



# Cisco Nexus1000V Release Notes, Release 4.0(4)SV1(1)

---

**Updated: May 10, 2013**  
**OL-19420-01**

This document describes the features, caveats, and limitations for the Cisco Nexus 1000V Release 4.0(4)SV1(1) software. Use this document in combination with documents listed in the [“Related Documentation”](#) section on page 14.

## Contents

This document includes the following information about Release 4.0(4)SV1(1) of the Cisco Nexus 1000V.

- [Introduction, page 1](#)
- [Software Compatibility, page 2](#)
- [Limitations and Restrictions, page 2](#)
- [Caveats, page 8](#)
- [MIB Support, page 14](#)
- [Related Documentation, page 14](#)

## Introduction

The Cisco Nexus 1000V provides a distributed, layer 2 virtual switch that extends across many virtualized hosts. The Cisco Nexus 1000V manages a Datacenter defined by the vCenter Server. Each server in the Datacenter is represented as a linecard in Cisco Nexus 1000V and can be managed as if it were a line card in a physical Cisco switch.

Cisco Nexus 1000V consists of the following two components:

- Virtual Supervisor Module (VSM), which contains the Cisco CLI, configuration, and high-level features



---

**Americas Headquarters:**  
**Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA**

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

- Virtual Ethernet Module (VEM), which acts as a Line Card and runs in each virtualized server to handle packet forwarding and other localized functions, and is compatible with any upstream physical access layer switch that is Ethernet standard's compliant. This includes Catalyst and Nexus switches from Cisco as well as switches from other network vendors.

## Software Compatibility

The servers running the Cisco Nexus 1000V VSM and VEM must be in the [VMware Hardware Compatibility List \(HCL\)](#) . This is a requirement for running the ESX 4.0 software, VMWare vSphere 4.0 Enterprise Plus.

## Limitations and Restrictions

- [Configuration Limits, page 3](#)
- [Access Lists, page 4](#)
- [Netflow, page 4](#)
- [Port Security, page 5](#)
- [Port Profile, page 5](#)
- [Telnet Enabled by Default, page 5](#)
- [SSH Support, page 5](#)
- [Cisco NX-OS Commands May Differ from Cisco IOS, page 5](#)
- [Layer 2 Switching, page 5](#)
- [Cisco Discovery Protocol, page 6](#)
- [DHCP Not Supported for the Management IP, page 7](#)
- [LACP, page 7](#)

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

## Configuration Limits

Use the following configuration limits with Cisco Nexus 1000V:

Component	Supported Limit	
<b>Maximum Modules:</b>	66	
Virtual Ethernet (VEM)	64	
Virtual Supervisor (VSM)	2	
Hosts	64	
Active VLANs across all VEMs	512	
MACs over VLAN within a VEM	1024 (1K)	
vEthernet interfaces per port profile	1024 (1K)	
PVLAN	256	
Distributed Virtual Switches (DVS) per vCenter	12	
	Per DVS	Per Host
vEthernet interfaces	2K	216
Port Profiles	256	
System Port Profiles		16
Port Channel	256	8
Physical Trunks	512	
Physical NICs		32
vEthernet Trunks	256	8
ACL	128	16 <sup>1</sup>
ACEs per ACL	128	16 <sup>1</sup>
ACL Interfaces	1024	128
NetFlow Policies	32	8
NetFlow Interfaces	256	32
SPAN/ERSPAN Sessions	64	4
QoS Policy-Map	128	16
QoS Class-Map	1024	128
QoS Interfaces	1024	128
Port Security	1024	216
MultiCast Groups	512	64

1. This number can be exceeded if VEM has available memory.

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

## VMware Lab Manager

VMware Lab Manager does not support using the Cisco Nexus 1000V.

## Access Lists

ACLs have the following limitations and restrictions:

### Limitations:

- IPV6 ACL rules are not supported.
- VLAN-based ACLs (VACLs) are not supported.

### Restrictions:

- IP ACL rules for TCP and UDP traffic cannot use logical operator *neq* (not equal to) to filter traffic based on port numbers.
- IP ACL rules do not support the following:
  - established TCP connections filtering option
  - fragments option
  - addressgroup option
  - portgroup option
  - interface ranges
- Control VLAN traffic between the VSM and VEM does not go through ACL processing.

## Netflow

The Netflow configuration has the following support, limitation, and restrictions:

- L2 match fields are not supported.
- Netflow Sampler is not supported.
- Netflow Exporter format V9 is supported
- Netflow Exporter format V5 is not supported.
- Multicast traffic type is not supported. Cache entries are created for multicast packets but packet/byte count does not reflect replicated packets.

The Netflow cache table has the following limitation:

- Immediate and Permanent cache types are not supported.



### Note

The cache size configured using the CLI defines the number of entries and not the size in bytes. The configured entries are allocated for each processor in the ESX host and the total memory allocated depends on the number of processors.

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

## Port Security

Port Security has the following support, limitations, and restrictions:

- The Port Security feature is enabled globally by default.  
The CLI command **feature/no feature port-security** is not supported.
- In response to a security violation, you can shut down the port
- Port Security Violation Actions that are supported on a Secure port are **Shutdown** and **Protect**. The **Restrict** Violation Action is not supported.
- Port Security is not supported on the PVLAN promiscuous ports.

## Port Profile

Port profiles have the following restrictions or limitations:

- If you attempt to remove a port profile that is in use, that is, one that has already been auto-assigned to an interface, the Cisco Nexus 1000V generates an error message and does not allow the removal.
- When you remove a port profile that is mapped to a VMware port group, the associated port group and settings within the vCenter Server are also removed.
- Policy names are not checked against the policy database when ACL/Netflow policies are applied through port profile. It is possible to apply a non-existent policy.

## Telnet Enabled by Default

The Telnet server is enabled by default.

For more information about Telnet, see the *Cisco Nexus 1000V Security Configuration Guide, Release 4.0(4)SV1(1)*.

## SSH Support

Only SSH version 2 (SSHv2) is supported.

For more information, see the *Cisco Nexus 1000V Security Configuration Guide, Release 4.0(4)SV1(1)*.

## Cisco NX-OS Commands May Differ from Cisco IOS

Be aware that the Cisco NX-OS CLI commands and modes may differ from those used in Cisco IOS.

For information about the CLI, see the *Cisco Nexus 1000V Command Reference, Release 4.0(4)SV1(1)*.

## Layer 2 Switching

This section lists the Layer 2 switching limitations and restrictions and includes the following topics:

- [No Spanning Tree Protocol, page 6](#)
- [MAC Address Table, page 6](#)

**[Send document comments to nexus1k-docfeedback@cisco.com.](mailto:nexus1k-docfeedback@cisco.com)**

- [Maximum Allowed VLANs and MAC Addresses per VLAN, page 6.](#)

For detailed information about Layer 2, see the *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0(4)SV1(1)*.

## No Spanning Tree Protocol

Its forwarding logic is designed to prevent network loops so the Cisco Nexus 1000V does not need to participate in Spanning Tree Protocol. Packets received from the network on any link connecting the host to the network are not forwarded back to the network by the Cisco Nexus 1000V.

## MAC Address Table

The following are limitations and restrictions for the MAC address table:

- The forwarding table for each VLAN in a VEM can store up to 1024 MAC addresses.

For detailed information about Cisco Nexus 1000V MAC address table, see the *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0(4)SV1(1)*.

## Maximum Allowed VLANs and MAC Addresses per VLAN

The following are the allowable number of VLANs and MAC addresses per VLAN that can be configured:

Feature	Maximum Limit
VLANs across all VEMs	512
MAC addresses per VLAN within a VEM	1024 (1K)

For detailed information about Cisco Nexus 1000V VLAN configuration, see the *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0(4)SV1(1)*.

## Cisco Discovery Protocol

Cisco Discovery Protocol (CDP) runs over the data link layer and is used by the Cisco Nexus 1000V to advertise information to all Cisco devices it attaches to, and, in turn, to discover and view information about those Cisco devices. CDP runs on all Cisco-manufactured equipment.

Cisco Discovery Protocol (CDP) has the following configuration guidelines and limitations:

- CDP can discover up to 256 neighbors per port if the port is connected to a hub with 256 connections.
- The CDP feature is enabled globally by default.
- If disabled globally on the Cisco Nexus 1000V, then CDP is also disabled for all interfaces.

For more information about Cisco Discovery Protocol, see the *Cisco Nexus 1000V System Management Configuration Guide, Release 4.0(4)SV1(1)*.

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

## DHCP Not Supported for the Management IP

DHCP is not supported for the management IP. The management IP must be configured statically.

## LACP

Link Aggregation Control Protocol is an IEEE standard protocol that aggregates Ethernet links into an Etherchannel.

Cisco Nexus 1000V has the following restrictions for enabling LACP on ports carrying the Control and Packet VLANs:



### Note

These restrictions do not apply to other data ports using LACP.

- At least two ports must be configured as part of the LACP channel.
- The upstream switch ports must be configured in spanning-tree portfast mode. The LACP negotiation causes upstream switchports to bounce as per protocol before starting the port aggregation process.

Without spanning-tree portfast on upstream switch ports, it takes ~30 seconds to recover these ports on the upstream switch, and since they are carrying Control and Packet VLANs, VSM loses connectivity to the VEM.

The following commands are available to use on Cisco upstream switch ports in interface configuration mode:

**spanning-tree portfast**

**spanning-tree portfast trunk**

**spanning-tree portfast edge trunk**

[Send document comments to nexus1k-docfeedback@cisco.com.](mailto:nexus1k-docfeedback@cisco.com)

## Caveats

The following are descriptions of the caveats in Cisco Nexus 1000V Release 4.0(4)SV1(1).

Bug ID	Caveat
CSCsq66077	<p><b>Headline:</b> Shutting an Ethernet interface in Cisco Nexus 1000V VSM is not reflected in the Cat6K.</p> <p><b>Symptom:</b> After shutting down an Ethernet interface for an uplink port on the VSM, the physical network interface attached to it on the VEM and the switch port attached to the physical NIC do not shut down as expected.</p> <p>The output of the following command shows the interface as shut down:</p> <p><b>show interface ethX/Y</b></p> <p>The output of the following command shows that the NIC is up:</p> <p><b>esxcfg-nics -l</b></p> <p><b>Conditions:</b> This happens on any ESX platform running Cisco NX-OS.</p> <p><b>Workaround:</b> Use the <b>shutdown</b> command on the interface of the upstream switch. This brings down both the link on the upstream switch and the ESX physical NIC.</p> <p><b>Further Problem Description:</b> The VEM sets the uplink port in the DOWN state, so that no traffic flows through that port. It is only the physical NIC attached to that DVS port which is not brought down.</p>
CSCsw32257	<p><b>Headline:</b> Shutting down a VSM VEthernet interface is not reflected in the VM.</p> <p><b>Symptom:</b> After you shut down a VEthernet port on the VSM, it does not appear to be down from the VM.</p> <p><b>Conditions:</b> The VM Guest OS, connected to the Cisco Nexus 1000V through a vEthernet port, does not see a link going down. A <b>shutdown</b> command on the VSM VEthernet interface shuts down the interface and stops traffic forwarding.</p> <p><b>Workaround:</b> You can use the GuestOS utilities to bring down the interface. In a Linux system, for example, use the <b>ifconfig down eth0</b> command.</p> <p><b>Further Problem Description:</b></p>
CSCsw49458	<p><b>Headline:</b> A change to the speed or duplex settings on a physical NIC causes module flap.</p> <p><b>Symptom:</b> The port bounces and the module flaps after changing speed or duplex settings on an Ethernet interface.</p> <p><b>Conditions:</b> The speed and duplex settings on Ethernet interfaces do not work on an interface carrying system VLANs.</p> <p><b>Workaround:</b> Avoid configuring speed or duplex settings in a VSM connected to an upstream switch.</p> <p><b>Further Problem Description:</b> None.</p>



***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

Bug ID	Caveat
CSCsx11210	<p><b>Headline:</b> Unable to add match criteria in a QoS class map after changing to <b>match-any</b>.</p> <p><b>Symptom:</b> Adding match criteria does not work after changing a class map to match-any.</p> <p><b>Conditions:</b> If you create a class map with <b>match-all packet length</b> as the only criteria, then changing the map to <b>match-any</b> prevents you from adding new match criteria.</p> <p><b>Example:</b></p> <pre>n1000v(config)# class-map c1 n1000v(config-cmap-qos)# match packet length 1028 n1000v(config-cmap-qos)# exit n1000v(config)# class-map match-any c1 n1000v(config-cmap-qos)# match packet length 1038 n1000v(config-cmap-qos)# show class-map c1 Type qos class-maps ===== class-map type qos match-any c1 match packet length 1028</pre> <p><b>Workaround:</b> Always verify a class map configuration using the <b>show class-map</b> command. If the map is not correct, delete the criteria using the <b>no match</b> command and then add it again.</p> <p><b>Example:</b></p> <pre>n1000v(config)# class-map match-any c1 n1000v(config-cmap-qos)# no match packet length 1028 n1000v(config-cmap-qos)# match packet length 1028, 1038</pre> <p><b>Further Problem Description:</b> None.</p>
CSCsx68200	<p><b>Headline:</b> Attempting to rename a port-group causes the port group to be deleted and a new port group to be created.</p> <p><b>Symptom:</b> When you use the following command to rename a port group, the existing port group is deleted and a new port group is created with the new name.</p> <pre>switch (config-port-prof)# vmware port-group new-name</pre> <p><b>Conditions:</b> Attempt to change the name of a port group.</p> <p><b>Workaround:</b> None</p> <p><b>Further Problem Description:</b> Deleting and creating a new port group may also cause the related NICs to be moved into the Quarantine port groups.</p> <p><b>Example:</b></p> <pre>2009 May 20 10:26:39 switch %VMS-3-DVPG_NICS_MOVED: '6' nics have been moved from port-group 'WebApp' to 'Unused_Or_Quarantine_Veth'</pre> <p>In this case, the NICs must again be re-associated with the port-group.</p>

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

Bug ID	Caveat
CSCsy25906	<p><b>Headline:</b> Error logged after changing the system port profiles for VMNIC with control VLAN</p> <p><b>Symptom:</b> If you change the system port profile for a physical adapter carrying the control VLAN, the following system logs are generated:</p> <pre>ETH_PORT_CHANNEL-3-COMPAT_CHECK_FAILURE PORT_PROFILE_CHANGE_VERIFY_REQ_FAILURE</pre> <p><b>Conditions:</b> Port profiles sys1 and sys2 are configured as channel-group and have the control VLAN configured as part of the system VLANs. Two physical adapters are attached to port profile sys1. From the vSphere client, change the adapters from port profile sys1 to port profile sys2 in a single step.</p> <p><b>Workaround:</b> Do not change the system port profile attached to a physical adapter in a single step. Instead, use the following procedure:</p> <ol style="list-style-type: none"> <li>1. From the vSphere client, remove the adapters from the port profile sys1.</li> <li>2. Click OK.</li> <li>3. Add the adapters to port profile sys2.</li> </ol> <p><b>Further Problem Description:</b> None</p>
CSCsy88176	<p><b>Headline:</b> Inaccurate <b>show policy-map interface</b> output if there are numerous policy-maps</p> <p><b>Symptom:</b> When displaying the statistics for multiple policy-maps, the results may not reflect accurate statistics.</p> <p><b>Conditions:</b> Multiple large policies are configured on a single VEM, especially if they use complex match-any conditions.</p> <p><b>Workaround:</b> Instead of displaying all policy maps at once, use the following command to display them one at a time: <b>show policy-map interface</b> <i>name</i></p> <p><b>Further Problem Description:</b> The syslog alerts you that too much data is being returned at one time.</p>
CSCsz03271	<p><b>Headline:</b> The same ACL cannot be used multiple times in a QoS policy-map.</p> <p><b>Symptom:</b> When an ACL policy is used more than once in a QoS policy-map, the system fails to apply it to an interface.</p> <p><b>Conditions:</b> Using the same ACL multiple times in a single policy-map.</p> <p><b>Workaround:</b> Do not use the same ACL in different class-maps that are referenced in a single QoS policy. Instead, create and reference a new ACL with same set of rules.</p> <p><b>Further Problem Description:</b> None</p>

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

Bug ID	Caveat
CSCsz15398	<p><b>Headline:</b> Performance impact when AIPC link goes down and comes back up</p> <p><b>Symptom:</b> Ports go into the errDisabled state and <b>EthPM timeout</b> system messages are generated after the control traffic link goes down and comes back up.</p> <p><b>Conditions:</b> A large number of interfaces (greater than 256 interfaces spread across 8 VEMs) are configured with ACL or QoS policies.</p> <p><b>Workaround:</b> View module states using the <b>show module</b> command. Once all modules in the system are active, do one of the following:</p> <ul style="list-style-type: none"> <li>• If the number of errDisabled interfaces is limited, enter the following command sequence: <ul style="list-style-type: none"> <li>– switch (config-if)# <b>shutdown</b></li> <li>– switch (config-if)# <b>no shutdown</b></li> </ul> </li> <li>• If there is a particular VEM (or a few VEMs) that have interfaces in the errDisabled state, force a module removal and re-insert (one by one for all affected modules). This can be done by shutting the port on the upstream switch that connects to the VEM uplink for 10 seconds.</li> <li>• If there is no difference between the switch running-configuration and startup-configuration, you can reload the VSM using the <b>reload</b> command.</li> </ul> <p><b>Further Problem Description:</b> None</p>
CSCsz21291	<p><b>Headline:</b> Port security configured in the port profile is pushing stale data (sw port-sec maximum <i>number</i>).</p> <p><b>Symptom:</b> The maximum secure MAC address count configured through the port profile is not applied to a vEthernet interface.</p> <p><b>Conditions:</b> One of the following is true for the vEthernet interface:</p> <ul style="list-style-type: none"> <li>• It is not up.</li> <li>• It does not have port security enabled.</li> </ul> <p><b>Workaround:</b> Re-configure the maximum secure address count when the interface is up and port security is enabled on the interface.</p> <p><b>Further Problem Description:</b> None</p>
CSCsz21693	<p><b>Headline:</b> Reload of iVISOR host removes the VIB package.</p> <p><b>Symptom:</b> Rebooting an ESXi host causes it to not show up in the VSM.</p> <p><b>Conditions:</b> This occurs if you load the N1KV VEM code onto the host before adding the host to vSphere Server.</p> <p><b>Workaround:</b> Do one of the following:</p> <ul style="list-style-type: none"> <li>• Add the ESXi host to vSphere before installing Cisco Nexus 1000V VEM software.</li> <li>• Reboot the host after installing Cisco Nexus 1000V VEM software but before adding the host to vSphere.</li> </ul> <p><b>Further Problem Description:</b> None</p>

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

**Bug ID****Caveat**

CSCsz24042

**Headline:** Static MAC entries for VMs are not updated upon VLAN change

**Symptom:** A private VLAN host virtual interface changed to a regular interface keeps the static MAC address in the running configuration.

**Conditions:** You configure a virtual interface as a private VLAN host and then change it to a regular interface.

**Example:**

```
switch# conf t
switch (config)# port-profile pvlan153
n1000v(config-port-prof)# vmware port-group
n1000v(config-port-prof)# switchport mode private-vlan host
n1000v(config-port-prof)# switchport private-vlan host-association 156 153
n1000v(config-port-prof)# no shutdown
n1000v(config-port-prof)# state enabled

switch# conf t
switch (config)# port-profile pvlan153
n1000v(config-port-prof)# switchport mode access
```

**Workaround:** When you change an interface from a private VLAN host interface to a regular interface, you must manually remove the static MAC from the configuration.

**Example:**

```
switch# conf t
switch (config)# no mac address-table static 0050.5692.1b66 vlan 156 interface
Vethernet1
```

**Further Problem Description:** None

CSCsz38042

**Headline:** show vlan private-vlan does not show promiscuous trunk information

**Symptom:** The show vlan private-vlan command output does not show private-vlan promiscuous trunk port information.

**Conditions:** The show vlan private-vlan command output does not list the interfaces associated with the private VLAN if it is configured as private VLAN promiscuous trunk port.

**Workaround:** Use show interface switchport or show running-config.

**Further Problem Description:** If the interface is configured as a private VLAN host port and private VLAN promiscuous access port, the show vlan private-vlan command output shows the secondary VLAN, the primary VLAN and interfaces associated with those private VLANs.

**Example:**

```
switch# show vlan private-vlan
Primary Secondary Type Ports
-----
157 community Eth3/4
152 158 isolated Eth3/4
156 153 community Eth3/4
156 154 community Eth3/4
156 155 isolated Eth3/4
```

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

Bug ID	Caveat
CSCsz48343	<p><b>Headline:</b> Link up message while powering up VM (vnic not up)</p> <p><b>Symptom:</b> As a guest operating system is booting, the vEthernet interface shows as link up for a few seconds, then down, and then it finally remains up.</p> <p><b>Conditions:</b> Powering up guest operating system.</p> <p><b>Workaround:</b> Use a flexible adapter and install VMware tools.</p> <p><b>Further Problem Description:</b> If the adapter type is flexible and you have VMware tools installed the Cisco Nexus 1000V VSM indicates the link for a vEthernet interface is up during the virtual machine boot.</p>
CSCsz63126	<p><b>Headline:</b> No switchport mode trunk - change modes to access and not to default</p> <p><b>Symptom:</b> After configuring the switchport mode on an Ethernet interface, the <b>no switchport mode command</b> leaves the switchport mode access config setting on the interface, overriding the policy inherited by the port-profile. There is no way to remove it.</p> <p><b>Conditions:</b> Configuring the switchport mode on an Ethernet interface.</p> <p><b>Workaround:</b> If the switchport mode setting has not been saved in the startup-config yet, the VSM can be reloaded to remove the setting. If it has been saved in the startup-config, there is no way to remove it.</p> <p><b>Further Problem Description:</b> None</p>
CSCsz99235	<p><b>Headline:</b> Installing a permanent license file does not add new licenses to the license pool.</p> <p><b>Symptom:</b> After installing a new license file on the VSM, the count of licenses is not increased to show that new licenses were added.</p> <p><b>Conditions:</b> Adding a permanent license file without removing an evaluation license file.</p> <p><b>Workaround:</b> Before installing a new license file, you must first transfer the evaluation licenses from the VEMs back to the VSM license pool and then uninstall the evaluation license from the VSM.</p> <p><b>Further Problem Description:</b> None</p>
CSCta05268	<p><b>Headline:</b> Modules do not come up for a VSM with a VEM port channel running in vPC-HM.</p> <p><b>Symptom:</b> The output of the <b>show l2 &lt;control VLAN number&gt;</b> command shows <i>dynamic</i> for the VM Eth0 MAC.</p> <p><b>Conditions:</b> After the VSM connects to the VEM, for example, when the VSM is reloaded.</p> <p><b>Workaround:</b> Migrate the VSM VM to a vSwitch or another host without vPC-HM.</p> <p><b>Further Problem Description:</b> When a VSM and VEM reconnect, the L2 table entries and port channel are deleted and the physical links carry the same VLANs. This causes broadcast packets from the VSM to go out through one upstream switch and come back through another. Therefore, the Eth0 MAC of the VSM VM is learned on a physical interface and the module never comes up.</p>

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

Bug ID	Caveat
CSCte28866	<p><b>Headline:</b> Configuring a Cisco Nexus 1000V with the <b>vlan dot1Q tag native</b> command does not result in the desired behavior.</p> <p><b>Symptom:</b> The traffic on the native VLAN is not tagged when sent across a trunk.</p> <p><b>Conditions:</b> Configuring a Cisco Nexus 1000V with the <b>vlan dot1Q tag native</b> command.</p> <p><b>Workaround:</b> There is currently no workaround. Disabling the native VLAN tagging on the upstream network infrastructure could alleviate the need to use the <b>vlan dot1Q tag native</b> command on the Cisco Nexus 1000V.</p>

## MIB Support

The Cisco Management Information Base (MIB) list includes Cisco proprietary MIBs and many other Internet Engineering Task Force (IETF) standard MIBs. These standard MIBs are defined in Requests for Comments (RFCs). To find specific MIB information, you must examine the Cisco proprietary MIB structure and related IETF-standard MIBs supported by the Cisco Nexus 1000V Series switch.

The MIB Support List is available at the following FTP site:

<ftp://ftp.cisco.com/pub/mibs/supportlists/nexus1000v/Nexus1000VMIBSupportList.html>

## Related Documentation

Cisco Nexus 1000V includes the following documents available on [Cisco.com](http://www.cisco.com):

### General Information

*Cisco Nexus 1000V Release Notes, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V and VMware Compatibility Information, Release 4.0(4)SV1(1)*

### Install and Upgrade

*Cisco Nexus 1000V Software Installation Guide, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V Virtual Ethernet Module Software Installation Guide, Release 4.0(4)SV1(1)*

### Configuration Guides

*Cisco Nexus 1000V License Configuration Guide, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V Interface Configuration Guide, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V Quality of Service Configuration Guide, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V Security Configuration Guide, Release 4.0(4)SV1(1)*

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

*Cisco Nexus 1000V System Management Configuration Guide, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V High Availability and Redundancy Reference, Release 4.0(4)SV1(1)*

## Reference Guides

*Cisco Nexus 1000V Command Reference, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V MIB Quick Reference*

## Troubleshooting and Alerts

*Cisco Nexus 1000V Troubleshooting Guide, Release 4.0(4)SV1(1)*

*Cisco Nexus 1000V Password Recovery Guide*

*Cisco NX-OS System Messages Reference*

---

This document is to be used in conjunction with the documents listed in the “[Related Documentation](#)” section.

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](http://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)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2010 Cisco Systems, Inc. All rights reserved.

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***