

Overview

This chapter contains the following sections:

- About Nexus 9000v Platform Family, on page 1
- Cisco Nexus 9300v Platform, on page 2
- Cisco Nexus 9500v Platform, on page 6
- Nexus 9000v Throughput, on page 12
- Nexus 9000v Feature Support, on page 13
- Nexus 9000v Platform MIB Support, on page 16
- Nexus 9000v Platform Guidelines and Limitations, on page 17

About Nexus 9000v Platform Family

The Nexus 9000v is a virtual platform family that is designed to simulate control plane aspects of a standalone switch running Nexus 9000 software. This platform family uses the same software image that runs the Nexus 9000 hardware platforms. Although the virtual platforms don't attempt to simulate any specific ASICs or hardware SKUs, they are aligned with their hardware counterparts. An optimized software data plane handles the traffic across the line card interfaces. The Nexus 9000v virtual platform family consists of two virtual platforms: Nexus 9300v and Nexus 9500v. The sections that follow describe the capabilities of these two platforms.

The virtual platforms in the Nexus 9000v platform family allows you to simulate their network in a cost-effective manner. These platforms also support cloning of virtual machines (VMs). Use the simulated network to validate configurations prior to their application on a production network. Use these platforms to rapidly develop and test network automation tools using Cisco NX-OS programmability interfaces.

Minimum Memory Requirement on Host

Nexus 9000v requires a minimum of 10GB of RAM to boot up. Ensure that the underlying host (or laptop) has additional memory available.

Cisco Nexus 9300v Platform

The Cisco Nexus 9300v platform simulates a single supervisor non-modular chassis with a single co-located line card. This virtual chassis closely aligns with the standalone Cisco Nexus 9300 hardware platform running in the 'lxc' mode. The following tables show the specifications for this virtual platform:

Form-Factor

Component/Parameter	Specification
Usage	Simulation
Binary	Same as NX-OS Hardware
Management Interface	1
Line Cards	1
Line Card Interfaces	64

Resource Requirements

Resource	Specification
Minimum RAM	10.0 G (basic bootup)
Recommended RAM	12.0 G (depending on the number of features)
Minimum vCPUs	4
Recommended vCPUs	4
Minimum vNICs	1
Maximum vNICs	65

Deployment Environment

- KVM/QEMU 4.2.0
- ESXI 8.0
- Vagrant 2.3.7

To deploy a Nexus 9300v platform, fetch the appropriate virtual artifacts from Cisco CCO. The following table documents the supported virtual artifacts. After deploying the virtual machine, it reports itself as a Nexus 9300v platform.

Note The Cisco Nexus 9300v platform supported only the 32-bit image until Cisco NX-OS Release 10.1(x). Beginning from Cisco NX-OS Release 10.2(1), it supports only 64-bit image.

The following table displays the virtual artifact(s):

Hypervisor	Virtual Artifact	Description
ESXI 8.0	nexus9300v.10.1.1.ova	Contains virtual disk, machine definition, and NXOS image.
KVM/QEMU 4.2.0	nexus9300v.10.1.1.qcow2	Contains virtual disk and NXOS image on bootflash.
Vagrant 2.3.7	nexus9300v.10.1.1.box	Contains a preinstalled NXOS image on a virtual disk along with a machine definition.

After the initial virtual machine deployment, you can upgrade the Cisco NX-OS image on the platform using the typical NX-OS workflow (example: **install all** <>).

Note When you upgrade an older Nexus 9000v to the current release, it's automatically transformed into the Nexus 9300v. Even after subsequent reloads and NX-OS image upgrades, the platform will continue to present itself as a Nexus 9300v.

Cisco Nexus 9300v Platform Components

The Cisco Nexus 9300v platform, like its reference hardware counterpart, consists of three key components: chassis, supervisor, and one line card. The following table presents the product identifications (PIDs) and the SNMP sysOID used associated with the platform

Component	Description	PID	sysOID
Chassis	Nexus9000 C9300v Chassis	N9K-C9300v	EntPhysicalVendorType = cevChassisN9KV9000
Supervisor	Supervisor Module	N9K-vSUP	**Note** To ensure backward compatibility, the sysOID used in the previous release Nexus 9000v platform is reused for the Nexus 9300v platform.
Line Card	Nexus 9000v 64-port Ethernet Module	N9K-X9364v	

Cisco Nexus 9300v Chassis

The following references sample chassis-related information outputs from relevant **show** commands.

```
switch# sh version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Documents: http://www.cisco.com/en/US/products/ps9372/tsd_products_support_seri
es_home.html
Copyright (c) 2002-2022, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.

Nexus 9000v is a demo version of the Nexus Operating System

Software
BIOS: version
NXOS: version 10.2(3) [build 10.2(2.185)] [Feature Release]
BIOS compile time:
NXOS image file is: bootflash:///nxos64-cs.10.2.2.185.F.bin
NXOS compile time: 3/30/2022 13:00:00 [03/31/2022 00:30:59]

Hardware
cisco Nexus9000 C9300v Chassis
Intel(R) Xeon(R) CPU E5-2658 v4 @ 2.30GHz with 20499656 kB of memory.
Processor Board ID 9GFDLI2JD0R
Device name: switch
bootflash: 4287040 kB

Kernel uptime is 1 day(s), 23 hour(s), 35 minute(s), 21 second(s)

Last reset
Reason: Unknown
System version:
Service:

plugin
Core Plugin, Ethernet Plugin

Active Package(s):

switch#
switch# sh module
Mod Ports Module-Type Model Status
--- ----- ------------------------------------ ---------------------- ---------
1 64 Nexus 9000v 64 port Ethernet Module N9K-X9364v ok
27 0 Virtual Supervisor Module N9K-vSUP active *

Mod Sw Hw Slot
--- ---------------------- ------ ----
1 10.2(2.185) 0.0 LC1
27 NA 0.0 SUP1

Mod MAC-Address(es) Serial-Num
--- -------------------------------------- ----------
1 00-ed-c6-a2-01-01 to 00-ed-c6-a2-01-40 9JAUKW2T51G
27 00-ed-c6-a2-1b-01 to 00-ed-c6-a2-1b-12 9GFDLI2JD0R

Mod Online Diag Status
--- ------------------
1 Pass
27 Pass

* this terminal session
switch#
```

```
switch# sh inventory
NAME: "Chassis", DESCR: "Nexus9000 C9300v Chassis"
PID: N9K-C9300v , VID: , SN: 9ZQKP299FIZ

NAME: "Slot 1", DESCR: "Nexus 9000v 64 port Ethernet Module"
PID: N9K-X9364v , VID: , SN: 9JAUKW2T51G

NAME: "Slot 27", DESCR: "Supervisor Module"
PID: N9K-vSUP , VID: , SN: 9GFDLI2JD0R

switch#
```

Cisco Nexus 9300v Line Card

Cisco Nexus 9300v platform supports a single virtual line card with 64 virtual interfaces. The line card automatically populates when the platform boots. You can't insert or remove the line card from this chassis. The line card boot process starts after the Supervisor successfully boots and reaches the "active" state. Like its hardware counterpart, the line card boot-up starts with the "present" state and becomes fully functional when it reaches the "ok" state.

vNIC Mapping

On an actual Cisco Nexus 9300 hardware platform, you can "plug in" fibers to the front panel ports on a line card. On a virtual platform, like the Nexus 9300v, you must export the required number of virtual network interface cards/interfaces (vNICs) from the hypervisor into the Nexus 9300v platform.

The Nexus 9300v platform uses a sequential vNIC mapping. It maps the first vNIC passed in by the hypervisor into the Nexus 9300v management port. Subsequent vNICs are mapped sequentially into the line card interfaces. For example, if you export two vNICs onto the Nexus 9300v, the first vNIC is mapped to the NX-OS "mgmt" interface. The second vNIC is mapped to the "Ethernet1/1" interface.

vNIC Mapping Informational Show Commands

Show Platform vNIC Commands

On the Cisco Nexus 9300v platform, CLI commands are available to show the current vNIC mapping scheme, the number of vNICs mapped, and the mapping of MAC addresses to vNICs. Using these commands, you can ensure that the correct number of vNICs were passed to their virtual machine, and you can see which interfaces have been mapped.

Example outputs of the show vNIC platform commands:

show platform vnic mapped

```
v-switch# show platform vnic mapped
    NXOS Interface    VNIC MAC-Address    Internal VNIC
    --------------    ----------------    -------------
    Ethernet1/1       00c0.c000.0101      phyEth1-1
    Ethernet1/2       00c0.c000.0102      phyEth1-2
    Ethernet1/3       00c0.c000.0103      phyEth1-3
    Ethernet1/4       00c0.c000.0104      phyEth1-4
    Ethernet1/5       00c0.c000.0105      phyEth1-5
    Ethernet1/6       00c0.c000.0106      phyEth1-6
    Ethernet1/7       00c0.c000.0107      phyEth1-7
    Ethernet1/8       00c0.c000.0108      phyEth1-8
    Ethernet1/9       00c0.c000.0109      phyEth1-9
    Ethernet1/10      00c0.c000.010a      phyEth1-10
    Ethernet1/11      00c0.c000.010b      phyEth1-11
    Ethernet1/12      00c0.c000.010c      phyEth1-12
```

```
    Ethernet1/13      00c0.c000.010d      phyEth1-13
    Ethernet1/14      00c0.c000.010e      phyEth1-14
    Ethernet1/15      00c0.c000.010f      phyEth1-15
    Ethernet1/16      00c0.c000.0110      phyEth1-16
```

show platform vnic info

```
v-switch# show platform vnic info
    VNIC Scheme: Sequential
        mgmt0 interface: eth1  (00c0.c000.aabb)
        Module       # VNICs Mapped
        ------       --------------
        16           16
        ------       --------------
        VNICs passed: 16
        VNICs mapped: 16
        VNICs unmapped: 0
```

Cisco Nexus 9500v Platform

The Cisco Nexus 9500v simulates a single-supervisor platform 16 slot modular chassis that supports dynamic line card insertion and removal. This virtual chassis closely aligns with the standalone Cisco Nexus 9500 hardware platform. This version of Nexus 9500v currently doesn't simulate the system controller or fabric card typically found on the modular hardware chassis. This platform supports four different form factors of the generic line cards. These line cards share the same Linux kernel and differ only in the supported number of interfaces. The following tables show the specifications for this virtual platform.

Form-Factor

Component/Parameter	Specification
Usage	Simulation
Binary	Same as NX-OS Hardware
Management Interface	1
Line Cards	Up to 16
Line Card Interfaces	Up to 400 interfaces in the KVM/QEMU environment

Resource Requirements

Resource	Specification
Minimum RAM	10.0G (basic bootup with one line card; 1.2G for each additional line card)
Recommended RAM	12.0G (depending on the number of features)
Minimum vCPUs	4 (if you configure 16 line cards, we recommend 6 vCPUs)
Minimum vNICs	1

Resource	Specification
Maximum vNICs	400 interfaces in the KVM/QEMU environment

Deployment Environment

- KVM/QEMU 4.2.0
- ESXI 8.0

To deploy a Cisco Nexus 9500v platform, fetch the appropriate virtual artifacts from Cisco CCO. The following table documents the supported virtual artifacts. Once you deploy the virtual machine, it reports itself as a Nexus 9500v.



Note The Cisco Nexus 9500v platform supports only the 64-bit image of the Cisco NX-OS Release 10.1(1).

The table displays the virtual artifact(s):

Hypervisor	Virtual Artifact	Description
ESXI 8.0	nexus9500v64.10.1.1.ova	Contains virtual disk, machine definition, and NXOS image The 64-bit .ova file boots the N9500v platform, which in turn boots up the 64-bit image of Cisco NX-OS Release 10.1(1) Software. **Note** The Supervisor is 64-bit, and the line card is 32-bit.
KVM/QEMU 4.2.0	nexus9500v64.10.1.1.qcow2	Contains virtual disk and NXOS image on bootflash.

After the initial virtual machine deployment, you can upgrade the Cisco NX-OS image on the platform using the typical NX-OS workflow (example: **install all** <>).



Note You can't transform a Cisco Nexus 9000v from a previous release into a Nexus 9500v platform. You can change it by applying the Nexus 9500v virtual artifact.

Cisco Nexus 9500v Platform Components

The Cisco Nexus 9500v platform, like its reference hardware counterpart, consists of three key components: chassis, supervisor, and line cards. The following table presents the product identifications (PIDs) and the SNMP sysOID used associated with the platform

Component	Description	PID	sysOID
Chassis	Nexus9000 C9500v Chassis	N9K-C9500v	EntPhysicalVendorType = cevChassisN9KC9500v
Supervisor	Supervisor Module	N9K-vSUP	
Line Card	Nexus 9000v 64-port Ethernet Module	N9K-X9564v	

Cisco Nexus 9500v Chassis

The following references sample chassis-related information outputs from relevant **show** commands, for a Nexus 9500v platform with a single line card.

```
switch# sh version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Documents: http://www.cisco.com/en/US/products/ps9372/tsd_products_support_serie
s_home.html
Copyright (c) 2002-2022, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.
Some parts of this software are covered under the GNU Public
License. A copy of the license is available at
http://www.gnu.org/licenses/gpl.html.

Nexus 9000v is a demo version of the Nexus Operating System

Software
BIOS: version
NXOS: version 10.2(3) [build 10.2(2.191)] [Feature Release]
BIOS compile time:
NXOS image file is: bootflash:///nxos64-cs.10.2.2.191.F.bin
NXOS compile time: 4/5/2022 11:00:00 [04/05/2022 22:45:26]

Hardware
cisco Nexus9000 C9500v Chassis ("Supervisor Module")
Intel(R) Xeon(R) CPU E5-2658 v4 @ 2.30GHz with 16395468 kB of memory.

...skipping 1 line
Device name: switch
bootflash: 4287040 kB

Kernel uptime is 0 day(s), 0 hour(s), 7 minute(s), 51 second(s)

Last reset
Reason: Unknown
System version:
Service:

plugin
Core Plugin, Ethernet Plugin

Active Package(s):

switch#
switch# sh module
Mod Ports Module-Type Model Status
--- ----- -------------------------------- ---------------------- ---------
1 64 Nexus 9000v 64 port Ethernet Module N9K-X9564v ok
```

```
27 0 Virtual Supervisor Module N9K-vSUP active *

Mod Sw Hw Slot
--- ---------------------- ------ ----
1 10.2(2.191) 0.0 LC1
27 10.2(2.191) 0.0 SUP1

Mod MAC-Address(es) Serial-Num
--- -------------------------------------- ----------
1 00-66-1e-aa-01-01 to 00-66-1e-aa-01-40 9LKOT3OP6A5
27 00-66-1e-aa-1b-01 to 00-66-1e-aa-1b-12 9V3ZW14YP0M

Mod Online Diag Status
--- ------------------
1 Pass
27 Pass

* this terminal session
switch# sh inventory
NAME: "Chassis", DESCR: "Nexus9000 C9500v Chassis"
PID: N9K-C9500v , VID: , SN: 9NNVBPJ9N8A

NAME: "Slot 1", DESCR: "Nexus 9000v 64 port Ethernet Module"
PID: N9K-X9564v , VID: , SN: 9LKOT3OP6A5

NAME: "Slot 27", DESCR: "Supervisor Module"
PID: N9K-vSUP , VID: , SN: 9V3ZW14YP0M

switch#
```

Cisco Nexus 9500v Line Cards

The Cisco Nexus 9500v platform can support up to 16 virtual line cards. The platform supports five different forms of line cards. The differences between these virtual line cards are the number of NX-OS interfaces they support. By default, the Nexus 9500v platform boots with a single line card.

You can insert or remove virtual line cards using a CLI command. When using the sequential vNIC mapping scheme, insert the line cards sequentially from module 1. The removal operation must be in the opposite order. However, when using mac-encoded vNIC mapping scheme, you can insert line cards in any order. This mode supports sparse population of the line card.

The line card boot process starts after the Supervisor successfully boots and reaches the "active" state. Like their hardware counterparts, line card boot up starts with the "present" state and becomes fully functional when it reaches the "ok" state.

To support line insertion, the **platform insert module** *number* **linecard** command is used. A line card can be removed by prefixing the command with **no**. Once the chassis is configured with line cards, the chassis configuration remains persistent across switch reboots.

```
switch# platform insert ?
  module  Insert a specific module

switch# platform insert module?
  <1-30>  Please enter the module number

switch# platform insert module 2?
  <CR>
  linecard  Linecard module

switch# platform insert module 2 linecard ?
  N9K-X9516v  Nexus 9000v 16 port Ethernet Module
  N9K-X9532v  Nexus 9000v 32 port Ethernet Module
  N9K-X9536v  Nexus 9000v 36 port Ethernet Module
```

```
N9K-X9548v  Nexus 9000v 48 port Ethernet Module
N9K-X9564v  Nexus 9000v 64 port Ethernet Module
```

vNIC Mapping

On actual Cisco Nexus 9500 hardware platforms, you can "plug in" fibers to the front panel ports on a line card. On a virtual platform, like the Nexus 9500v, export the required number of virtual network interface cards/interfaces (vNICs) from the hypervisor into the Nexus 9500v platform.

The Nexus 9500v platform supports two vNIC mapping schemes. Depending on the specified scheme, the system maps the vNIC passed in by the hypervisor into the appropriate Nexus 9500v NX-OS interface. The following sections describe the capabilities of both the vNIC mapping schemes.

Sequential vNIC Mapping Scheme

In this scheme, vNICs acquired from the hypervisor are mapped into the NX-OS interfaces sequentially. For example, if you export two vNICs to the Nexus 9500v, the first vNIC is mapped to the NX-OS “mgmt” interface. The second vNIC is mapped to the “Ethernet1/1” interface. This is the default mapping mode when the virtual switch boots up for the first time.

Other limitations are:

- This mode doesn't support sparse population of line cards or interfaces.
- The line cards must be inserted in sequence from module 1 and removed in the opposite order.

  When a line card is removed, the vNICs remain within the system and are automatically remapped into the appropriate NX-OS interface upon reinsertion of the line card.

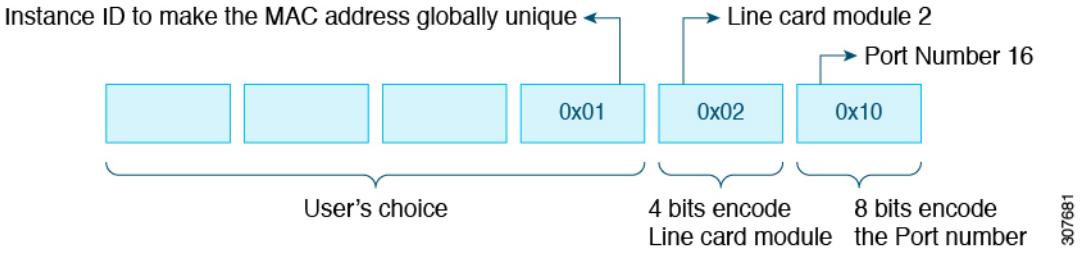
Use the **platform vnic scheme** command to select the vNIC mapping scheme. Once you select a mapping scheme, it remains persistent through a switch reload. By default, the Nexus 9500v switch boots up in the sequential vNIC mapping scheme. Changing this scheme requires a switch reload.

```
 switch# platform vnic ?
  scheme  Virtual Network Interface Card allocation scheme

switch# platform vnic scheme ?
  mac-encoded  MAC address encoded allocation of VNICs to linecard modules
  sequential   Sequential allocation of VNICs to linecard modules
```

MAC-Encoded vNIC Mapping Scheme

In this scheme, vNICs acquired from the hypervisor are mapped to NX-OS interfaces based on the MAC address configured on the vNIC (at the hypervisor level). This mode allows you to map any vNIC to any NX-OS line card interface. To use this mode, add the line card module and port number into the last 2 bytes of the vNIC MAC address. This MAC address configuration must be performed at the hypervisor level before powering up the Nexus 9500v virtual switch. The following diagram presents the required vNIC MAC address format:



Other features:

- This mode supports sparse population of line cards and interfaces.
- While the Nexus 9500v is in this mode, you can insert and remove the line cards in any order.
- When a line card is removed, the vNICs remain within the system and are automatically remapped into the appropriate NX-OS interface upon reinsertion of the line card.

Use the **platform vnic scheme** command to select the vNIC mapping scheme. Once you select a mapping scheme, it remains persistent through the switch reload. By default, the Nexus 9500v switch boots up in the sequential vNIC mapping scheme. Changing this scheme requires a switch reload.

```
 switch# platform vnic ?
  scheme  Virtual Network Interface Card allocation scheme

switch# platform vnic scheme ?
  mac-encoded  MAC address encoded allocation of VNICs to linecard modules
  sequential   Sequential allocation of VNICs to linecard modules
```



Note This scheme allows line card modules to be inserted and removed in any order. However, once the line cards have been inserted, the mapping scheme must not be changed to sequential unless absolutely necessary. If the mapping scheme is changed to "sequential", you must remove all line cards in non-sequential order and insert them back starting with module 1. If you insert line cards in non-sequential order, change the scheme to sequential, and reboot the switch, none of the line cards will come online. The switch, booted in the sequential scheme, expects any existing line cards to be present in serial order, starting with module 1. An error message similar to the following displays the mismatch in the vNIC scheme and the line cards inserted:

```
2020 Jul 15 14:44:03 N9Kv_3 %$ VDC-1 %$ %PLATFORM-2-MOD_INSERTION_FAILED:
Failed to insert module 6 (Nexus 9000v 64 port Ethernet Module - VNIC allocation scheme
is set to sequential, modules must be inserted in sequence)
```

You can validate the scheme by entering the **show platform vnic info** command. To recover from the above state, change the vNIC scheme back to MAC-encoded by entering the **platform vnic scheme mac-encoded** command and reboot the switch. If you require the vNIC scheme to be sequential, remove all line cards first before changing the scheme to sequential.

vNIC Mapping Informational Show Commands

Show Platform vNIC Commands

On the Cisco Nexus Nexus 9500v platform, CLI commands are available to show the current vNIC mapping scheme, the number of vNICs mapped, and the mapping of MAC addresses to vNICs. Using these commands, you can ensure that the correct number of vNICs were passed to their virtual machine, and you can see which interfaces have been mapped.

Example outputs of the show vNIC platform commands:

show platform vnic mapped

```
v-switch# show platform vnic mapped
    NXOS Interface    VNIC MAC-Address    Internal VNIC
    --------------    ----------------    -------------
    Ethernet1/1       00c0.c000.0101      phyEth1-1
    Ethernet1/2       00c0.c000.0102      phyEth1-2
    Ethernet1/3       00c0.c000.0103      phyEth1-3
    Ethernet1/4       00c0.c000.0104      phyEth1-4
```

```
Ethernet1/5       00c0.c000.0105      phyEth1-5
Ethernet1/6       00c0.c000.0106      phyEth1-6
Ethernet1/7       00c0.c000.0107      phyEth1-7
Ethernet1/8       00c0.c000.0108      phyEth1-8
Ethernet1/9       00c0.c000.0109      phyEth1-9
Ethernet1/10      00c0.c000.010a      phyEth1-10
Ethernet1/11      00c0.c000.010b      phyEth1-11
Ethernet1/12      00c0.c000.010c      phyEth1-12
Ethernet1/13      00c0.c000.010d      phyEth1-13
Ethernet1/14      00c0.c000.010e      phyEth1-14
Ethernet1/15      00c0.c000.010f      phyEth1-15
Ethernet1/16      00c0.c000.0110      phyEth1-16
```

show platform vnic info

```
v-switch# show platform vnic info
    VNIC Scheme: Mac-Encoded
    mgmt0 interface: eth1  (00c0.c000.aabb)
    Module     # VNICs Mapped
    ------     --------------
    1          16
    ------     --------------
    VNICs passed: 32
    VNICs mapped: 16
    VNICs unmapped: 16
```

Nexus 9000v Throughput

This section describes approximate throughput values for the Nexus 9000v virtual platform and the resource requirements to achieve the stated values.

4vCPU/16GB

The following table identifies the approximate throughput values for the Nexus 9000v with four virtual CPUs and 16 GB of RAM.

Feature(s)	Throughput on Stateless Traffic	Throughput on Stateful Traffic	
	UDP Protocol Traffic by Third-party Packet-generator-tool	TCP Protocol Traffic by file transfer	TCP Protocol Traffic by iPerf-tool
L2 switching, unicast forwarding	~700 Mbps	~50 Mbps	500 Mbps
L3 routing, unicast forwarding	~60 Mbps	~5 Mbps	50 Mbps

The following are the details of the throughput measurement:

Two Linux-VMs with iPerf-tools and one Third-party device were used for traffic generation and throughput measurement.

Topology used

- Nexus-9000v runs on a ESXI 8.0 hypervisor enabled UCS-device that were connected back-to-back to a Third-party Packet-generator-device.
- Nexus-9000v was connected back-to-back to Linux-VMs and all VMs were run on a ESXI 8.0 hypervisor enabled UCS-device.

Stateless (unidirectional) traffic testing

User defined UDP packets generated by Third-party device were sent through Nexus-9000v.

Stateful (bidirectional) traffic testing

- FTP based file (2GB sized) transfer was performed by Third-party device through Nexus-9000v.
- TCP traffic generated by iPerf-tools server/client VMs were sent through Nexus-9000v.

Nexus 9000v Feature Support

The Cisco Nexus 9000v platform family simulates a broad set of Nexus features. The forwarding plane of these features is implemented on a Cisco proprietary software data plane. Therefore, there can be some behavior differences. For example, the amount of system throughput is different between the virtual simulation platform and its hardware counterpart.

The following tables list the Cisco NX-OS features that have been tested on the Nexus 9000v platforms. You can configure and simulate untested NX-OS platform-independent features on the Nexus 9000v platforms. However, consider these features as unsupported. As more features are tested on the platform, the following table will be updated.

It's important to note that some of the chassis form-factor dependent feature commands can be available only on the corresponding Nexus 9000v platform. For example, NAT commands will be enabled only on the N9300 hardware platform and not on the N9500 hardware platform. It is also important to note that availability of a command does not imply that the feature is supported on the data plane. Please refer to the following feature tables for the supported features.

Layer 2 Features

The following table lists layer 2 feature support for the Nexus 9300v and Nexus 9500v platforms.

Feature	Nexus 9300v Support	Nexus 9500v Support
802.1AB LLDP	Yes	Yes
802.1Q VLANs/Trunk	Yes	Yes
802.1s RST	Yes	Yes
802.3ad LACP	Yes	Yes
L2 Multicast	Yes (as broadcast)	Yes (as broadcast)
vPC/MLAG	Yes	Yes
Port Channel	Yes	Yes

Feature	Nexus 9300v Support	Nexus 9500v Support
VLANs	Yes	Yes

Layer 3 Features

The following table lists layer 3 feature support for the Nexus 9300v and Nexus 9500v platforms.

Feature	Nexus 9300v Support	Nexus 9500v Support
OSPF	Yes	Yes
OSPFv3	Yes	Yes
BGP	Yes	Yes
MP-BGP	Yes	Yes
IS-IS	Yes (as broadcast)	Yes (as broadcast
RIPv2	Yes	Yes
Equal Cost Multipath Routing (ECMP)	Yes	Yes
PIM-SM	Yes	Yes
HSRP	Yes	Yes
VRRP	Yes	Yes
MPLS	Yes	Yes
EIGRP	Yes	Yes
CDP	Yes	Yes
L3 SVI	Yes	Yes
Sub Interfaces	Yes	Yes
IPsec	Yes	Yes

VXLAN and Segment Routing Features

The following table lists VXLAN and segment routing feature support for the Nexus 9300v and Nexus 9500v platforms.

Feature	Nexus 9300v Support	Nexus 9500v Support
VXLAN flood and Learn BUM Replication (PIM/ASM, IR)	(Yes, Yes)	(Yes, Yes)
VXLAN EVPN BUM Replication (PIM/ASM, BIDIR, IR)	(Yes, No, Yes)	(Yes, No, Yes)
VXLAN EVPN Routing	Yes	Yes
VXLAN EVPN Bridging	Yes	Yes
VXLAN EVPN Anycast GW	Yes (as broadcast)	Yes (as broadcast)
VXLAN Tunnel Endpoint	Yes	Yes
VXLAN ARP Suppression	Yes	Yes
VXLAN EVPN Multi-Site BGW	Yes (with non-vPC on border gateway (BGW))	Yes (with non-vPC on border gateway (BGW))
VXLAN EVPN TRM	No	No
VXLAN IPv6 Underlay	Yes	Yes
MPLS Segment Routing (SRv4)	No	No
Downstream VNI	Yes	Yes
vPC with Fabric Peering [1]	Yes	Yes
L3VNI without VLAN	No	No

[1] The vPC Fabric Peering peer-link is established over the transport network (the spine layer of the fabric). DSCP is not supported on N9000v/N9300v/N9500v; hence, overloading the spine with traffic burst may result in disruptions on the vPC state of the leaf switches.

Programmability Features

The following table lists programmability feature support for the Nexus 9300v and Nexus 9500v platforms.

Feature	Nexus 9300v Support	Nexus 9500v Support
Bash shell access and scripting	Yes	Yes
RPM support	Yes	Yes
Programmatic access to system state (Python)	Yes	Yes

Feature	Nexus 9300v Support	Nexus 9500v Support
Guest Shell within OS	Yes	Yes
Docker within OS	Yes	Yes
NXAPI	Yes	Yes
DME	Yes	Yes
RESTCONF	Yes	Yes
NETCONF	Yes	Yes
YANG Models	Yes	Yes
Telemetry	Yes	Yes
GNMI	Yes	Yes
NxSDK	Yes	Yes

Guestshell Support

The Cisco Nexus 9000v supports Nexus Guestshell. However, this feature isn't enabled by default. You can use Guestshell functionality by following proper Nexus Guestshell documentation to enable it.



Note The Cisco Nexus 9000v virtual artifacts currently have enough bootflash size to fit two binary images. However, Guestshell requires more bootflash disk size once enabled. There may not be enough space in bootflash to boot two binary images respectively in released virtual artifacts. Plan to clear enough disk space if you need to enable Guestshell.

Nexus 9000v Platform MIB Support

The Nexus 9000v platforms support the Simple Network Management Protocol (SNMP) along with many of the Cisco NX-OS SNMP MIBs. Note that some of the managed objects may not be relevant to the simulation platform and may not be supported.

The following list shows supported platform-related MIBs:

- CISCO entity Asset MIB
- ceEXTEntityLEDTable
- ciscoEntityExtMIB
- ciscoRFMIB
- ciscoTSMIB
- ciscoEntityFRUControlMIB

- ciscoSyslogMIB

Nexus 9000v Platform Guidelines and Limitations

The following guidelines and limitations apply to the Nexus 9000v platforms:

- Hardware consistency checker isn't supported
- Nondisruptive ISSU isn't supported
- Link carrier status of NX-OS interface depends on the capability of the underlying hypervisor.

Nested VM use cases are supported. However, take care about deploying Nexus 9000v VMs in different environments, based on requirements. In a nested VM environment, performances in terms of bootup time, admin operation, and feature application, are degraded depending on the available vCPU and memory resources.

- If you want to simulate large number of nodes, or more than 10 Nexus 9000v nodes, use a Bare metal environment. VIRL2 is currently not supported.
- Bootup time takes longer when simulating large number of nodes.
- Sparse mode is currently not supported in the VIRL environment.
- Beginning with Cisco NX-OS Release 10.1(1), when multiple VMs are needed to boot on an ESXi at the same time, it is recommended to boot a cluster of 2 to 3 VMs at one time, then, after providing some time interval, the next cluster of 2 to 3 VMs should be booted. Otherwise, the VMs may become unresponsive.
- Beginning with Cisco NX-OS Release 10.2(3)F, fcoe/fcoe-npv feature set support is added on Cisco Nexus 9300v platform switches, but not on 9500v platform switches.
- Beginning with Cisco NX-OS Release 10.2(3)F, a lite image is introduced with a reduced memory footprint, resulting in a much smaller size image than the earlier images for Cisco Nexus 9300v and 9500v.
- Telemetry is not supported on N9kv lite image, so gNMI testing is not possible on N9kv lite image. However, beginning with Cisco NX-OS Release 10.4(2)F, this is supported.
- Beginning with Cisco NX-OS Release 10.6(1)F, MLD snooping is supported on Cisco N9000V switch. For more information, see Configuring MLD (Cisco Nexus 9000 Series NX-OS Multicast Routing Configuration Guide).
- If the user upgrades ESXi version to 8.0 or later, modify the existing value (1) of Net.BlockGuestBDPU to zero (0) in the **Manage > System > Advance** setting of ESXi Host Management UI to ensure that LACP works on N9kv.