]	2	9336C–FX2 93240YC–FX2 9332C 9336C–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 93600CD–GX
NXA-PAC-1100W-PI2	1100-W AC power supply with port–side intake airflow	2	93240YC–FX2

Product ID	Description	Quantity	Cisco Nexus Switches
	(burgundy coloring)		9332C 9316D-GX 9336C-FX2 93600CD-GX
NXA-PAC-1100W-PI	Cisco Nexus 9000 PoE 1100W AC PS, port-side intake	2	93108TC-FX3P
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
N9K-PAC-1200W	1200-W AC power supply with port-side intake airflow (burgundy coloring)	2	93120TX
N9K-PAC-1200W-B	1200-W AC power supply with port-side exhaust airflow (blue coloring)	2	93120TX
NXA-PAC-1200W-PE	1200-W AC power supply with port-side exhaust airflow (blue coloring)	2	9272Q 93360YC-FX2 9364C
NXA-PAC-1200W-PI	1200-W AC power supply with port-side intake airflow (burgundy coloring)	2	9272Q 93360YC-FX2 9364C
N9K-PUV-1200W	1200-W Universal AC/DC power supply with bidirectional airflow (white coloring)	2	92160YC-X 9236C 92300YC 92304QC 9272Q[1] 93108TC-EX 93108TC-FX 93360YC-FX2 93180YC-FX3S 93120TX 93128TX 93180LC-EX 93180YC-EX 93180YC-FX 9364C
NXA-PDC-930W-PE	930-W DC power supply with port-side exhaust airflow (blue coloring)	2	9272Q 93108TC-EX 93180YC-EX 93360YC-FX2 93180YC-FX3S 93120TX 93180YC-FX 9364C 92160YC-X
NXA-PDC-930W-PI	930-W DC power supply with port-side intake airflow (burgundy coloring)	2	9272Q 93108TC-EX 93180YC-EX 93360YC-FX2 93180YC-FX3S 93120TX 93180YC-FX 9364C

Product ID	Description	Quantity	Cisco Nexus Switches
			92160YC-X
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
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
UCSC-PSU-930WDC	930-W DC power supply with port-side intake (green coloring)	2	92160YC-X 9236C 92304QC 9272Q 93108TC-EX 93120TX 93128TX 93180YC-EX 9332PQ 9372PX 9372PX-E 9372TX 9372TX-E 9396PX 9396TX
UCS-PSU-6332-DC	930-W DC power supply with port-side exhaust (gray coloring)	2	92160YC-X 9236C 92304QC 9272Q 93108TC-EX 93120TX 93128TX 93180YC-EX 9332PQ 9372PX 9372PX-E 9372TX 9372TX-E 9396PX 9396TX
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 (burgundy coloring)	2	93240YC-FX2 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

Product ID	Description	Quantity	Cisco Nexus Switches
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

Table 15.   Cisco Nexus 9200 and 9300 Switches

Cisco Nexus Switch	Description
N9K-C92160YC-X	1-RU Top-of-Rack switch with 48 10-/25-Gigabit SFP+ ports and 6 40-Gigabit QSFP+ ports (4 of these ports support 100-Gigabit QSFP28 optics).
N9K-C92300YC	1.5-RU Top-of-Rack switch with 48 10-/25-Gigabit SFP28 ports and 18 fixed 40-/100-Gigabit QSFP28 ports.
N9K-C92304QC	2-RU Top-of-Rack switch with 56 40-Gigabit Ethernet QSFP+ ports (16 of these ports support 4x10 breakout cables) and 8 100-Gigabit QSFP28 ports.
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 mpps. 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-C9236C	1-RU Top-of-Rack switch with 36 40-/100-Gigabit QSFP28 ports (144 10-/25-Gigabit ports when using breakout cables)
N9K-C9272Q	2-RU Top-of-Rack switch with 72 40-Gigabit Ethernet QSFP+ ports (35 of these ports also support 4x10 breakout cables for 140 10-Gigabit 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-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-C93120TX	2-RU Top-of-Rack switch with 96 1/10GBASE-T (copper) ports and 6 40-Gigabit QSFP+

Cisco Nexus Switch	Description
	ports
N9K-C93128TX	3-RU Top-of-Rack switch with 96 1/10GBASE-T (copper) ports and an uplink module up to 8 40-Gigabit QSFP+ ports
N9K-C9316D-GX	1-RU switch with 16x400/100/40-Gbps ports.
N9K-C93180LC-EX	1-RU Top-of-Rack switch with 24 40-/50-Gigabit QSFP+ downlink ports and 6 40/100-Gigabit uplink ports. You can configure 18 downlink ports as 100-Gigabit QSFP28 ports or as 10-Gigabit SFP+ ports (using breakout cables).
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-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-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) 6 10/25/40/50/100-Gigabit QSFP28 ports (ports 49-54)
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-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-C9332C	1-RU fixed switch with 32 40/100-Gigabit QSFP28 ports and 2 fixed 1/10-Gigabit SFP+ ports.
N9K-C9332PQ	1-RU switch with 32 40-Gigabit Ethernet QSFP+ ports (26 ports support 4x10 breakout cables and 6 ports support QSFP-to-SFP adapters)
N9K-C93360YC-FX2	2-RU switch with 96 10-/25-Gigabit SFP28 ports and 12 40/100-Gigabit QSFP28 ports
N9K-C9336C-FX2	1-RU switch with 36 40-/100-Gb Ethernet QSFP28 ports.
N9K-C9348GC-FXP	Nexus 9300 with 48p 100M/1 G, 4p 10/25 G SFP+ and 2p 100 G QSFP
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.

Cisco Nexus Switch	Description
	– Ports 49 to 64 support MACsec encryption.
	Ports 65 and 66 support 1/10 Gigabit speeds.
N9K-C9364C-GX	2-RU fixed-port switch with 64 100-Gigabit SFP28 ports.
N9K-C9372PX	1-RU Top-of-Rack switch with 48 1-/10-Gigabit SFP+ ports and 6 40-Gigabit QSFP+ ports
N9K-C9372PX-E	An enhanced version of the Cisco Nexus 9372PX-E switch.
N9K-C9372TX	1-RU Top-of-Rack switch with 48 1-/10GBASE-T (copper) ports and 6 40-Gigabit QSFP+ ports
N9K-C9372TX-E	An enhanced version of the Cisco Nexus 9372TX-E switch.
N9K-C9396PX	2-RU Top-of-Rack switch with 48 1-/10-Gigabit Ethernet SFP+ ports and an uplink module with up to 12 40-Gigabit QSFP+ ports
N9K-C9396TX	2-RU Top-of-Rack switch with 48 1/10GBASE-T (copper) ports and an uplink module with up to 12 40-Gigabit QSFP+ ports

Table 16.   Cisco Nexus 9000 Series Uplink Modules

Cisco Nexus Switch	Description
N9K-M4PC-CFP2	Cisco Nexus 9300 uplink module with 4 100-Gigabit Ethernet CFP2 ports. For the Cisco Nexus 93128TX switch, only two of the ports are active. For the Cisco Nexus 9396PX and 9396TX switches, all four ports are active.
N9K-M6PQ	Cisco Nexus 9300 uplink module with 6 40-Gigabit Ethernet QSFP+ ports for the Cisco Nexus 9396PX, 9396TX, and 93128TX switches.
N9K-M6PQ-E	An enhanced version of the Cisco Nexus N9K-M6PQ uplink module.
N9K-M12PQ	Cisco Nexus 9300 uplink module with 12 40-Gigabit Ethernet QSPF+ ports.

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 Network Insights for Data Center

Cisco NX-OS Release 9.3(7) supports the Cisco Network Insights Advisor (NIA) and Cisco Network Insights for Resources (NIR) on Cisco Nexus 9200, 9300-EX, and 9300-FX platform switches and 9500 platform switches with -EX/FX line cards. For more information, see the Cisco Network 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 9.3(x). For information about an In Service Software Upgrade (ISSU), see the [Cisco NX-OS ISSU Support Matrix.](#)

Exceptions

Cisco Nexus 9200, 9300-EX, and 9300-FX Platform Switches

The following features are not supported for the Cisco Nexus 9200, 9300-EX, and 9300-FX platform switches:

- 64-bit ALPM routing mode
- Cisco Nexus 9272PQ and Cisco Nexus 92160YC platforms do not support the PXE boot of the Cisco NX-OS image from the loader.
- ACL filters to span subinterface traffic on the parent interface
- Egress port ACLs
- Egress QoS policer (not supported for Cisco Nexus 9200 platform switches). The only policer action supported is drop. Remark action is not supported on the egress policer.
- FEX (not supported for Cisco Nexus 9200 platform switches)
- GRE v4 payload over v6 tunnels
- IP length-based matches
- IP-in-IP (not supported on the Cisco Nexus 92160 switch)
- Maximum Transmission Unit (MTU) checks for packets received with an MPLS header
- NetFlow (not supported on Cisco Nexus 9200 platform switches)
- Packet-based statistics for Traffic Storm Control (only byte-based statistics are supported)
- PVLANs (not supported on Cisco Nexus 9200 platform switches)
- PXE boot of the Cisco NX-OS image from the loader (not supported for Cisco Nexus 9272PQ and 92160YC switches)
- Q-in-VNI (not supported on Cisco Nexus 9200 platform switches)
- Q-in-Q for VXLAN (not supported on Cisco Nexus 9200 and 9300-EX platform switches)
- Q-in-VNI (not supported on Cisco Nexus 9200 platform switches)
- Resilient hashing for port channels
- Rx SPAN for multicast if the SPAN source and destination are on the same slice and no forwarding interface is on the slice
- SVI uplinks with Q-in-VNI (not supported for Cisco Nexus 9300-EX platform switches)
- Traffic Storm Control for copy-to-CPU packets
- Traffic Storm Control with unknown multicast traffic
- Tx SPAN for multicast, unknown multicast, and broadcast traffic

- VACL redirects for TAP aggregation

Cisco Nexus 9300-FX3 Platform Switches

The following features are not supported for the Cisco Nexus 9300-FX3 Platform switches:

- ACL with DSCP Wildcard Mask
- ARP Suppression with Reflective Relay
- Dynamic ACL - Named ACL support for applying blacklist/limited VLAN access for devices
- ECMP Hashing based on GRE Inner IP Header
- Enhanced ISSU
- Enhanced Policy-Based Routing (ePBR)
- ePBR Multi-Hop
- ePBR with Probes
- ePBR with User-Defined Probes
- IPv6 MIB support (IP-MIB)
- Multicast Service Reflection (Ingress, PIM-border, Egress)
- Multiple LLDP neighbors per physical interface
- Secure VXLAN EVPN Multi-Site using CloudSec
- Selective Q-in-VNI + Advertise PIP on a VTEP
- Selective Q-in-VNI + VXLAN VLAN on the same port
- Standard ISSU
- Symmetric Hashing - ECMP (Inner DA)
- Unidirectional Ethernet (UDE)
- VXLAN EVPN with downstream VNI
- VXLAN over parent interface that also carries sub-interfaces

Cisco Nexus 9300-GX Platform Switches

The following features are not supported for the Cisco Nexus 9300-GX platform switches:

- Asymmetric PFC
- Autonegotiation on all ports
- FC-FEC for Cisco Nexus 9316D-GX and 93600CD-GX switches is not supported on the second lane of the 50x2 breakout port.
- FEX
- Multicast over GRE

Cisco Nexus N9K-X9408PC-CFP2 Line Card and 9300 Platform Switches

The following features are not supported for Cisco Nexus 9500 platform switches with the N9K-X9408PC-CFP2 line card and Cisco Nexus 9300 platform switches with generic expansion modules (N9K-M4PC-CFP2):

- 802.3x
- Breakout ports
- FEX (supported on some Cisco Nexus 9300 platform switches)
- Flows other than 40G
- Multichassis EtherChannel Trunk (MCT)
- NetFlow
- Port-channel (No LACP)
- PFC/LLFC
- Precision Time Protocol (PTP)
- PVLAN (supported on Cisco Nexus 9300 platform switches)
- Shaping support on 100g port is limited
- SPAN destination/ERSPAN destination IP
- Traffic Storm Control
- vPC
- VXLAN access port

FEX Modules

The following features are not supported for FEX modules:

- Active-Active FEX and straight-through FEX are not supported on the Cisco Nexus 92348GC switch.
- For Cisco Nexus 9500 platform switches, 4x10-Gb breakout for FEX connectivity is not supported.

Cisco Nexus N9K-X96136YC-R Line Card

The following features are not supported for Cisco Nexus 9500 platform switches with the N9K-X96136YC-R line card:

- Breakout
- PTP and gPTP

Cisco Nexus N9K-X9736C-FX Line Card

The following feature is not supported for Cisco Nexus 9500 platform switches with the N9K-X9736C-FX line card:

- Ports 29-36 do not support 1 Gbps speed.

Cisco Nexus 9500 Cloud Scale (EX/FX) Line Cards

The following features are not supported for Cisco Nexus 9500 platform switches with -EX/FX line cards:

- FEX

- IPv6 support for policy-based routing

- LPM dual-host mode

- SPAN port-channel destinations

Related Content

Cisco Nexus 9000 Series documentation: Cisco Nexus 9000 Series Switches

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 9.3(x)

Cisco Nexus 9000 Series FPGA/EPLD Upgrade Release Notes: Cisco Nexus 9000 Series FPGA/EPLD Upgrade Release Notes, Release 9.3(7)

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

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