



Troubleshooting the Cisco Nexus 9000v

This chapter contains the following sections:

- [Troubleshooting the Cisco Nexus 9000v Platform, on page 1](#)
- [Troubleshooting the Cisco Nexus 9000v Dataplane, on page 6](#)

Troubleshooting the Cisco Nexus 9000v Platform

General Troubleshooting/Debugging

The following CLI command provides troubleshooting help for both the Nexus 9300v and Nexus 9500v platforms:

show tech-support nexus9000v

The following is an example output of this command:

```
switch# show tech-support nexus9000v

----- Virtual Chassis Manager Debugs -----

#####
# /cmn/pss/virt_cmgr.log
#####
[19-12-10 20:42:34.160609]: virt_cmgr_startup_init called
[19-12-10 20:42:34.161351]: virt_cmgr_validate_file returned success
[19-12-10 20:42:34.161390]: Version 1, VNIC_scheme 2
[19-12-10 20:42:34.161404]: VM sup1: Module no 26, upg_version 1, type 1, card_i
ndex 0, image loc None
...
...
...
```

Common Issues for All Hypervisors

Boot when VM drops into "loader >" prompt

Generally, the initial boot is successful. However, the system boot could fail and drop into the "loader >" prompt on the VGA console or serial console, depending on how you provisioned the VM.

Example:

```

Loader Version 5.9

Loader > dir

bootflash::

  .rpmstore
  nxos.9.3.2.20.bin
  bootflash_sync_list
  .swtam
  eem_snapshots
  virtual-instance
  scripts
  platform-sdk.cmd

loader > boot nxos.9.3.2.20.bin

```

To continue the boot, enter the **boot nxos.9.3.2.20.bin** command at the "loader >" prompt

Prevent VM from dropping into "loader >" prompt

After you set up your Cisco Nexus 9000v (and following the set-up of the POAP interface), configure the boot image in your system to avoid dropping to the "loader >" prompt after reload/shut down.

Example:

```

nx-osv9000-2# config t
Enter configuration commands, one per line. End with CNTL/Z.
nx-osv9000-2(config)# boot nxos bootflash:nxos.9.3.2.20.bin
Performing image verification and compatibility check, please wait....
nx-osv9000-2(config)# copy running-config startup-config

```

Bootup Warning Message

During bootup, you may get a warning message similar to the following:

```

Checking all filesystems. **Warning** : Free memory available in bootflash is
553288 bytes
need at least 2 GB space for full image installation ,run df -h

```

This message generally indicates that the Nexus 9000v bootflash doesn't have enough memory space for holding another image. To eliminate this warning message, free up bootflash space to allow for the download of another binary image.

Nexus 9000v Mac-Encoded Mode Network Mapping Check

This check is only relevant if you explicitly enter the **platform vnic scheme mac-encoded** command on Nexus 9500v platform. This command enables the vNIC mac-encoded scheme. If any data traffic passes, or vNIC-mapped interfaces show the "Link not connected" state, refer to the Nexus 9000v informational show commands to verify correct vNIC mapping.

ESXi Hypervisor Issues

Nexus 9000v boot not seen after powering on the VM

The likely cause of this issue is that the EFI boot isn't set in the VM configuration. To resolve this issue, refer to the ESXi deployment guide to change "BIOS" to "EFI" in **Edit virtual machine settings > VM Options > Boot Options** after deployment using the distributed OVA virtual artifacts.

Bootup logs not seen after VGA output

A common problem during ESXi bootup is that the VGA console displays output similar to the following:

```
Sysconf checksum failed. Using default values
console (dumb)

Booting nxos.9.3.2.6.bin...
Booting nxos.9.3.2.bin
Trying diskboot
  Filesystem type is ext2fs, partition type 0x83
Image valid

Image Signature verification for Nexus9000v is not performed.

Boot Time: 12/5/2019 10:38:41
```

The issue is that, in the VGA console, there's no following activity in the bootup process. It's often misunderstood as a switch bootup process hang. To see the output of a switch bootup, connect to the provisioned serial console based on steps provided in the ESXi hypervisor deployment guide.

If nothing happens in the serial console, or you see the "telnet: Unable to connect to remote host: Connection refused" error message, it indicates one or more of the following issues:

- The serial console provisioning is incorrect in the VM configuration. Read and follow the instructions for serial console connectivity in the ESXi deployment guide.
- ESXi 8.0 deployment is the only version supported. Make sure that you have a valid license for ESXi vCenter and a valid UCS server license.
- Make sure that the "Security Profile" in the server has "VM serial port connected over network", both for incoming connections and outgoing connections.

No access to "loader>" prompt after powering down the VM

This issue occurs if you power on the VM and it boots up as expected, but the serial console wasn't correctly provisioned. Then the "config t; boot nxos bootflash:nxos.9.3.2.20.bin" configure is performed and saved. Powering up the VM again results in a drop to the VGA console.

The following recommendations help to avoid this issue in the ESXi hypervisor.

EFI BIOS defaults all input/output to the VM console. When a VM drops to the "loader >" prompt, go to the vSphere client or VGA console to access the "loader >" prompt to boot the image in the hard disk. You can change this behavior by adding an extra configuration in the ESXi VM editing mode. Use one of the following methods:

1. In the vSphere client Configuration Parameters window, add one row in the configuration (Edit Settings > VM Options > Advanced > Edit Configuration).
2. Add `efi.serialconsole.enabled = "TRUE"` to the `.vmx` file once the VM is created.

The vCenter or UCS server connectivity is lost as soon as the Cisco Nexus 9000v is up

**Caution**

When connecting a vNIC to a vSwitch or bridge, an incorrect network connection might result in losing the connectivity to your hypervisor server or vCenter on ESXi.

The Cisco Nexus 9000v uses vNICs entered from a graphical representation on ESXi for networking, either externally or internally within a hypervisor server. The first NIC is always used as the Cisco Nexus 9000v management interface.

The first NIC in the Cisco Nexus 9000v VM is the management interface. Connect it directly to your lab LAN physical switch or vSwitch (VM Network). Don't connect any data port vNIC to any physical switch conflicting with your server management connectivity.

Cisco Nexus 9000v data port isn't passing traffic in the ESXi server

To ensure a smooth operation, specific configuration settings on the vSwitch must be enabled:

- Ensure that all instances of the vSwitch connecting to the Cisco Nexus 9000v are in "Promiscuous Mode" = "Accept", and pointing to the UCS server. You can access this option through "Configuration > Properties > Edit" from the vSphere Client.
- Ensure that all instances of vSwitch pass through all VLANs. You can access this option through "Configuration > Properties > Edit" from the vSphere Client.

ESXi 8.0 hypervisor often defaults the network interfaces adapter to the "E1000E" type which isn't supported in the Nexus 9000v platform. After deployment, make sure that all Network adapter types are "E1000".

KVM/QEMU Hypervisor Issues

Understanding the KVM/QEMU command line options requires a basic Linux background. In order to deploy the Nexus 9000v in this hypervisor, follow the deployment instruction and pay attention to the following areas:

- Make sure that the user guide recommends bios.bin.
- If the command line supports multiple disk inputs, check that the bootable disk is set to bootindex=1 so that the VM doesn't try to boot from other devices.
- If you're attempting to implement a complicated command line, follow basic KVM/QEMU deployment instruction to bring up a simple switch instance first to verify the user environment.

Multicast on KVM or QEMU Hypervisor

The multicast feature on the Cisco Nexus 9000v is supported as broadcast. To make this feature to work properly, disable IGMP multicast snooping in this environment on all bridge interfaces.

The following example shows how to disable vxlan_br1, vxlan_br2, vxlan_br3, and vxlan_br4 from the linux prompt:

```
echo 0 > /sys/devices/virtual/net/vxlan_br1/bridge/multicast_snooping
echo 0 > /sys/devices/virtual/net/vxlan_br2/bridge/multicast_snooping
echo 0 > /sys/devices/virtual/net/vxlan_br3/bridge/multicast_snooping
echo 0 > /sys/devices/virtual/net/vxlan_br4/bridge/multicast_snooping
```

Follow the Linux bridge mask setup in the KVM/QEMU deployment guide, for passing L2 packets such as LLDP, LACP, and others.

Vagrant/VirtualBox Issues

Networking on VirtualBox/Vagrant

To use the dataplane interfaces on VirtualBox/Vagrant, ensure the following:

- The interfaces must be in "Promiscuous" mode.
- In the VirtualBox network settings, select "Allow All" for the Promiscuous mode.
- Ensure all instances of Cisco Nexus 9000v in your topology have unique MAC addresses by using the **show interface mac** command.

VM normal bootup on VirtualBox/Vagrant:

```
Bringing machine 'default' up with 'virtualbox' provider...
==> default: Clearing any previously set forwarded ports...
==> default: Clearing any previously set network interfaces...
==> default: Preparing network interfaces based on configuration...
      default: Adapter 1: nat
==> default: Forwarding ports...
      default: 22 (guest) => 2222 (host) (adapter 1)
==> default: Booting VM...
==> default: Waiting for machine to boot. This may take a few minutes...
      default: SSH address: 127.0.0.1:2222
      default: SSH username: vagrant
      default: SSH auth method: private key
```

The configured shell (config.ssh.shell) is invalid and unable to properly execute commands. The most common cause for this is using a shell that is unavailable on the system. Please verify you're using the full path to the shell and that the shell is executable by the SSH user.

The **vagrant ssh** command will access the Nexus 9000v switch prompt after the successful normal bootup.

The following is an example of one possible VM bootup failure:

```
Bringing machine 'default' up with 'virtualbox' provider...
==> default: Importing base box 'base'...
==> default: Matching MAC address for NAT networking...
==> default: Setting the name of the VM: n9kv31_default_1575576865720_14975
==> default: Clearing any previously set network interfaces...
==> default: Preparing network interfaces based on configuration...
      default: Adapter 1: nat
==> default: Forwarding ports...
      default: 22 (guest) => 2222 (host) (adapter 1)
==> default: Booting VM...
==> default: Waiting for machine to boot. This may take a few minutes...
      default: SSH address: 127.0.0.1:2222
      default: SSH username: vagrant
      default: SSH auth method: private key
```

Timed out while waiting for the machine to boot. This means that Vagrant was unable to communicate with the guest machine within the configured ("config.vm.boot_timeout" value) time period.

If you look above, you should be able to see the error(s) that Vagrant had when attempting to connect to the machine. These errors are usually good hints as to what may be wrong.

If you're using a custom box, make sure that networking is properly working and you're able to connect to the machine. It is a common problem that networking isn't setup properly in these boxes. Verify that authentication configurations are also setup properly,

as well.

If the box appears to be booting properly, you may want to increase the timeout ("config.vm.boot_timeout") value.

To troubleshoot this failure, check the following:

- Ensure that enough resources, such as memory and vCPU, are available. Close all applications that consume a significant amount of memory in your PC or server. Check the available free memory.
- Power down VM by entering **vagrant halt -f**
- Go to the VirtualBox GUI after powering down the VM. Enable the VM serial console to observe the boot up process and to view possible issues through "Ports" -> "Enable Serial Port".

Alternatively, use the following VBox command to enable this guest serial console. Find your VM name:

```
VBoxManage list vms
    "n9kv_default_1575906706055_2646" {0b3480af-b9ac-47a4-9989-2f5e3bdf263f}
```

Then enable serial console:

```
VBoxManage modifyvm n9kv_default_1575906706055_2646 --uart1 0x3F8 4
```

- Power up the VM again by entering “vagrant up” from the same terminal, where you entered the original “vagrant up”.
- To access the serial console, enter “telnet localhost 2023” from another terminal on your computer.
- Check the bootup issue by observing the output from the serial console.
- Turn off the serial console if the guest serial console is no longer needed. Either use the following VBox command or go to the VirtualBox GUI setting and de-select “Enable Serial Port”.

```
VBoxManage modifyvm n9kv_default_1575906706055_2646 --uart1 off
```

Troubleshooting the Cisco Nexus 9000v Dataplane

The debug and show commands in this section are available to troubleshoot both Cisco Nexus 9300v and Cisco Nexus 9500v platforms. These commands must be executed on the line card/module.

Debug Commands

- **debug l2fwder event**
- **debug l2fwder error**
- **debug l2fwder fdb**
- **debug l2fwder pkttrace**

To run any of these commands, attach to the line card by following this example:

```
switch# sh mod | inc Mod
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	Serial-Num
Mod	MAC-Address (es)			

```

Mod Online Diag Status
switch# attach mod 1
Attaching to module 1 ...
To exit type 'exit', to abort type '$.'
module-1# debug l2fwder ?
  error      Configure debugging of l2fwder control and data path errors
  event      Configure debugging of l2fwder events over ipc
  fdb        Configure debugging of l2fwder events over fdb
  ha         Configure debugging of l2fwder events from sysmgr
  logfile    Enable file logging to /logflash/l2fwder.debug
  packet     Configure debugging of l2fwder packet forwarding information
  pkttrace   Configure debugging of l2fwder packet trace

module-1# debug l2fwder

```

Event History Commands

- show system internal l2fwder event-history events
- show system internal l2fwder event-history errors
- show system internal l2fwder event-history fdb

Show Commands

show system internal l2fwder table bd

```

v-switch# show system internal l2fwder table bd

vlan 1 member 3, 4, 5, untagged 3, 4, 5, STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan no
vlan 80 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes
vlan 90 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes
vlan 110 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes
vlan 210 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes
vlan 310 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes
vlan 410 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes
vlan 510 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes
vlan 550 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan no
vlan 560 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan no
vlan 610 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes
vlan 650 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan no
vlan 660 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan no
vlan 710 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes
vlan 810 member 3, 4, 5, untagged none STP ports 3, 4, 5, dis none blk_lis none
lrn none fwd 3, 4, 5, tid 1, 2, vxlan yes

```

show system internal l2fwder table if

v-switch# show system internal l2fwder table if

If_name	If_index	gport	fd	untagged	vlanid	Trunk	SVP Info	Native vlan
Ethernet1/1	0x1a000000	0x8000801	14	1	4095	0x0	none	4095
Ethernet1/2	0x1a000200	0x8000802	15	1	4095	0x0	none	4095
Ethernet1/3	0x1a000400	0x8000803	16	0	4045	0x1	none	40451
Ethernet1/4	0x1a000600	0x8000804	17	0	810	0x2	none	810
Ethernet1/5	0x1a000800	0x8000805	18	0	810	0x0	none	810
Ethernet1/6	0x1a000a00	0x8000806	0	1	4095	0x0	none	4095
Ethernet1/7	0x1a000c00	0x8000807	0	1	4095	0x0	none	4095
Ethernet1/8	0x1a000e00	0x8000808	0	1	4095	0x0	none	4095
Ethernet1/9	0x1a001000	0x8000809	0	1	4095	0x0	none	4095
Ethernet1/10	0x1a001200	0x800080a	0	1	4095	0x0	none	4095
Ethernet1/11	0x1a001400	0x800080b	0	1	4095	0x0	none	4095

show system internal l2fwder table port-channel

v-switch# show system internal l2fwder table port-channel

Port-channel	Count	Member-list
0x1	1	0x8002004
0x4	2	0x8005001 0x8000805
0x5	2	0x8002001 0x8000801

Port-channel	Count	Local member-list6
0x1	0	
0x4	1	0x8000805
0x5	1	0x8000801

show system internal l2fwder table vxlan peer

v-switch# show system internal l2fwder table vxlan peer

```
VXLAN Tunnel:
  src_ip: 6.6.6.6, Is VxLAN enabled = TRUE
  multisite: no, nve_tun_dci_sip: 0.0.0.0
VXLAN PEER: No of tunnels = 7
  peer_ip: 224.1.1.2, vxlan_port_id: 0x0,
    tunnel_id: 0x4c000000, is_dp: 0 is_dci: 0
  peer_ip: 224.1.1.4, vxlan_port_id: 0x0,
    tunnel_id: 0x4c000002, is_dp: 0 is_dci: 0
  peer_ip: 224.1.1.6, vxlan_port_id: 0x0,
    tunnel_id: 0x4c000004, is_dp: 0 is_dci: 0
  peer_ip: 224.1.1.8, vxlan_port_id: 0x0,
    tunnel_id: 0x4c000006, is_dp: 0 is_dci: 0
  peer_ip: 224.1.1.9, vxlan_port_id: 0x0,
    tunnel_id: 0x4c000008, is_dp: 0 is_dci: 0
  peer_ip: 224.1.1.10, vxlan_port_id: 0x0,
    tunnel_id: 0x4c00000a, is_dp: 0 is_dci: 0
  peer_ip: 6.5.5.5, vxlan_port_id: 0x80002db8,
    tunnel_id: 0x4c00050a, is_dp: 0 is_dci: 0
Tunnel_id entry:
  peer_ip: 224.1.1.2, tunnel_id: 0x4c000000
  peer_ip: 224.1.1.4, tunnel_id: 0x4c000002
  peer_ip: 224.1.1.6, tunnel_id: 0x4c000004
  peer_ip: 224.1.1.8, tunnel_id: 0x4c000006
  peer_ip: 224.1.1.9, tunnel_id: 0x4c000008
  peer_ip: 224.1.1.10, tunnel_id: 0x4c00000a
  peer_ip: 6.5.5.5, tunnel_id: 0x4c00050a
```



```
Vxlan_gport ucast-entry:
  peer_ip: 6.5.5.5, vxlan_port_id: 0x80002db8
```

show system internal l2fwder table vxlan vni

```
v-switch# show system internal l2fwder table vxlan vni
```

VNI	VLAN	DF
81000	810	no
51000	510	no
5001	1001	no
5002	1002	no
5003	1003	no
5004	1004	no
21000	210	no
71000	710	no
9000	90	no
41000	410	no
11000	110	no
61000	610	no
31000	310	no

show system internal l2fwder acl info

```
v-switch# show system internal l2fwder acl info
```

Inactive List:

Entry ID: 14596 Qualify: DstTrunk 4, Action: RedirectTrunk 5 Prio: 4

Active List:

Inactive List:

Active List:

Entry ID: 15873 Qualify: EtherType ARP ForwardingVlanId 110, 610, 710, 1001, 1003, Action: CopyToCpu SET Drop SET Prio: 1

show system internal l2fwder mac

```
v-switch# show system internal l2fwder mac
```

Legend:

- * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
- + - primary entry using vPC Peer-Link,
- (T) - True, (F) - False, C - ControlPlane MAC

VLAN	MAC Address	Type	Secu	NTF	Del	Ports	Station_id
* 1	008b.860d.1b08	static	F	F	0	0xc000005	0
G -	008b.860d.1b08	static	F	F	0	sup-eth1 (R)	508,
* 210	0000.4545.6767	dynamic	F	F	0	0xc000004	0
G 710	008b.bc90.1b08	static	F	F	0	sup-eth1 (R)	0
G 310	008b.bc90.1b08	static	F	F	0	sup-eth1 (R)	0
G -	0002:0002:0002	static	F	F	0	sup-eth1 (R)	1,
* 210	008b.860d.1b08	static	F	F	0	0xc000005	0
G 410	008b.bc90.1b08	static	F	F	0	sup-eth1 (R)	0
* 1003	008b.2b34.1b08	dynamic	F	F	1	nve (0x80002db9)	0
* 1002	008b.2b34.1b08	dynamic	F	F	1	nve (0x80002db9)	0
* 1001	008b.2b34.1b08	dynamic	F	F	1	nve (0x80002db9)	0

```
* 1004 008b.2b34.1b08 dynamic F F 1 nve(0x80002db9) 0
* 810 008b.860d.1b08 static F F 0 0xc000005 0
G 510 008b.bc90.1b08 static F F 0 sup-eth1(R) 0
* 610 008b.2b34.1b08 dynamic F F 1 nve(0x80002db9) 0
G 1 008b.bc90.1b08 static F F 0 sup-eth1(R) 0
G - 008b:bc90:1b08 static F F 0 sup-eth1(R) 511,
```

show system internal l2fwder port egress info

```
v-switch# show system internal l2fwder port egress info
```

```
Ingress port :          Blocked egress ports
+-----+-----+
0x8002001      1          5
0x8000801      1          5
0x8020821      1          5
```

show system internal l2fwder vpc info

```
v-switch# show system internal l2fwder vpc info

VPC role : Primary
```

Packet Capture Commands

The Cisco Nexus N9000v supports Ethalyzer similarly to the standalone Nexus 9000 hardware switch.