



Cisco Nexus 9000 Series NX-OS Release Notes, Release 9.3(13)

Introduction

This document describes the features, issues, and exceptions of Cisco NX-OS Release 9.3(13) 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.

| Date | Description |
|-------------------|---|
| April 25, 2024 | Added CSCwh50989 to Open Issues. |
| December 18, 2023 | Cisco NX-OS Release 9.3(13) became available. |

New and Enhanced Software Features

There are no new or enhanced software and hardware features introduced in Cisco NX-OS Release 9.3(13).

Open Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug.

| Bug ID | Description |
|----------------------------|--|
| CSCvx62127 | <p>Headline: LICMGR crashes continuously and device reboots post issuing cli " sh lic all" .</p> <p>Symptoms: On upgrading the licmgr crashed continuously four times and hap_reset is triggered causing DUT to reboot. This happened on execution of CLI : " show license all"</p> <p>2021 Mar 9 09:01:38 switch %\$ VDC-1 %\$ %SYSMGR-2-SERVICE_CRASHED: Service " licmgr" (PID 20292) hasn't caught signal 11 (core will be saved).</p> <p>2021 Mar 9 09:01:38 TECTAE-PGBL-89 %\$ VDC-1 %\$ %SYSMGR-2-SERVICE_CRASHED: Service " licmgr" (PID 28524) hasn't caught signal 11 (core will be saved).</p> <p>2021 Mar 9 09:01:38 TECTAE-PGBL-89 %\$ VDC-1 %\$ %CALLHOME-2-EVENT: SW_CRASH</p> <p>2021 Mar 9 09:01:39 TECTAE-PGBL-89 %\$ VDC-1 %\$ %SYSMGR-2-SERVICE_CRASHED: Service " licmgr" (PID 28606) hasn't caught signal 11 (core will be saved).</p> <p>Workarounds: Change the host name so that it does not have any special characters.</p> |
| CSCwf54392 | <p>Headline: The " radius-server key <>" missing after ND ISSU.</p> <p>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 the change in default boot mode.</p> <p>Workarounds: Remove and reapply the missing CLI string to reconfigure and ensure it is consistent between configuration and DME.</p> |

| Bug ID | Description |
|----------------------------|--|
| CSCwf08452 | <p>Headline: Lag in entering commands in 93360YC-FX2.</p> <p>Symptoms: Most of the ports on these switches are populated with GLC-TE Transceivers. Any changes to the port configurations usually cause very long delays (minutes). This could be from a 'switchport' or 'shut/no shut' command.</p> <p>Have experimented with 'spanning-tree port type edge' and explicitly setting the speed on the port, but issues are still occurring. The switch also seems to take a very long time to recognize a new physical connection or GLC-TE reset. GLC-TEs are supported on this model according to the matrix page.</p> <p>Workarounds: Use 'no QOS statistics'.</p> |
| CSCwf37269 | <p>Headline: Nexus 9000 Cloud Scale switches with mix of static and dynamic queue-limit drops low-rate traffic.</p> <p>Symptoms:</p> <ol style="list-style-type: none"> 1. Egress interface reports incrementing output discards despite a low rate of traffic destined for that interface. 2. Incrementing occupancy drops under PG-Drop pool (slot X show hardware internal tah buffer counter). 3. The queue depth remains at 0 byte for the affected Egress interface (show queuing interface), which means that there is no actual congestion for that interface. 4. When sufficiently *reducing* the static queue-limit configuration or the number of ports configured with static queue-limit, under the same traffic conditions, output discards for the affected interface stop. <p>Workarounds: Use 'dynamic queue-limits'.</p> |
| CSCwh02830 | <p>Headline: Manual power cycle required after upgrading N9K-C9500 switch NXOS EPLD with SUP A+/SUP B+ V02</p> <p>Symptoms:</p> <ol style="list-style-type: none"> 1. Nexus 9500 chassis with single or dual SUP A+ or B+ HW V02 2. Upgrade of Sup or SC EPLD (Just NXOS upgrade will not trigger this issue.) 3. After the reload, the active SUP module does not boot up all the way and appears to be stuck with no console output and a continuously blinking amber STS LED. <p>Workarounds:</p> <ol style="list-style-type: none"> 1. Physically re-seat the module (cold power-cycle). 2. Remove all power to the chassis to allow the SUP module to power up again. |
| CSCwh72222 | <p>Headline: ARP request dropped by ARP Suppression in traditional VLAN</p> <p>Symptoms: ARP requests of hosts in a Non-VXLAN VLAN are dropped in the VTEP switch due to ARP suppression.</p> <p>Additionally, ARP request packets can be captured in CPU with Ethalyzer tool, even though the VLAN does not have an SVI.</p> <p>Workarounds: Configure an SVI for the affected VLAN. SVI can be shut down and does not need an IP address or reload the switch.</p> |

| Bug ID | Description |
|----------------------------|---|
| CSCwd13172 | <p>Headline: Nexus 9000 FX2 multicast-tx span is not working.</p> <p>Symptom:</p> <ol style="list-style-type: none"> 1. Multicast traffic is not mirrored from the configured source interface or 2. Multicast traffic would be mirrored from incorrect source interface or 3. Not a configured interface. <p>Workarounds: None</p> |
| CSCwi45695 | <p>Headline: During BIOS downgrade, N9K-X96XX modules may experience time-out.</p> <p>Symptom: BIOS downgrade for N9K-X96XX linecard modules may report failure due to time-out.</p> <p>Module 3: Refreshing compact flash and upgrading bios/loader/bootrom.</p> <p>Warning: please do not remove or power off the module at this time.</p> <p>[#####] 100% -- FAIL.</p> <p>Return code 0x40820011 (Job timedout).</p> <p>CAUTION: The BIOS/loader/bootrom of above module may be in corrupted state. Please try programming it again.</p> <p>Workarounds: If 'show install all impact ...' indicates a lower BIOS version than that currently installed on the module, avoid using the 'bios-force' flag during install operations.</p> |
| CSCwi44292 | <p>Headline: Pre-upgrade check on nxos64-msll images does not display new BIOS version correctly during downgrade.</p> <p>Symptom: During downgrade from impacted NXOS version to a previous version, pre-upgrade check displays New BIOS version as blank on following impacted platforms</p> <p>Impacted Software version:10.4(2),9.3(13)</p> <p>Impacted Platforms:</p> <ul style="list-style-type: none"> -N3K-C36180YC-R -N9K-X96XX with N9K-X9636Q-R,N9K-X9636C-RX, N9K-X9624D-R2,N9K-X96136YC-R line cards <p>Workarounds: None</p> |
| CSCwi36291 | <p>Headline: QSFP-100G40G-BIDI higher linkup time issue</p> <p>Symptom: Sometimes link up time may be higher on N9K-C9364C (port 49-64) or N9K-C9332C (port 25-32) when using QSFP-100G40G-BIDI optics.</p> <p>Workarounds: Use other ports on these PIDs or switch to another type of optics.</p> |
| CSCwh50989 | <p>Headline: Custom COPP causing transit traffic to be punted to the CPU on Nexus 9300-GX2</p> <p>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.</p> <p>Workarounds: Custom COPP policy using src/dst match mitigates punt for transit traffic.</p> |

Resolved Issues

Click the bug ID to access the Bug Search tool and see additional information about the bug.

| Bug ID | Description |
|----------------------------|--|
| CSCwa86549 | <p>Headline: Nexus 9000 M1100 sees I/O error, goes read-only then leaf crashes</p> <p>Symptoms: Dmesg reports I/O errors followed by remounting the filesystem as read-only:<3>[xxx] Buffer I/O error on device sda4, logical block 0 <4>[xxx] lost page write due to I/O error on sda4 <2>[xxx] EXT4-fs error (device sda4): ext4_journal_check_start:56: Detected aborted journal <2>[xxx] EXT4-fs (sda4): Remounting filesystem read-only <3>[xxx] EXT4-fs (sda4): previous I/O error to superblock detected <6>[xxx] sd 1:0:0:0: [sda] Unhandled error code <6>[xxx] sd 1:0:0:0: [sda] <6>[xxx] Result: hostbyte=DID_BAD_TARGET driverbyte=DRIVER_OK The switch can then proceed to crash due to going read-only</p> <p>Workarounds: Reloading the switch should bring it back to full operation</p> |
| CSCwd53591 | <p>Headline: Silent Reload due to watchdog timeout</p> <p>Symptoms: Silent Reload due to watchdog timeout</p> <p>Workarounds: No valid workaround so far</p> |
| CSCwf10364 | <p>Headline: kmalloc fail due to SSD writes causing memory frag/higher order page depletion</p> <p>Symptoms: Kmalloc failures due to excessive fragmentation</p> <p>Workarounds: Observe from buddyinfo:cat /proc/buddyinfo run bashsync && echo 3 > /proc/sys/vm/drop_caches echo 1 > /proc/sys/vm/compact_memory if the above doesn't work then reload box</p> |
| CSCwf40750 | <p>Headline: N9K C9336C-FX2 9.3(8) crashed while adding port channel to physical port</p> <p>Symptoms: Unexpected reload happened when adding port channel to physical port. Tahusd core file decode mentions functions related to adding a member to a port channel. The decode is not matching any from our known bugs.</p> <p>Workarounds: No workarounds known</p> |
| CSCwf48692 | <p>Headline: N9300-FX3S/FX3 may randomly timeout syncE peer</p> <p>Symptoms: Random syncE peer timeouts.</p> <p>Workarounds: N/A</p> |
| CSCwf58507 | <p>Headline: FEX 2348UPQ brings hosts link too fast after powercycle causing traffic blackholing for around 1min.</p> <p>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.</p> <p>Workarounds: Use LACP or any other SFP or upgrade the version.</p> |
| CSCwf82223 | <p>Headline: On N9K switch SNMP polling unable to retrieve QSFP DOM values.</p> <p>Symptoms: A third party DOM supported QSFP will not provide any DOM information when polled via SNMP.</p> <p>Workarounds: Physically reset QSFP28. OIR.</p> |

| Bug ID | Description |
|----------------------------|---|
| CSCwf93175 | <p>Headline: FIB and RIB become inconsistent after EIGRP route switchover</p> <p>Symptoms: When there are two EIGRP processes, if the route switchover occur due to a linkflap, the FIB and RIB become inconsistent.</p> <p>Workarounds: clear ip route x.x.x.x/x</p> |
| CSCwf97335 | <p>Headline: N9k Host routes transition with Punt index in v4host table after route change</p> <p>Symptoms: Multiple Symptoms may be reported.1) Traffic to certain IP addresses may end up having high response times or2) 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 Ethalyzer.</p> <p>Workarounds: Introduce back the AM or /32 Host routes</p> |
| CSCwh17302 | <p>Headline: Attached HMM /32 vrf leaking not working with maximum-paths mixed</p> <p>Symptoms: Trying to leak a /32 prefix from source VRF-A to target VRF-B on a pair of N9ks- Source-vrf is seeing the /32 prefix locally from HMM-Target VRF is not leaking the /32 prefix- Looking for the /32 prefix on the RIB only less specific route is leaked (less specific prefix is coming redistribute-direct of SVI subnet where /32 is attached)Source-vrf: VRF-ATarger-vrf: VRF-BSource vrf SVI: vlan100: 10.10.1.32/27Prefix on the source-vrf: 10.10.1.38/32vrf context VRF-A vni XXXX rd auto address-family ipv4 unicast route-target export AAAA:AAAA route-target export AAAA:AAAA evpnrouter bgp 65000 vrf VRF-A log-neighbor-changes address-family ipv4 unicast redistribute direct route-map allow <<< redistributing SVI subnet maximum-paths mixed 8 <<< maximum-paths mixed configuredvrf context VRF-B vni XXXX rd auto address-family ipv4 unicast route-target import AAAA:AAAA route-target import AAAA:AAAA evpnSource VRF prefix: Nexus-B# show ip route 10.10.1.38 vrf VRF-A IP Route Table for VRF "VRF-A" '*' denotes best ucast next-hop '*' denotes best mcast next-hop '[x/y]' denotes [preference/metric] '%<string>' in via output denotes VRF <string> 10.10.1.38/32, ubest/mbest: 1/0, attached *via 10.10.1.38, Vlan100, [190/0], 00:09:44, hmmChecking on target VRF, /32 prefix is not there, it is only leaking the /27 prefix (from the SVI/direct redistribution)Nexus-B# show ip route 10.10.1.38 vrf VRF-B IP Route Table for VRF "VRF-B" '*' denotes best ucast next-hop '*' denotes best mcast next-hop'[x/y]' denotes [preference/metric] '%<string>' in via output denotes VRF <string>10.10.1.32/27, ubest/mbest: 3/0, attached, all-best (0xc0a80102) <<<<<<< missing the /32, this is the /27 from the vlan100 subnet *via x.x.x.x%default, [20/0], 00:00:46, bgp-65000, external, tag 65000, segid: 2214 encap: VXLAN *via x.x.x.x%default, [200/0], 00:01:15, bgp-65000, internal, tag 65000, segid: 2214 (Asymmetric) tunnelid: 0xc0a80101 encap: VXLAN *via x.x.x.x%VRF-A, Vlan100, [20/0], 00:00:46, bgp-65000, external, tag 65000</p> <p>Workarounds: Option 1: Clearing the less specific prefix on the target vrf resolves the problemWrong prefix on the targert VRF (this isn't leaking the /32, it is taking the less specific prefix)Nexus-B# show ip route 10.10.1.38 vrf VRF-B IP Route Table for VRF "VRF-B" '*' denotes best ucast next-hop '*' denotes best mcast next-hop'[x/y]' denotes [preference/metric] '%<string>' in via output denotes VRF <string>10.10.1.32/27, ubest/mbest: 3/0, attached, all-best (0xc0a80102) <<<<< missing the /32, this is the /27 from the vlan100 subnet *via x.x.x.x%default, [20/0], 00:00:46, bgp-65000, external, tag 65049, segid: 2214 encap: VXLAN *via x.x.x.x%default, [200/0], 00:01:15, bgp-65000, internal, tag 65049, segid: 2214 (Asymmetric) tunnelid: 0xc0a80101 encap: VXLAN *via x.x.x.x%VRF-A, Vlan100, [20/0], 00:00:46, bgp-65000, external, tag 65049Nexus-B#clear routing vrf VRF-B 10.10.1.32/27Clearing 10.10.1.32/27Now /32 prefix is being leaked correctly: Nexus-B# show ip route 10.10.1.38 vrf VRF-B IP Route Table for VRF "VRF-B" '*' denotes best ucast next-hop '*' denotes best mcast next-hop '[x/y]' denotes [preference/metric] '%<string>' in via output denotes VRF <string> 10.10.1.38/32, ubest/mbest: 1/0, attached *via 10.10.1.38%VRF-A, Vlan100, [190/0], 00:17:03, attached-exportOption2: Remove the <maximum-paths mixed> in the source VRF:router bgp xxxxx vrf VRF-A log-neighbor-changes address-family ipv4 unicast redistribute direct route-map allow no maximum-paths mixed 8 <<<<<<</p> |

| Bug ID | Description |
|----------------------------|---|
| CSCwh18918 | <p>Headline: N9K- python3 crashes being observed after upgrade</p> <p>Symptoms: Several python3 core files are created after upgrading to 10.2(x) and higher releases from 9.3(x) You can use the command show cores to see the cores generated. n9k-switch# show cores</p> <pre>VDC Module Instance Process-name PID Date(Year-Month-Day Time)--- ----- --- ----- -----1 1 1 python3 --</pre> <p>Additional issue that can be seen on upgrade from 10.2(x) and higher is that the displayed timezone in "show clock" output may not match the configuration. For example: n9k-switch# show run in clock</p> <pre>clock protocol none vdc 1 n9k-switch# n9k-switch# show clock 23:57:38.129 BST Wed Sep 20 2023 <<<<<< should have been UTC; BST comes from residual info left behind for summer-time config Time source is Hardware Calendar n9k-switch#</pre> <p>Workarounds: Configure clock timezone from the device: conf t clock timezone PST -8 0 end After having completed this configuration, you need to go back and check from shell that localtime file has been populated.</p> <pre>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 TZif2 PST TZif2 tT bash-4.4# strings /etc/localtime TZif2 TZif2 bash-4.4#</pre> <p>The "clock timezone" config may be removed if needed after verifying the file.</p> |
| CSCwh84282 | <p>Headline: After reload of 93108TC-FX3P random copper/RJ45 interfaces might not come up</p> <p>Symptoms: After reload of 93108TC-FX3P device some RJ45 interfaces might not come up. This can happen with any reload reason (power up/down, upgrade, SW reset, crash). Issue is related only to the front panel interfaces Ethernet 1/1-48 (MGMT port is not affected)</p> <p>Workarounds: FPGA upgrade</p> |
| CSCwh94363 | <p>Headline: BGP routes stuck in stale for shutdown BGP neighbor</p> <p>Symptoms: BGP neighbor is in Shut (Admin) State BGP table shows all routes in a Stale and Best state from the neighbor that was shutdown URIB shows route still learned via the BGP peer that is shutdown</p> <p>Workarounds: 1- bringing back up the BGP neighbor - Routes are refreshed and work2 - restart BGP process. DO NOT WORK - clear bgp ipv4 un * softclear ip route X.X.X.X/Y</p> |
| CSCwi28438 | <p>Headline: Rapid mac flaps caused by Server NIC teaming failover will cause ARP table and CAM table out of sync</p> <p>Symptoms: Server is connected to a pair of Nexus 9K Broadcom switches over Non-VPC HIF FEX ports. When a Nexus 9000 has 10 mac flaps under 1 second. Mac learning for that vlan will be disabled for 120 seconds. The egress interface entry for ARP and mac address table may point to different ports. When mac learning is re-enabled, the egress interface for the ARP entry may not update to the new interface. This may occur when a host with nic teaming transitions between active-standby to active-active and back.</p> <p>Workarounds: 1) Clear the ARP entry 2) Send an arp request on the previous port and then another one on the new port</p> |
| CSCwf89997 | <p>Headline: Unexpected reload ipfib hap reset</p> <p>Symptoms: Symptoms: During AM adjacency update, sometimes, the reprogramming of MPLS adjacencies goes in a loop and the system crashes.</p> <p>Workarounds: None</p> |

| Bug ID | Description |
|----------------------------|--|
| CSCwh00127 | <p>Headline: Kernel Panic due to HR Timeout & SMC</p> <p>Symptoms: Following kernel panic logs are seen on one of the N3K-C3548P-XL boxes. \$ %KERN-0-SYSTEM_MSG: [11269926.144849] NMI watchdog: BUG: soft lockup - CPU#2 stuck for 22s! [swapper/2:0] - kernel%KERN-0-SYSTEM_MSG: [11269926.144869] NMI watchdog: BUG: soft lockup - CPU#0 stuck for 22s! [swapper/0:0] - kernelVDC-1 %\$ %VPC-2-PEER_KEEP_ALIVE_RECV_FAIL: In domain 1, VPC peer keep-alive receive has failedVDC-1 %\$ %KERN-0-SYSTEM_MSG: [11269926.144849] NMI watchdog: BUG: soft lockup - CPU#2 stuck for 22s! [swapper/2:0] - kernelVDC-1 %\$ %KERN-0-SYSTEM_MSG: [11269926.144869] NMI watchdog: BUG: soft lockup - CPU#0 stuck for 22s! [swapper/0:0] - kernelVDC-1 %\$ %KERN-0-SYSTEM_MSG: [11269926.148346] nxos_panic: Kernel panic - not syncing: softlockup: hung tasks - kernelVDC-1 %\$ %KERN-0-SYSTEM_MSG: [11269926.148415] ttyS console device is disabled - kernelVDC-1 %\$ %KERN-0-SYSTEM_MSG: [11269926.150327] END: PANIC REPORT GENERATED AT 1689769096 - kernel</p> <p>Workarounds: None.</p> |
| CSCwf25135 | <p>Headline: MSDP peer is flapping</p> <p>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</p> <p>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</p> |
| CSCwh42690 | <p>Headline: Netstack crash when punting IPV6 packet with IPV6 destination option larger than 56 bytes</p> <p>Symptoms: In an MPLS environment, the device that punts the IPV6 MPLS packet (with an extension header of greater than 56 bytes) to SUP experiences netstack crash or device reload</p> <p>Workarounds: Stop sending packets that get punted to CPU such as traceroute probes that might have these options enabled.</p> |
| CSCwa96146 | <p>Headline: Cisco NX-OS Release 10.1(1) does not set TACACS+ Single-Connect Flag with CLI enabled.</p> <p>Symptoms: Customer is trying to implement TACACS+ "single-connect" mode, and it is not working correctly.</p> <p>Workarounds:</p> <ul style="list-style-type: none"> - Use Cisco TACACS+ servers (ISE) which have a check box for "enable single connect mode" and "Legacy devices". This will make the server function in single connect mode, and will not check for the flag in TACACS+ headers. - Disable Single-Connection mode. - Remove CLI from NX-OS |

Known Issues and Limitations

| Bug ID | Description |
|----------------------------|---|
| CSCwb90675 | <p>If a EOR gets stuck while configuring monitor scale at flow system configuration, try the following workaround:</p> <ul style="list-style-type: none"> - Save current running config to bootflash, other local file system or USB. - Reload ascii - Reload switch. - Do a SUP system switchover (SSO). |
| CSCwe74517 | <p>When you remove template peer-policy under I2vpn EVPN address-family it deletes the prefixes. The workaround is to reconfigure the affected neighbor or restart BGP.</p> |
| CSCwh49061 | <p>VLAN Mapping: strict incompatibility during downgrade from 10.3.x to 9.3.x. VLAN Mapping commands need to be removed when you downgrade from 10.3.x to 9.3.x versions. Workaround is to remove switchport vlan mapping command and add it again after the downgrade.</p> |
| CSCwi35215 | <p>If a single mac address moves more than 10 times in a second, then mac learning would be disabled on corresponding vlan for 120 seconds. This is a known limitation.</p> |
| CSCwi15483 | <p>Interface is up without cable, only QSFP is inserted: Interface will not go down when the link is shut on the remote side and even after removing the cable interface is still up. This is a known limitation.</p> |
| CSCwi99525 | <p>On Cisco Nexus N2K-C2348TQ HIFs fail to utilize redundant Port-Channel links, to NIF, during link failover events.</p> |

Device Hardware

The following tables list the Cisco Nexus 9000 Series hardware that Cisco NX-OS Release 9.3(13) 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-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 |
| 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 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 |

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 |

| Product ID | Description | Maximum Quantity | | |
|--------------|--|------------------|------------------|------------------|
| | | Cisco Nexus 9504 | Cisco Nexus 9508 | Cisco Nexus 9516 |
| 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 card | 4 | 8 | 16 |
| 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 |
| 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-FAN2 | Fan 2 module with port-side intake airflow (burgundy coloring) | 3 | 93128TX |
| N9K-C9300-FAN2-B | Fan 2 module with port-side exhaust airflow (blue coloring) | 3 | 93128TX |
| 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 |
| 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 9348GC-FXP ^a |
| NXA-FAN-30CFM-F | Fan module with port-side exhaust airflow (blue coloring) | 3 | 92160YC-X 9236C ^a |

| Product ID | Description | Quantity | Cisco Nexus Switches |
|------------------|--|----------|---|
| | | | 93108TC-EX 93108TC-FX ^a 93180LC-EX ^a 93180YC-EX 93180YC-FX ^a 9332PQ 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 |

^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 |
| N9K-PAC-650W-B | 650-W AC power supply with port-side exhaust (blue coloring) | 2 | 9332PQ |
| NXA-PAC-650W-PE | 650-W power supply with port-side exhaust (blue coloring) | 2 | 92160YC-X 9236C 92300YC 93180YC-FX3S 92304QC 93108TC-EX |

| Product ID | Description | Quantity | Cisco Nexus Switches |
|-------------------|---|----------|--|
| | | | 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 (burgundy coloring) | 2 | 93240YC-FX2 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 |

| Product ID | Description | Quantity | Cisco Nexus Switches |
|------------------|--|----------|---|
| | | | 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 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 |
| 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 |

| Product ID | Description | Quantity | Cisco Nexus Switches |
|------------------|--|----------|--------------------------|
| 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 |
| 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

| Product ID | 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-C93108TC-EX | 1-RU Top-of-Rack switch with 48 10GBASE-T (copper) ports and 6 40-/100-Gigabit QSFP28 ports |

| Product ID | Description |
|--------------------|--|
| 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+ 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) |
| N9K-C93180YC-FX3S | 6 10/25/40/50/100-Gigabit QSFP28 ports (ports 49-54) |
| N9K-C93216TC-FX2 | 48 1/10/25 Gigabit Ethernet SFP28 ports (ports 1-48) |
| N9K-C93240YC-FX2 | 6 10/25/40/50/100-Gigabit QSFP28 ports (ports 49-54) |
| 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 |

| Product ID | Description |
|-----------------|--|
| 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 <ul style="list-style-type: none"> - 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-C9364C-GX | 2-RU fixed-port switch with 64 100-Gigabit SFP28 ports |
| N9K-9372PX-E | 1RU switch with 48 x 1/10-Gbps SFP+ and 6 x 40-Gbps fixed QSFP+ ports |
| N9K-9372TX-E | 1RU switch with 48 x 100M/1/10GBASE-T and 6 x 40-Gbps fixed QSFP+ 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

Cisco NX-OS Release 9.3(13) 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 sub interface 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
- gPTP

Note: One-step PTP is supported only on Cisco Nexus 9500-R series.

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\(11\)](#)

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, 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.