



Cisco Nexus 9000 Series NX-OS Release Notes, Release 6.1(2)I3(5)

This document describes the features, caveats, and limitations for Cisco NX-OS Release 6.1(2)I3(5) software for use on the Cisco Nexus 9000 Series switches and the Cisco Nexus 3164Q switch. Use this document in combination with documents listed in *Related Documentation*.

Table 1 shows the online change history for this document.

Table 1. Online History Change

Date	Description
August 21, 2015	Created the release notes for Release 6.1(2)I3(5).
January 11, 2016	Added a link to the ALE limitations in the <i>Limitations</i> section.

Contents

This document includes the following sections:

Introduction	2
System Requirements	2
New and Changed Information	6
Installation Notes	7
Upgrade Instructions	7
Downgrade Instructions	7
Software Maintenance Upgrades	7
Limitations	7
Guidelines and Limitations for Private VLANs	9
Unsupported Features	13
Caveats	14
Related Documentation	18
Obtaining Documentation and Submitting a Service Request	18

Introduction

Cisco NX-OS software is a data center-class operating system designed for performance, resiliency, scalability, manageability, and programmability at its foundation. The Cisco NX-OS software provides a robust and comprehensive feature set that meets the requirements of virtualization and automation in mission-critical data center environments. The modular design of the Cisco NX-OS operating system makes zero-impact operations a reality and enables exceptional operational flexibility.

The Cisco Nexus 9000 Series uses an enhanced version of Cisco NX-OS software with a single binary image that supports every switch in the series, which simplifies image management.

System Requirements

This section includes the following sections.

- Supported Device Hardware
- Supported Optics
- Supported FEX Modules

Supported Device Hardware

Table 2 lists the Cisco Nexus 9000 Series hardware that Cisco NX-OS Release 6.1(2)I3(5) supports. For additional information about the supported hardware, see the Hardware Installation Guide for your Cisco Nexus 9000 Series device.

Table 2. Cisco Nexus 9000 Series Hardware.

Product ID	Hardware	Quantity
N9K-C9516	Cisco Nexus 9516 16-slot chassis	1
N9K-C9516-FM	Cisco Nexus 9500 Series fabric module	3-6 depending on the line card
N9K-C9516-FAN	Cisco Nexus 9516 fan trays	3
N9K-C9508	Cisco Nexus 9508 8-slot chassis	1
N9K-C9508-FM	Cisco Nexus 9508 Series fabric module	3-6 depending on the line card
N9K-C9508-FAN	Cisco Nexus 9508 fan trays	3

System Requirements

N9K-X9564PX	Cisco Nexus 9500 Series 48-port, 1-/10-Gbps SFP+ plus 4-port QSFP I/O module	<ul style="list-style-type: none"> ■ Up to 4 in the Cisco Nexus 9504 ■ Up to 8 in the Cisco Nexus 9508 ■ Up to 16 in the Cisco Nexus 9516
N9K-X9564TX	Cisco Nexus 9500 Series 48-port, 1-/10-Gbps BASE-T plus 4-port QSFP I/O module	<ul style="list-style-type: none"> ■ Up to 8 in the Cisco Nexus 9508 ■ Up to 4 in the Cisco Nexus 9504 ■ Up to 16 in the Cisco Nexus 9516
N9K-X9536PQ	Cisco Nexus 9500 36-port, 40 Gigabit Ethernet QSFP aggregation module	<ul style="list-style-type: none"> ■ Up to 4 in the Cisco Nexus 9504 ■ Up to 8 in the Cisco Nexus 9508 ■ Up to 16 in the Cisco Nexus 9516
N9K-X9636PQ	Cisco Nexus 9500 Series 36-port 40-Gigabit QSFP I/O module <i>Note:</i> Not supported on the Cisco Nexus 9516 switch (N9K-C9516).	<ul style="list-style-type: none"> ■ Up to 4 in the Cisco Nexus 9504 ■ Up to 8 in the Cisco Nexus 9508
N9K-X9464PX	Cisco Nexus 9500 Series 48-port 10-Gigabit SFP+ plus 4-port QSFP I/O module	<ul style="list-style-type: none"> ■ Up to 4 in the Cisco Nexus 9504 ■ Up to 8 in the Cisco Nexus 9508 ■ Up to 16 in the Cisco Nexus 9516
N9K-X9464TX	Cisco Nexus 9500 Series 48-port 10-GBASE-T plus 4-port QSFP I/O module	<ul style="list-style-type: none"> ■ Up to 4 in the Cisco Nexus 9504 ■ Up to 8 in the Cisco Nexus 9508 ■ Up to 16 in the Cisco Nexus 9516

System Requirements

N9K-X9432PQ	Cisco Nexus 9500 Series 32-port 40-Gigabit QSFP I/O module <i>Note:</i> The Cisco Nexus X9432PQ I/O module supports static breakout.	<ul style="list-style-type: none"> ■ Up to 4 in the Cisco Nexus 9504 ■ Up to 8 in the Cisco Nexus 9508 ■ Up to 16 in the Cisco Nexus 9516
N9K-SC-A	Cisco Nexus 9500 Series System Controller Module	2
N9K-SUP-A	Cisco Nexus 9500 Series supervisor module	2
N9K-SUP-B	Cisco Nexus 9500 Series supervisor B module	2
N9K-PAC-3000W-B	Cisco Nexus 9500 Series 3000 W AC power supply	Up to 4 in the Cisco Nexus 9504 Up to 8 in the Cisco Nexus 9508 Up to 10 in the Cisco Nexus 9516
N9K-C9504	Cisco Nexus 9504 4-slot chassis	1
N9K-C9504-FM	Cisco Nexus 9504 fabric module	3 to 6 depending on line card
N9K-C9504-FAN	Cisco Nexus 9504 fan trays	3
N9K-C9396PX	Cisco Nexus 9300 48-port, 1-/10-Gigabit Ethernet SFP+ and 12-port, 40-Gigabit Ethernet QSFP switch	1
N9K-C9396TX	Cisco Nexus 9300 48-port, 1-/10-Gigabit Ethernet BASE-T and 12-port, 40-Gigabit Ethernet QSFP switch	1
N9K-C9372PX	Cisco Nexus 9300 48-port, 1-/10-Gigabit Ethernet SFP+ and 6-port, 40-Gigabit Ethernet QSFP switch	1
N9K-C9372TX	Cisco Nexus 9300 48-port, 1-/10-Gigabit Ethernet BASE-T and 6-port, 40-Gigabit Ethernet QSFP switch	1
N9K-C9332PQ	Cisco Nexus 9300 32-port, 40-Gigabit Ethernet QSFP switch	1
N9K-C93128TX	Cisco Nexus 9300 switch with 96 1-/10-Gigabit BASE-T ports and eight 40-Gigabit Ethernet QSFP ports (The 1-/10-Gigabit BASE-T ports also support a speed of 100 Megabits.)	1
N9K-PAC-650W	Cisco Nexus 9300 650 W AC power supply, hot air out (red) <i>Note:</i> For use with the Cisco Nexus 9396 switch (N9K-C9396PX).	2 or less

System Requirements

N9K-PAC-650W-B	Cisco Nexus 9300 650 W AC power supply, cold air in (blue) Note: For use with the Cisco Nexus 9396 switch (N9K-C9396PX).	2 or less
N9K-PAC-1200W	Cisco Nexus 9300 1200 W AC power supply, hot air out (red) Note: For use with the Cisco Nexus 93128 switch (N9K-C93128TX).	2 or less
N9K-PAC-1200W-B	Cisco Nexus 9300 1200 W AC power supply, cold air in (blue) Note: For use with the Cisco Nexus 93128 switch (N9K-C93128TX).	2 or less
N9K-C9300-FAN1	Cisco Nexus 9300 fan 1, hot air out (red) Note: For use with the Cisco Nexus 9396 switch (N9K-C9396PX).	3
N9K-C9300-FAN1-B	Cisco Nexus 9300 fan 1, cold air in (blue) Note: For use with the Cisco Nexus 9396 switch (N9K-C9396PX).	3
N9K-C9300-FAN2	Cisco Nexus 9300 fan 2, hot air out (red) Note: For use with the Cisco Nexus 93128 switch (N9K-C93128TX).	3
N9K-C9300-FAN2-B	Cisco Nexus 9300 fan 2, cold air in (blue) Note: For use with the Cisco Nexus 93128 switch (N9K-C93128TX).	3
NXA-FAN-30CFM-F	Cisco Nexus 9300 fan, port-side exhaust Note: For use with the Cisco Nexus 9332PQ, 9372PX, and 9372TX switches (N9K-C9332PQ, N9K-C9372PX, and N9K-9372TX).	4
NXA-FAN-30CFM-B	Cisco Nexus 9300 fan, port-side intake Note: For use with the Cisco Nexus 9332PQ, 9372PX, and 9372TX switches (N9K-C9332PQ, N9K-C9372PX, and N9K-9372TX).	4
N9K-M12PQ	Cisco Nexus GEM 9300 uplink module, 12-port, 40-Gigabit Ethernet QSPF Note: The front-panel ports on these GEM modules do not support auto negotiation with copper cables. Manually configure the speed on the peer switch.	1 (required)
N9K-M6PQ	Cisco Nexus GEM 6-port 40-Gigabit Ethernet uplink module for the Cisco Nexus 9396PX, 9396TX, and 93128TX switches Note: The front-panel ports on these GEM modules do not support auto negotiation with copper cables. Manually configure the speed on the peer switch.	1

Table 3 lists the Cisco Nexus 3164Q switch hardware that Cisco NX-OS Release 6.1(2)I3(5) supports.

Table 3. Cisco Nexus 3164Q Switch Hardware.

Product ID	Hardware	Quantity
N3K-C3164Q-40GE	Cisco Nexus 3164Q switch	1
N9K-C9300-FAN3	Cisco Nexus 3164Q fan module	3
N9K-PAC-1200W	Cisco Nexus 3164Q 1200W AC power supply	2

For additional information about the supported hardware, see the *Cisco Nexus 3000 Series Hardware Installation Guide*.

Supported Optics

See the [Cisco 10-Gigabit Ethernet Transceiver Modules Compatibility Matrix](#) for a list of supported optical components.

Supported FEX Modules

The following is a list of FEX modules the Cisco NX-OS Release 6.1(2)I3(5) supports with Cisco Nexus 9372PX and 9396PX switches:

- Cisco Nexus 2224TP
- Cisco Nexus 2232PP
- Cisco Nexus 2232TM and 2232TM-E
- Cisco Nexus 2248PQ
- Cisco Nexus 2248TP and 2248TP-E
- Cisco Nexus B22Dell
- Cisco Nexus B22HP

New and Changed Information

This section lists the following topics:

- New Hardware Features in Cisco NX-OS Release 6.1(2)I3(4b)
- New Software Features in Cisco NX-OS Release 6.1(2)I3(4b)

New Hardware Features in Cisco NX-OS Release 6.1(2)I3(5)

Cisco NX-OS Release 6.1(2)I3(5) supports no new hardware features.

New Software Features in Cisco NX-OS Release 6.1(2)I3(5)

Cisco NX-OS Release 6.1(2)I3(5) supports no new software features.

Installation Notes

Only one software image (called nx-os) is required to load the Cisco NX-OS operating system. This image runs on all Cisco Nexus 9000 Series switches and the Cisco Nexus 3164Q switch. For installation instructions, see the *Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide*.

Upgrade Instructions

To perform a software upgrade, follow the installation instructions in the Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide.

Note: Cisco NX-OS Release 6.1(2)I3(5) only supports a fast reload upgrade to 7.0(3)I1(2) and later releases.

Note: Upgrading with PVLANS configured is only supported with 7.0(3)I1(2) and later releases.

Note: For a boot-variable change and reload to a 7.0(3)I1(1x) release, the PVLAN process is not brought up, and the PVLAN ports are kept down.

Downgrade Instructions

To perform a software downgrade to an earlier release, disable the Guest Shell.

Note: Downgrading with PVLANS configured is only supported with the 6.1(2)I3(4) and 6.1(2)I3(4a) releases.

Note: For a boot-variable change to the 6.1(2)I3(3) release and earlier, an ASCII replay will be tried, but feature PVLANS and other PVLAN configurations will fail.

Software Maintenance Upgrades

For information about software maintenance upgrades, see the “Performing Software Maintenance Upgrades” section in the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#).

Note: If you perform a software maintenance upgrade (SMU) and later upgrade your device to a new Cisco NX-OS software release, the new image will overwrite both the previous Cisco NX-OS release and the SMU package file.

Limitations

This section lists limitations related to Cisco NX-OS Release 6.1(2)I3(5).

- Layer 3 routed traffic for missing Layer 2 adjacency information is not flooded back onto VLAN members of ingress units when the source MAC address of routed traffic is a non-VDC MAC address. This limitation is for hardware flood traffic and can occur when the SVI has a user-configured MAC address.
- The neighbor-down fib-accelerate command is supported in a BGP-only environment.

Limitations

- The uplink module should not be removed from a Cisco Nexus 9300 Series switch that is running Cisco NX-OS Release 6.1(2)I3(5). The ports on the uplink module should be used only for uplinks.
- The PortLoopback and BootupPortLoopback tests are not supported.
- The ASIC Memory-NS test is applicable only for the N9K-X9564PX and N9K-X9564TX line cards.
- Priority flow control (PFC) is supported on Cisco Nexus 9500 Series switches with the N9K-X9636PQ line card. It is not yet supported on Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with the N9K-X9564PX or N9K-X9564TX line card.
- If you configure the FEX with 100/full-duplex speed and you do not explicitly configure the neighboring device with 100/full-duplex speed, the data packet traffic does not pass properly even though the link appears to be “up.”
 - no speed: Auto negotiates and advertises all speeds (only full duplex).
 - speed 100: Does not auto negotiate; pause cannot be advertised. The peer must be set to not auto negotiate (only 100 Mbps full duplex is supported).
 - speed 1000: Auto negotiates: Auto negotiates and advertises pause (advertises only for 1000 Mbps full duplex).
- Eight QoS groups are supported only on modular platforms with the following Cisco Nexus 9500 Series line cards: N9K-X9636PQ, N9K-X9464PX, N9K-X9464TX, and N9K-X9432PQ.
- Cisco NX-OS Release 6.1(2)I2(2b) supports flooding for Microsoft Network Load Balancing (NLB) unicast mode on Cisco Nexus 9500 Series switches but not on Cisco Nexus 9300 Series switches. NLB is not supported in max-host system routing mode. NLB multicast mode is not supported on Cisco Nexus 9500 or 9300 Series switches.

Note: To work around the situation of Unicast NLB limitation, we can statically hard code the ARP and MAC address pointing to the correct interface. Please refer to bug ID CSCuq03168 in detail in the “**Open Caveats—Cisco NX-OS Release 6.1(2)I3(5)**” **section**.

- The no negotiate auto command, which disables auto-negotiation, is only applicable for the following port types:
 - Base-T ports (9464TX,9564TX LCs. 9396TX, 93128TX, 9372TX TORs) with speed 100 configuration
 - 1G Fiber SFP Ports (9464PX, 9564PX LCs. 9396PX, 9372PX TORs with 1G Fiber SFP inserted) with speed 1000 configuration
 - 40G QSFP Ports with copper QSFP (9436PQ, 9536PQ, 9636PQ LCs. 3164Q, 9332PQ) first 26 ports

Note: For 40G QSFP ports, auto-negotiation is disabled by configuring speed 40000 in the 6.1(2)I3(3a) release and earlier. To continue to support backward compatibility, if speed 40000 is configured, no negotiate auto will be added internally by the software.

- DHCP relay to subnet broadcast address is not supported.
- An ACL or source VLAN on a SPAN session affects traffic on other sessions.
- Traffic is not spanned if SPAN ACL and policy-based routing are enabled.
- When routed ACL is applied to multiple SVIs (switched virtual interfaces) in the egress direction, TCAM resources are not shared.

Guidelines and Limitations for Private VLANs

- When VACL (VLAN ACL) is applied to multiple VLANs, TCAM resources are not shared.
- N9K hardware does not support range checks (layer 4 operators) in egress TCAM. Because of this, ACL/QoS policies with layer 4 operations-based classification need to be expanded to multiple entries in the egress TCAM. Egress TCAM space planning should take this limitation into account.
- If the same QoS policy and ACL is applied on multiple interfaces, the label will be shared only when the qos-policy is applied with the no-stats option.
- The following switches support the QSFP+ with QSA adapter (40G to 10G QSA): N9K-C9396PX, N9K-C9396TX, N9K-C93128TX, N9K-C9372PX, N9K-C9372TX, and N9K-C9332PQ

Note: The Nexus 9300 support for the QSFP+ with QSA adapter has the following limitations:

- Only 10G can be supported using QSA on 40G uplink ports on N9300 switches in NX-OS.
- 1G with QSA is not supported.
- For the Cisco Nexus 9332PQ switch, all ports except 13-14 and 27-32 can support breakout
- All ports in the QSA speed group must operate at the same speed (see the configuration guide)
- The following switches support the breakout cable (40G ports to 4x10G ports): N9K-C9332PQ, N9K-X9436PQ, and N9K-X9536PQ.
- Limitations for ALE uplink ports are listed at the following URL:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/ale_ports/b_Limitations_for_ALE_Uplink_Ports_on_Cisco_Nexus_9000_Series_Switches.html

Guidelines and Limitations for Private VLANs

This section provides guidelines and limitations for configuring private VLANs.

- Configuring Private VLANs
- Secondary and Primary VLAN Configuration
- Private VLAN Port Configuration
- Limitations with Other Features

Configuring Private VLANs

Private VLANs have the following configuration guidelines and limitations:

- You must enable private VLANs before the device can apply the private VLAN functionality.
- You must enable the VLAN interface feature before the device can apply this functionality.
- Shut down the VLAN network interface for all VLANs that you plan to configure as secondary VLANs before you configure these VLANs.
- When a static MAC is created on a regular VLAN and then that VLAN is converted to a secondary VLAN, the Cisco NX-OS maintains the MAC that was configured on the secondary VLAN as the static MAC.

- Private VLANs support PVLAN port modes as follows:
 - Promiscuous.
 - Promiscuous trunk.
 - Isolated host.
 - Isolated host trunk.
 - Community host.
- When configuring PVLAN promiscuous trunks or PVLAN isolated trunks, it is recommended to allow non-private VLANs in the list specified by the switchport private-vlan trunk allowed id command.
- Private VLANs are mapped or associated depending on the PVLAN trunk mode.
- Private VLANs support PVLAN and PACLs/RACLs.
- Private VLANs support PVLAN and SVIs as follows:
 - SVI allowed only on primary VLANs.
 - Primary and secondary IP's on the SVI.
 - HSRP on the primary SVI.
- Private VLANs support PVLAN and Layer 2 forwarding.
- Private VLANs support PVLAN and STP as follows:
 - RSTPs
 - MSTs
- Private VLANs support PVLAN across switches as follows:
 - Through a regular trunk port.
 - Through a regular trunk port-channel.
 - Through a regular vPC-port.
- Private VLANs do not provide support for Cisco Fabric Extenders (FEXs).
- Private VLANs port mode is not supported on 40G interfaces of the Cisco Nexus C9396PX or Cisco Nexus C93128TX.
- Private VLANs port mode is not supported on the Nexus 3164Q.
- Private VLANs do not provide port mode support for port channels.
- Private VLANs do not provide port mode support for virtual port channels (vPCs) interfaces.
- Private VLANs do not provide support on breakout.
- Private VLANs do not provide support for IP multicast or IGMP snooping.
- Private VLANs do not provide support for DHCP snooping.

Guidelines and Limitations for Private VLANs

- Private VLANs do not provide support for PVLAN QoS.
- Private VLANs do not provide support for VACLs.
- Private VLANs do not provide support for VTP.
- Private VLANs do not provide support for tunnels.
- Private VLANs do not provide support for VXLANs.
- Private VLANs do not provide support for SPAN when the source is a PVLAN VLAN.
- You cannot configure a shared interface to be part of a private VLAN. For more details, see the *Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide*.
- Although Cisco NX-OS CLI allows the configuration of multiple isolated VLAN configurations per PVLAN group, such a configuration is not supported. A PVLAN group can have at most one isolated VLAN.

Secondary and Primary VLAN Configuration

Follow these guidelines when configuring secondary or primary VLANs in private VLANs:

- You cannot configure the default VLAN (VLAN1) or any of the internally allocated VLANs as primary or secondary VLANs.
- You must use VLAN configuration (config-vlan) mode to configure private VLANs.
- A primary VLAN can have multiple isolated and community VLANs associated with it. An isolated or community VLAN can be associated with only one primary VLAN.
- Although private VLANs provide host isolation at Layer 2, hosts can communicate with each other at Layer 3.
- A PVLAN group can have at most one isolated VLAN. Multiple isolated VLAN configurations per primary VLAN configurations are not supported.
- When a secondary VLAN is associated with the primary VLAN, the STP parameters of the primary VLAN, such as bridge priorities, are propagated to the secondary VLAN. However, STP parameters do not necessarily propagate to other devices. You should manually check the STP configuration to ensure that the spanning tree topologies for the primary, isolated, and community VLANs match exactly so that the VLANs can properly share the same forwarding database.
- For normal trunk ports, note the following:
 - There is a separate instance of STP for each VLAN in the private VLAN.
 - STP parameters for the primary and all secondary VLANs must match.
 - The primary and all associated secondary VLANs should be in the same MST instance.
- For nontrunking ports, note the following:
 - STP is aware only of the primary VLAN for any private VLAN host port; STP runs only on the primary VLAN for all private VLAN ports.

Note: We recommend that you enable BPDU Guard on all ports that you configure as a host port; do not enable this feature on promiscuous ports.

- For private VLAN promiscuous trunk ports, note the following:
 - You can configure a maximum of 16 private VLAN primary and secondary VLAN pairs on each promiscuous trunk port.
- For private VLAN isolated trunk ports, note the following:
 - You can configure a maximum of 16 private VLAN primary and secondary VLAN pairs on each isolated trunk port.
 - The native VLAN must be either a normal VLAN or a private VLAN secondary VLAN. You cannot configure a private VLAN primary port as the native VLAN for a private VLAN isolated trunk port.
- To downgrade a system that has private VLAN ports configured, you must unconfigure these ports.
- Before you configure a VLAN as a secondary VLAN, you must shut down the VLAN network interface for the secondary VLAN.
- DHCP relay on secondary VLANs will not work if there are no ports in the primary VLAN on the line card ASIC, **which has the DHCP relay packet's** incoming interface.

Private VLAN Port Configuration

Follow these guidelines when configuring private VLAN ports:

- Use only the private VLAN configuration commands to assign ports to primary, isolated, or community VLANs.
- The Layer 2 access ports that are assigned to the VLANs that you configure as primary, isolated, or community VLANs are inactive while the VLAN is part of the private VLAN configuration. Layer 2 trunk interfaces, which may carry private VLANs, are active and remain part of the STP database.
- If you delete a VLAN used in the private VLAN configuration, the private VLAN ports (promiscuous ports or host ports, not trunk ports) that are associated with the VLAN become inactive.

Limitations with Other Features

Consider these configuration limitations with other features when configuring private VLANs:

Note: In some cases, the configuration is accepted with no error messages, but the commands have no effect.

- User has to ensure consistent pvlan type, states and configuration across vpc peers. There is no pvlan consistency check for vpc currently. Inconsistent pvlan configs across vpc peers may end up in incorrect forwarding and impacts.
- A private VLAN port can be configured as a SPAN source port.
- Private VLAN host or promiscuous ports cannot be a SPAN destination port.
- A destination SPAN port cannot be an isolated port. (However, a source SPAN port can be an isolated port.)
- After you configure the association between the primary and secondary VLANs, the dynamic MAC addresses that learned the secondary VLANs are aged out.
- After you configure the association between the primary and secondary VLANs, you cannot create static

Unsupported Features

- MAC addresses for the secondary VLANs.
- After you configure the association between the primary and secondary VLANs, if you delete the association, all static MAC addresses that were created on the primary VLANs remain on the primary VLAN only.
- In private VLANs, STP controls only the primary VLAN.

Note: See the *Cisco Nexus 9000 Series NX-OS Security Configuration Guide* for information on configuring static MAC addresses.

Unsupported Features

This section lists features that are not supported in the current release:

- [VXLAN](#)
- [PVLAN](#)
- [DHCP](#)
- [FEX](#)
- [Other Unsupported Features](#)

VXLAN

This section lists VXLAN features that are not supported.

- VXLAN routing is not supported.
- The default Layer 3 gateway for VXLAN VLANs should be provisioned on a different device.
- Switched Port Analyzer (SPAN) Tx for VXLAN traffic is not supported for the access-to-network direction.
- Ingress router access control lists (ACLs) are not supported on Layer 3 uplinks for VXLAN traffic. Egress VACLs cannot be used on decapsulated packets in the network-to-access direction on the inner payload. As a best practice, use PACLS or VACLs for the access-to-network direction.
- QoS classification is not supported for VXLAN traffic in the network-to-access direction.
- The QoS buffer-boost feature is not applicable for VXLAN traffic.
- ACL and QoS for VXLAN traffic in the network-to-access direction is not supported.
- Native VLANs for VXLAN are not supported. All traffic on VXLAN Layer 2 trunks needs to be tagged.
- Consistency checkers are not supported for VXLAN tables.
- Just one network virtualization edge (NVE) interface is allowed on the switch.
- Because the NVE (VXLAN) process is not restartable, patching support is not supported for VXLAN.
- vPC type-1 consistency checkers are not supported for VXLAN configurations.
- Dynamic re-IP of an NVE tunnel is not supported. Tunnels must be shut down prior to live IP address changes.

Caveats

VXLAN Topology Restrictions

- A device cannot be a VXLAN gateway (vxlan-vlan flows) and a VXLAN bridge (vxlan-vxlan flows) for the same multicast groups, which are also called the bud-node topology. As a best practice, use the device as either a bridging device or a gateway device, but not both.
- Due to bud-node restrictions, a VXLAN tunnel endpoint (VTEP) cannot reach the rendezvous point (RP) through another VTEP. Because of this limitation, there can be no direct Layer 3 links between two VTEPs, unless one of the VTEPs is the RP.
- FEX host interface ports are not supported for VLANs extended with VXLAN.

VXLAN ACL Limitations

The following ACL related features are not supported:

- Ingress RACL that is applied on an uplink Layer 3 interface that matches on the inner or outer payload in the network-to-access direction (decapsulated path)
- Egress RACL that is applied on an uplink Layer 3 interface that matches on the inner or outer payload in the access-to-network direction (encapsulated path)
- Egress VACL for decapsulated VXLAN traffic

We recommend that you use a PACL or VACL on the access side to filter out traffic entering the overlay network.

PVLAN

PVLAN port modes are not supported on Port-channels, vPCs, and Nexus 2000 interfaces.

DHCP

DHCP subnet broadcast is not supported.

FEX

- FEX is supported only on the Cisco Nexus 9372PX and 9396PX switches. It is not supported on the other Cisco Nexus 9300 Series switches or the Cisco Nexus 9500 Series switches.

Other Unsupported Features

The following lists other features not supported in the current release:

- Due to a Poodle vulnerability, SSLv3 is no longer supported.
- The Cisco Nexus 9300 Series switches and the Cisco Nexus 3164Q switch do not support the 64-bit ALPM routing mode.

Caveats

This section contains lists of open and resolved caveats.

Caveats

- [Resolved Caveats—Cisco NX-OS Release 6.1\(2\)I3\(5\)](#)
- [Open Caveats—Cisco NX-OS Release 6.1\(2\)I3\(5\)](#)
- Known Behaviors—Cisco NX-OS Release 6.1(2)I3(5)

Resolved Caveats—Cisco NX-OS Release 6.1(2)I3(5)

Table 4 lists the Resolved Caveats in the Cisco NX-OS Release 6.1(2)I3(5) release. Click the bug ID to access the Bug Search tool and see additional information about the bug.

Table 4 Resolved Caveats in Cisco NX-OS Release 6.1(2)I3(5)

Bug ID	Description
CSCuo49979	The current Delay Restore/Auto-recover values should be visible.
CSCup55648	A buildup of MTS messages for SAP 199 may lead to an IPFIB crash on multiple line cards.
CSCuq60026	The Cisco Fabric Services (CFS) leaks 12KB with the restart pim command.
CSCus71452	Certain IPs are unreachable when sending traffic through a Nexus 9300.
CSCut08582	Snmpd crashed when a user issued an snmpgetbulk request with multiple OIDs.
CSCut48634	When executing a copy file start +fastreload/reload with a customer copp policy, the following error message is displayed: "Non atomic ACL/QOS policy update done for Copp". This has no functional impact and is just an error message.
CSCut68375	When using a Nexus 9508 and vPC peer-links created on different modules, the STP status of the vPCs change to BLK and then FWD when one of the modules is shut down.
CSCut81254	A lot of LPM routes are not programmed on Fabric Manager.
CSCuu91639	10G active optical cable should not be displayed as ACU with Cisco QSAs.
CSCuv05263	VTP advertisement within a vPC domain is causing high CPU usage. The VTP packet storm could hit adjacent devices causing instability downstream or upstream.
CSCuv06108	The following message appears in the syslog file: “%IPFIB-SLOT1-4-UFIB_ROUTE_CREATE: Unicast route create failed for VRF.”
CSCuv10431	ifTable and ifXTable counter objects toggle between 0 and actual value between successive polling.
CSCuv46792	The N9K acts as a DHCP relay agent with a link in the Spanning-tree Alternate Blocking State. This causes: <ul style="list-style-type: none"> ■ High CPU usage in dhcp_snoop and pktmgr ■ Excessive drops in the DHCP CoPP class ■ Control plane instability
CSCuv47053	Nexus 9000 may crash and generate a bcm_usd core. The reset- reason shows “Reset Requested due to Fatal Module Error.”

Open Caveats—Cisco NX-OS Release 6.1(2)I3(5)

Table 5 lists the open caveats in the Cisco NX-OS Release 6.1(2)I3(5) release. Click the bug ID to access the Bug Search tool and see additional information about the bug.

Table 5 Open Caveats in Cisco NX-OS Release 6.1(2)I3(5)

Bug ID	Description
CSCui95544	PVLAN as a source VLAN is not supported on span/ERspan sessions.
CSCun26726	HSRP packet decoding fails with an assertion error.
CSCun34856	All VLANs are suspended if one has a QoS policy but the TCAM is not configured.
CSCuo86036	Spurious error messages appear when an MCT port-channel member is shut.
CSCup35239	The Ethalyzer does not see packets that egress on a virtual interface.
CSCup55774	No drop counters are displayed when a FEX HIF is congested.
CSCup60475	Under rare circumstances, when a VPC node reloads and comes back online, some MACs may be missing in the hardware of the reloaded switch and cause flooding for traffic destined to those MACs.
CSCuq03168	Microsoft NLB traffic being routed into the destination VLAN is experiencing packet loss.
CSCuq36330	BFD support on Cisco ALE port sub interfaces.
CSCuq86978	VXLAN traffic will experience a decap failure on a standalone device. The traffic does not recover at all. The mcast tree/routes and unicast seem to fine.
CSCur22618	The show queuing interface command returns empty output when executed for FEX HIF interfaces.
CSCur30555	The show policy-map type queuing does not show statistics for FEX HIF interfaces.
CSCur63227	BGP prefixes can experience temporary traffic drop during supervisor switchover when BGP prefixes have the Nexthop learned over BGP (Recursive Nexthop) in the presence of a default route in the system.
CSCur37816	When QoS Lite TCAM is configured, policer violated statistics shown as part of the show policy-map interface command are reported as 0 instead of NA (Not-Applicable).
CSCur46879	When copying the tunnel configuration file to running, the tunnel may flap before stabilizing.
CSCur59482	Policer action is not supported when QoS policy type qos is applied with no-stats keyword.
CSCur61647	Even though there are no QoS classification policies currently active on any of the FEX HIF interfaces, the show incompatibility command still reports FEX QoS incompatibility during downgrade from 3.2 to earlier versions of software.
CSCur87839	Traffic cannot be routed using policy-based routing if the next-hop reachability is across the vPC peer link and the local vPC leg is down.
CSCus06693	ERPSAN sessions with destination on port-channel sub-interface is not supported.

Caveats

Bug ID	Description
CSCus64028	A PVLAN deleted on one vPC peer and present on the other, will see macs pointing to the vpc-peer-link on that suspended VLAN at the other peer.
CSCus92589	VLAN shut/no shut causes stale macs and sometimes traffic blackhole.
CSCus92777	The linkUp/linkDown/cielinkUp/cielinkDown traps are not sent Po subinterfaces.
CSCus94907	The cpvlanVlanTable does not show normal VLAN.
CSCus99951	The show hardware capacity for IPv6 does not show 100% utilization when a table is full.
CSCut10188	The dot1Q tag native functionality does not kick in on CLI configuration. Need interface flap to kick in.
CSCut10588	The no negotiate auto configuration requires speed configuration but does not print the warning clearly.
CSCut15002	Removing a private-vlan association after a PVLAN deletion will not go through.
CSCut24756	VLAN qos policy will not take affect if there is only a PVLAN Isolated trunk or Promiscuous trunk port config on a given T2 instance of the module.
CSCut27696	When a regular VLAN is converted to a secondary VLAN, the MAC learned on the secondary VLAN is re-learned on the primary VLAN, and the MAC learned before the conversion ages out.
CSCut32424	Traceroute should give a destination IP for the last hop. Currently it is giving an incoming interface IP.
CSCut36860	The speed 1000 configuration works same way as speed auto 100 1000 by advertising both 100M and 1G.
CSCuu81903	By default, Nexus 9000 switches support the ALPM mode of routing on the fabric modules, and switches with a " spine-leaf" architecture support the ALPM mode on the leaf units of the ASIC. Also, ALPM routing mode can be turned on by configuration. An ALPM sanity checking tool can be used to identity programming errors for ALPM routes. This is a development bug to track ALPM sanity checking tool.

Known Behaviors–Cisco NX-OS Release 6.1(2)I3(5)

Table 6 lists the known behaviors in the Cisco NX-OS Release 6.1(2)I3(5) release. Click the bug ID to access the Bug Search tool and see additional information about the bug.

Table 6 Known Behaviors in Cisco NX-OS Release 6.1(2)I3(5)

Bug ID	Description
CSCu118670	The show license usage command shows the incorrect license if a valid license is installed over an honor license.

Bug ID	Description
CSCus91619	Private-vlan association command on a Primary VLAN always does a union of the already available VLANs. Need to explicitly do no private-vlan association <vlan-id> or private-vlan association remove <vlan-id> if the association needs to be removed.
CSCut03046	VLAN Multicast counters (displayed by show vlan [id <vlan number>] counters) do not work for VLANs with IGMP snooping disabled on Modulat Platforms (N950X).
CSCut12867	When downgrading to a lower release of the 6.1(2) train by changing the boot-variable, the PVLAN configurations will be retained and unsupported behavior may occur. Install all will block downgrade if PVLAN is enabled.

Related Documentation

The entire Cisco Nexus 9000 Series NX-OS documentation set is available at the following URL:

<http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/tsd-products-support-series-home.html>

The Cisco Nexus 3164Q Switch - Read Me First is available at the following URL:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus3164/sw/6x/readme/b_Cisco_Nexus_3164Q_Switch_Read_Me_First.html

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.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation and gathering additional **information**, see the **monthly What's New in Cisco Product Documentation**, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

Open a service request online at:

<https://tools.cisco.com/ServiceRequestTool/create/launch.do>

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Obtaining Documentation and Submitting a Service Request

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Cisco Nexus 9000 Series NX-OS Release Notes, Release 6.1(2)I3(5)

© 2013 - 2015 Cisco Systems, Inc. All rights reserved.