

Send documentation comments to nexus5kdocs@cisco.com



Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Release Notes, Release 4.0(1a)N2(1)

Current Release: 4.0(1a)N2(1) -February 26, 2009
Part Number: OL-16601-01 H0

This document describes the features, caveats, and limitations for Cisco Nexus 5000 Series switches and the Cisco Nexus 2000 Series Fabric Extenders. Use this document in combination with documents listed in the “[Related Documentation](#)” section on page 28.



Note

Release notes are sometimes updated with new information about restrictions and caveats. See the following website for the most recent version of the Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Release Notes:
http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/release/notes/Nexus_5000_Release_Notes.html



Note

[Table 1](#) shows the online change history for this document.

Table 1 **Online History Change**

Part Number	Revision	Date	Description
OL-16601-01	A0	June 03, 2008	Created release notes.
OL-16601-01	B0	June 16, 2008	Added information for Release 4.0(0)N1(1a).
OL-16601-01	C0	June 30, 2008	Added information for Cisco Fabric Manager Release 3.4(1a).
OL-16601-01	D0	July 22, 2008	Added information for Release 4.0(0)N1(1a).
OL-16601-01	E0	August 13, 2008	Added information for Release 4.0(0)N1(2).
OL-16601-01	F0	September 29, 2008	Added information for Release 4.0(0)N1(2a).



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

Send documentation comments to nexus5kdocs@cisco.com

Table 1 Online History Change

Part Number	Revision	Date	Description
OL-16601-01	G0	December 03, 2008	Added information for Release 4.0(1a)N1(1).
OL-16601-01	H0	February 26, 2009	Added information for Release 4.0(1a)N2(1).

Contents

This document includes the following sections:

- [Introduction, page 2](#)
- [New and Changed Features in Cisco NX-OS Release 4.0\(1a\)N2\(1\), page 5](#)
- [Cisco NX-OS Release 4.0\(0\)N-Based Releases Upgrade/Downgrade Issues, page 5](#)
- [Changes to the FCoE Model and Related Configuration, page 8](#)
- [Limitations, page 11](#)
- [Caveats, page 14](#)
- [Cisco Fabric Manager, page 27](#)
- [Related Documentation, page 28](#)
- [Obtaining Documentation and Submitting a Service Request, page 29](#)

Introduction

This section includes the following topics:

- [Cisco Nexus 5000 Series Switches, page 2](#)
- [Cisco Nexus 2000 Series Fabric Extenders, page 3](#)

Cisco Nexus 5000 Series Switches

The Cisco Nexus 5000 Series switches comprise a family of line-rate, low-latency, lossless 10-Gigabit Ethernet, Cisco Data Center Ethernet, and Fibre Channel over Ethernet (FCoE) switches for data center applications. The Cisco Nexus 5000 Series includes the Cisco Nexus 5020 switch and the Cisco Nexus 5010 switch.

The Cisco Nexus 5000 Series switch hardware is described in the following topics:

- [Cisco Nexus 5020 Switch, page 3](#)
- [Cisco Nexus 5010 Switch, page 3](#)

Send documentation comments to nexus5kdocs@cisco.com

Cisco Nexus 5020 Switch

The Cisco Nexus 5020 is a 56-port switch. It is a two rack unit (2 RU), 10-Gigabit Ethernet, Cisco Data Center Ethernet, FCoE, and Fibre Channel switch that provides 1.04 terabits per second (Tbps) throughput with very low latency.

It has the following features:

- Forty fixed 10-Gigabit Ethernet, Cisco Data Center Ethernet, and FCoE Small Form Factor Pluggable Plus (SFP+) ports. Sixteen of the forty fixed ports support both Gigabit Ethernet and 10-Gigabit Ethernet. The default is 10-Gigabit Ethernet.
- Two expansion module slots that can be configured to support up to 12 additional 10-Gigabit Ethernet, Cisco Data Center Ethernet, and FCoE SFP+ ports, up to 16 Fibre Channel switch ports, or a combination of both.
- Serial console port and an out-of-band 10/100/1000-Mbps Ethernet management port.
- 1+1 redundant, hot-pluggable power supplies.
- 4+1 redundant, hot-pluggable fan modules to provide highly reliable front-to-back cooling.

For additional information about the Cisco Nexus 5020 switch, see the *Cisco Nexus 5000 Series Hardware Installation Guide*.

Cisco Nexus 5010 Switch

The Cisco Nexus 5010 is a 28-port switch. It is a 1 RU, 10-Gigabit Ethernet, Cisco Data Center Ethernet, FCoE, and Fibre Channel switch that provides more than 500-Gbps throughput with very low latency. It has the following features:

- Twenty fixed 10-Gigabit Ethernet, Cisco Data Center Ethernet, and FCoE SFP+ ports. Eight of the twenty fixed ports support Gigabit Ethernet and 10-Gigabit Ethernet speed.
- One expansion module slot that can be configured to support up to 6 additional 10-Gigabit Ethernet, Cisco Data Center Ethernet, and FCoE SFP+ ports, up to 8 Fibre Channel switch ports, or a combination of 4 additional 10-Gigabit Ethernet, Cisco Data Center Ethernet, and FCoE SFP+ ports with 4 additional Fibre Channel switch ports.
- Serial console port and an out-of-band 10/100/1000-Mbps Ethernet management port.
- 1+1 redundant, hot-pluggable power supplies.
- 1+1 redundant, hot-pluggable fan modules to provide highly reliable front-to-back cooling.

For additional information about the Cisco Nexus 5010 switch, see the *Cisco Nexus 5000 Series Hardware Installation Guide*.



Note

The Cisco Nexus 5020 switch and the N5K-M1404 and N5K-M1600 gatos expansion modules (GEMs) use a release 4.0(0)N1(1) or later image. The Cisco 5010 switch and the N5K-M1008 GEM use a release 4.0(1a)N1(1) or later image.

Cisco Nexus 2000 Series Fabric Extenders

The Cisco Nexus 2000 Series Fabric Extender is a highly scalable and flexible server networking solution that works with the Cisco Nexus 5000 Series switches to provide high-density and low-cost connectivity for server aggregation.

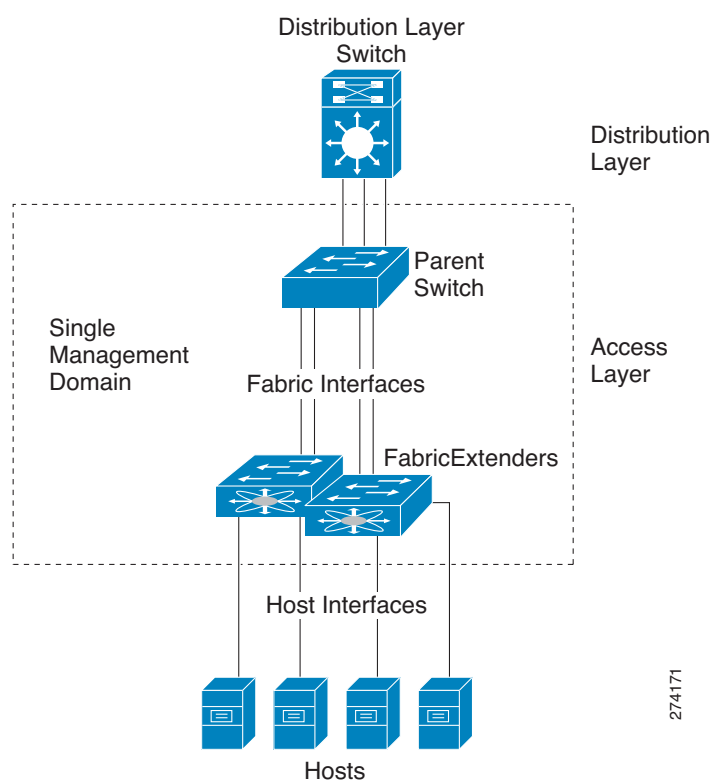
Send documentation comments to nexus5kdocs@cisco.com

Scaling across a multitude of 1-Gigabit Ethernet, 10-Gigabit Ethernet, unified fabric, rack, and blade server environments, the Fabric Extender is designed to simplify data center architecture and operations.

The Fabric Extender integrates with its parent switch, allowing zero-touch provisioning as well as automatic configuration. This integration allows large numbers of servers and hosts to be supported using the same feature set as the parent Nexus 5000 Series switch, including security and quality of service (QoS) configuration parameters, with a single point of management as shown in [Figure 1-1](#). Spanning Tree Protocol (STP) is not required between the Fabric Extender and its parent switch, because the Fabric Extender and its parent switch allow you to enable a large multi-path, loop-free, active-active topology.

Since the Fabric Extender is designed to connect to servers directly, by default, all Fabric Extender host ports are edge ports. In addition, BPDU guard and BPDU filters are also enabled on Fabric Extender host ports by default.

Figure 1-1 Single Management Domain



This section describes the 2148T Fabric Extender. It includes the following topic:

- [Cisco Nexus 2148T Fabric Extender, page 4](#)

Cisco Nexus 2148T Fabric Extender

The first product in the Cisco Nexus 2000 Series is the Nexus 2148T Fabric Extender, a 1 RU chassis designed for rack mounting. The chassis supports redundant hot-swappable fans and power supplies.

Send documentation comments to nexus5kdocs@cisco.com

The Cisco Nexus 2148T Fabric Extender forwards all traffic to a parent Cisco Nexus 5000 Series switch over 10-Gigabit Ethernet fabric uplinks, allowing all traffic to be inspected by policies established on the Cisco Nexus 5000 Series switch. No software is included with the Nexus 2148T. Software is downloaded and upgraded from its parent Cisco Nexus 5000 Series switch.

The Nexus 2148T has 48 1-Gigabit Ethernet host interfaces for its downlink connection to servers or hosts and 4 10-Gigabit Ethernet fabric interfaces with SFP+ interface adapters for its uplink connection to the parent switch.

New and Changed Features in Cisco NX-OS Release 4.0(1a)N2(1)

This release includes the following new or changed features:

- Support for the Cisco Nexus 2000 Series Fabric Extender (part number N2K-C2148T-1GE)
- Port based CoS assignment
- STP bridge assurance

For more information about the features listed, see the Cisco Nexus 5000 Series and the Cisco Nexus 2000 Series documentation listed in the [“Related Documentation”](#) section on page 28.

Cisco NX-OS Release 4.0(0)N-Based Releases Upgrade/Downgrade Issues

This section describes issues you may encounter when you upgrade from or downgrade to Cisco NXOS 4.0(0)N based releases on the Cisco Nexus 5000 Series switch. It provides examples of configuration syntax differences for some of these changes.

This section includes the following topics:

- [EtherChannel Upgrade/Downgrade Changes, page 5](#)
- [Fibre Channel Port Shutdown, page 7](#)
- [Switched Port Analyzer, page 8](#)
- [Example of Virtual Interface Configuration Changes, page 10](#)
- [Upgrading from Cisco NX-OS 4.0\(0\)N-Based Releases, page 11](#)
- [Downgrading to Cisco NX-OS Release 4.0\(0\)N-Based Releases, page 11](#)

EtherChannel Upgrade/Downgrade Changes

The following table describes the changes to the security access control lists (ACLs) for EtherChannel members.

Release	Description
4.0(0)N-based releases	Configuration is allowed on member ports (but not used while the port is a member of the EtherChannel).
4.0(1a)N-based releases and later releases	No configuration is allowed on member ports. All member ports follow the configuration on the EtherChannel.

Send documentation comments to nexus5kdocs@cisco.com

Upgrade	Member port ACL configuration, if any, is lost; EtherChannel configuration is preserved. No impact on the functional behavior while the port is a member of the EtherChannel. When the member port leaves the EtherChannel, you have to recreate the ACL configuration on the physical interface.
Downgrade	No issue.

The following table describes the changes to the interface-level QoS service policy.

Release	Description
4.0(0)N-based releases	Interface-level QoS service policy is not aware of the EtherChannel. QoS strictly follows a per-member port configuration. The only service policy supported on the interface level is egress queue scheduling.
4.0(1a)N-based releases and later releases	No configuration is allowed on member ports. All member ports follow the configuration on the EtherChannel.
Upgrade	If any member port has a modified egress scheduling policy and the EtherChannel is explicitly configured, then the egress scheduling configuration for the port is lost. A EtherChannel has a default egress scheduling policy.
Downgrade	This capability needs to be reconfigured after the downgrade. Otherwise, egress queue scheduling configuration on the EtherChannel will be lost after the downgrade.

The following table describes the changes to the priority flow-control configuration.

Release	Description
4.0(0)N-based releases	Priority flow-control is an interface-level configuration. This CLI option is used to override results of DCBX negotiation. Priority flow-control function is not aware of the EtherChannel. It strictly follows a per-member port configuration.
4.0(1a)N-based releases and later releases	No configuration allowed on member ports. All member ports follow the configuration on the EtherChannel.
Upgrade	If a member port has modified priority flow-control configuration and the EtherChannel is explicitly configured, then port configuration is lost. The EtherChannel has default priority flow-control configuration.
Downgrade	This capability needs to be reconfigured after the downgrade. Otherwise, the priority flow-control configuration on the EtherChannel will be lost.

The following table describes the changes to the syntax of the Ethernet load-balancing commands.

Release	Description
4.0(0)N-based releases	The following command is used to set the load-balancing method in a channel-group bundle: <pre>switch(config)# port-channel load-balance ethernet source-destination-? source-destination-ip Source & Destination IP address source-destination-mac Source & Destination MAC address source-destination-port Source & Destination TCP/UDP port</pre> <p>Note The keyword source-destination is used in this release.</p>

Send documentation comments to nexus5kdocs@cisco.com

4.0(1a)N-based releases and later releases	<p>The following command is used to set the load-balancing method in a channel-group bundle:</p> <pre>switch(config)# port-channel load-balance ethernet source-dest-? source-dest-ip Source & Destination IP address source-dest-mac Source & Destination MAC address source-dest-port Source & Destination TCP/UDP port</pre> <p>Note The keyword source-dest is used in this release.</p>
Upgrade	The load-balancing configuration is lost.
Downgrade	This capability needs to be reconfigured after the downgrade. Otherwise, the load-balancing configuration on the EtherChannel will be lost.

When an Ethernet interface joins an EtherChannel, the following interface-level parameters are disabled:

bandwidth	Set bandwidth informational parameter
delay	Specify interface throughput delay
duplex	Enter the port duplex mode
flowcontrol	Configure interface flowcontrol
ip	Configure IP features
ipv6	Configure IPv6 features
mac	MAC configuration commands
priority-flow-control	Configure interface priority-flowcontrol
service-policy	Configure QoS service policy
spanning-tree	Spanning Tree Subsystem
speed	Enter the port speed
storm-control	Configure Interface storm control

Fibre Channel Port Shutdown

The following table describes the changes to the interface **shutdown** command syntax.

Release	Description
4.0(0)N-based releases	The system default switchport shutdown command causes all Fibre Channel ports, virtual or physical, to default to shut down.
4.0(1a)N-based releases and later releases	The system default switchport shutdown command now configures the administrative state for all Ethernet ports as down. To configure the Fibre Channel ports to the shutdown state, use the system default switchport shutdown san command.
Upgrade	Ethernet ports will default to the shutdown state instead of Fibre Channel ports.
Downgrade	This capability needs to be reconfigured after the downgrade. Otherwise, the switch will flag the system default switchport shutdown san command as invalid.

[Send documentation comments to nexus5kdocs@cisco.com](mailto:nexus5kdocs@cisco.com)

Switched Port Analyzer

The following table describes the changes to Switched Port Analyzer (SPAN) sessions.

Release	Description
4.0(0)N-based releases	The default is to keep a session in an open state. To shut a session, use the following command: <code>switch(config)# monitor session session-number suspend</code>
4.0(1a)N-based releases and later releases	The default is to keep a session shut state. To open a session, use the following command: <code>switch(config)# no monitor session session-number shut</code> Note The syntax of the command has also changed. The suspend keyword has been changed to shut .
Upgrade	SPAN session will be in the shut state after the upgrade.
Downgrade	If you do not execute a specific shut or no-shut command on the SPAN session, the SPAN session will be in a no suspend state after the downgrade.

Changes to the FCoE Model and Related Configuration

In the previous Cisco NX-OS 4.0(0)N-based releases, the FCoE model allowed Ethernet and FCoE to coexist on the interface. The virtual interfaces, virtual Ethernet, and virtual Fibre Channel (VFC), did not affect each other. For example, if the virtual Ethernet interface was errdisabled, the VFC interface could still be up.

With Cisco NX-OS 4.0(1a)N-based releases and later releases, the CLI implementation was changed to provide forward compatibility with the forthcoming T11 FCoE Initialization Protocol (FIP). In this new FCoE mode, the following is changed:

- FCoE traffic is passed over Ethernet; the Ethernet STP controls port status.
- FCoE traffic for a VSAN is transported over a single dedicated FCoE-enabled VLAN.

Although the CLI implementation requires a VLAN-to-VSAN mapping, the FCoE frames are expected to be untagged or priority-tagged. The FCoE VLAN is not expected to be carried in the FCoE frames. This will change in a future release once FIP-based FCoE is supported.

The changes to the virtual interfaces are described in the following topics:

- [Virtual Interface Groups, page 9](#)
- [Virtual Ethernet Interfaces, page 9](#)
- [Virtual Fibre Channel Interfaces, page 9](#)
- [VSAN-to-VLAN Mapping, page 9](#)

Send documentation comments to nexus5kdocs@cisco.com

Virtual Interface Groups

In previous Cisco NX-OS 4.0(0)N-based releases, a virtual interface group allowed you to bind virtual interfaces to a physical Ethernet interface, as shown in the following example:

```
switch# configure terminal
switch(config)# interface vig 1
switch(config-if)# bind interface ethernet 1/1
```

Virtual interface groups have been deprecated in Cisco NX-OS 4.0(1a)N-based releases and later releases.

Virtual Ethernet Interfaces

Cisco NX-OS 4.0(1a)N-based releases and later releases do not support virtual Ethernet interfaces. All Ethernet features previously configured at the virtual Ethernet interface now need to be configured at the bound Ethernet interface.

The following configuration statements must be explicitly configured on the Ethernet interface to keep the behavior the same as on a virtual Ethernet interface:

```
spanning-tree bpduguard enable
spanning-tree port type edge trunk
```

All other features (for example, ACLs, SPAN, and so forth.) that were previously configured at the virtual Ethernet interface would need to be applied to the bound Ethernet interface.

Virtual Fibre Channel Interfaces

In previous Cisco NX-OS 4.0(0)N-based releases, a virtual Fibre Channel interface was attached to a virtual interface group, which then bound it to the physical Ethernet interface.

Each virtual Fibre Channel interface is bound directly to an FCoE-enabled physical Ethernet interface in Cisco NX-OS 4.0(1a)N based releases and later releases. This change simplifies the **interface vfc** command as shown in the following example:

```
switch# configure terminal
switch(config)# interface vfc 1
switch(config-if)# bind interface ethernet 1/1
```

The Ethernet interface that you bind the virtual Fibre Channel interface to must be configured as follows:

- It must be a trunk port (use the **switchport mode trunk** command).
- The FCoE VLAN that corresponds to the virtual Fibre Channel's VSAN must be in the allowed The FCoE VLAN must not be configured as the native VLAN of the trunk port.
- The Ethernet interface must be configured as PortFast (use the **spanning-tree port type edge trunk** command).

VSAN-to-VLAN Mapping

In previous Cisco NX-OS 4.0(0)N-based releases, FCoE had no dependence on the VLANs defined on the switch.

Send documentation comments to nexus5kdocs@cisco.com

In Cisco NX-OS 4.0(1a)N based releases and later releases, each virtual Fibre Channel interface is associated with only one VSAN. Any VSAN with associated virtual Fibre Channel interfaces must be mapped to a dedicated FCOE-enabled VLAN as shown in the following example:

```
switch# configure terminal
switch(config)# vlan 200
switch(config-vlan)# fcoe vsan 2
switch(config-vlan)# exit
switch(config)# interface ethernet 1/1
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk native vlan 1
switch(config-if)# switchport trunk allowed vlan 1,200
switch(config-if)# exit
switch(config)# interface vfc 1
switch(config-if)# bind interface ethernet 1/1
switch(config-if)# exit
switch(config)# vsan database
switch(config-vsan)# vsan 2 interface vfc 1
```



Note

The FCoE VLAN must be exclusively reserved for FCoE traffic; it must not be used for non-FCoE Ethernet forwarding.

Example of Virtual Interface Configuration Changes

This section includes the following topics:

- [Cisco NX-OS 4.0\(0\)N-Based Release Configuration, page 10](#)
- [Cisco NX-OS 4.0\(1a\)N-Based Releases and Later Releases Converted Configuration, page 10](#)

Cisco NX-OS 4.0(0)N-Based Release Configuration

```
interface vig 1
bind interface Ethernet 1/1

interface vethernet 1/1
switchport access vlan 2

interface vfc 1/1
no shutdown

vsan database
vsan 1 interface vfc 1/1
```

Cisco NX-OS 4.0(1a)N-Based Releases and Later Releases Converted Configuration

```
vlan 101
fcoe vsan 1

interface Ethernet 1/1
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 2,101
spanning-tree bpduguard enable
spanning-tree port type edge trunk
```

Send documentation comments to nexus5kdocs@cisco.com

```
interface vfc 1
no shutdown

vsan database
vsan 1 interface vfc 1
```

Upgrading from Cisco NX-OS 4.0(0)N-Based Releases

When you upgrade your Cisco Nexus 5000 Series switch from Cisco NX-OS 4.0(0)N based releases, all virtual Fibre Channel and virtual Ethernet interface configuration will be lost because the applicable CLI is incompatible with the previous releases. We recommend that you backup your startup-config file prior to performing the upgrade. If your switch has an FCoE configuration you will need to reconfigure FCoE using the new FCoE CLI. Alternatively, you can contact Cisco Customer Support for help with upgrading and converting your configuration to the new format.

Downgrading to Cisco NX-OS Release 4.0(0)N-Based Releases

Your FCoE configuration will be lost when you downgrade to Cisco NX-OS Release 4.0(0)N-based Releases. We recommend that you back up your original configuration for Cisco NX-OS 4.0(0)N-based releases. The backed up startup-config can be used to restore your original configuration after a downgrade.

Downgrade Steps for FCoE Configurations

To restore your FCoE configuration after a downgrade to Cisco NX-OS 4.0(0)N-based releases, follow these steps:

- | | |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Downgrade the switch to your Cisco NX-OS 4.0(0)N-based image by using the install all command. |
| Step 2 | After the downgrade completes, clear the switch configuration by using the write erase command and then the reload command. |
| Step 3 | Reconfigure the switch by running the feature fcoe and copy running startup commands. For a switch that runs in NPV mode use the npv enable command. Reload the switch by using the reload command. |
| Step 4 | Restore your pre-upgrade configuration using the copy 4.0(0)N based image-config running-config command. |

Limitations

This section describes the limitations in Nexus 5000 Series switches and the Cisco Nexus 2000 Series Fabric Extenders, Release 4.0(1a)N2(1).

- In large scale configurations some Cisco Nexus 2000 Series Fabric Extenders may take up to 3 minutes to appear online after a **reload** command is issued. A configuration can be termed large scale when the maximum permissible Cisco Nexus 2000 Series Fabric Extenders are connected to a Cisco Nexus 5000 Series switch, all host facing ports are connected and each host facing interface has large configuration (supporting the maximum permissible ACEs per interface).

Send documentation comments to nexus5kdocs@cisco.com

- The Cisco Nexus 2000 Fabric Extender does not support PVLANs over VLAN trunks used to connect to another switch. The PVLAN trunks are only used on inter-switch links but the FEX ports are only meant to connect to servers. Since it is not a valid configuration to have an isolated secondary VLAN as part of a Fabric Extender port configured as a VLAN trunk, all frames on isolated secondary VLANs are pruned from going out to a FEX.
- Egress scheduling is not supported across drop/no-drop class. Each Fabric Extender host port does not support simultaneous drop and no drop traffic. Each Fabric Extender host port can support drop or no drop traffic.
- Traffic going out the Ethernet SPAN destination is always tagged. The SPAN destination can be in the access or trunk mode and frames on the SPAN source port can be tagged or untagged. Frames are always tagged internally as they travel through the system. Information about whether the frame was originally tagged or untagged, as it appeared in the SPAN source, is not preserved in the SPAN destination. The spanned traffic exiting the SPAN destination port always has the VLAN tag on it. The correct VLAN tag is applied on the frame as it goes out the SPAN destination. The only exception is if frames ingress on a SPAN source port on an invalid VLAN. In this case, **vlan 0** will be applied on a spanned frame.
- RADIUS and AAA startup configuration is lost when you upgrade from 4.0(0)N based releases to 4.0(1a)N based releases and later releases or when you downgrade to 4.0(0)N based releases. Save the startup configuration to bootflash memory before an upgrade or a downgrade, and restore it from bootflash memory after the upgrade or downgrade.
- Spanned Fibre Channel over Ethernet (FCoE) frames do not preserve original SMAC and DMAC fields. The Ethernet header gets modified as the frame is spanned to the destination. The modified header fields are displayed when monitored on the SPAN destination.
- The CoS value in spanned Fibre Channel over Ethernet (FCoE) frames on the Ethernet SPAN destination port does not match with the CoS value in the SPAN FCoE source frame. The CoS value on the captured SPAN FCoE frame should be ignored.
- Ethernet and Fibre Channel frames with bad CRC or invalid SOF and EOF are not dropped. The Cisco Nexus 5000 Series switch operates in cut-through switching mode, so frames are forwarded through the system before they are completely received. The Ethernet and Fibre Channel CRCs are overwritten in the frame and the EOF code is set to EOFa. A downstream switch or the destination end station will drop the bad frames.
- The class-fcoe cannot be removed even if Fibre Channel is not enabled on a switch.

Send documentation comments to nexus5kdocs@cisco.com

- VACLs of more than one type on a single VLAN are unsupported. NX-OS software supports only a single type of VACL (either MAC, IPv4, or IPv6) applied on a VLAN. When a VACL is applied to a VLAN, it replaces the existing VACL if the new VACL is a different type. For instance, if a MAC VACL is configured on a VLAN and then an IPv6 VACL is configured on the same VLAN, the IPv6 VACL gets applied and the MAC VACL is removed.
- A MAC ACL is applied only on non-IP packets. Even if there is a **match eth type = ipv4** statement in the MAC ACL, it does not match an IP packet. To overcome this situation, use IP ACLs to apply access control to IP traffic instead of using a MAC ACL that matches the Ethertype to Ipv4 or Ipv6.
- If a port drains traffic at a rate less than 100 Kbps, it is errdisabled in 10 seconds to avoid buffer exhaustion. However, if the drain rate is larger than 100 Kbps, the port may not be consistently errdisabled within 10 seconds. This could cause ingress buffers to be exhausted leading to frames being discarded. Use the **shut** command to disable the slow-draining port.
- The multicast storm control functionality in the Cisco Nexus 5000 Series hardware does not distinguish between IP, non-IP, registered, or unregistered multicast traffic. All multicast traffic is subject to a single multicast storm control policer when configured.
- Multiple **boot kickstart** statements in the configuration are not supported.
- If you remove an expansion module with Fibre Channel ports, and the cable is still attached, the following FCP_ERRFCP_PORT errors are displayed:

```
2008 May 14 15:55:43 switch %KERN-3-SYSTEM_MSG: FCP_ERRFCP_PORT:
gat_fcp_isr_ip_fcmac_sync_intr@424, jiffies = 0x7add9a:Unknown intr src_id 42 - kernel
2008 May 14 15:55:43 switch %KERN-3-SYSTEM_MSG: FCP_ERRFCP_PORT:
gat_fcp_isr_ip_fcmac_sync_intr@424, jiffies = 0x7add9a:Unknown intr src_id 41 - kernel
```

These messages are informational only, and result in no loss of functionality.

SPAN Limitations on Fabric Extender Ports

- Ports on a Fabric Extender (FEX) can be configured as a tx-source in one session only.

If two ports on the same FEX are enabled to be tx-source, the ports need to be in the same session. If you configure a FEX port as a tx-source and another port belonging to the same FEX is already configured as a tx-source on a different SPAN session, then an error is displayed on the CLI.

In the following example, Interface Ethernet100/1/1 on a FEX 100 is already configured as a tx-source on SPAN session-1:

```
swor28(config-monitor)# show running-config monitor
version 4.0(1a)N2(1)
monitor session 1
source interface Ethernet100/1/1 tx
destination interface Ethernet1/37
no shut
```

If you add an interface Ethernet100/1/2 as a tx-source to a different SPAN session (session-2) the following error is displayed:

```
swor28(config)# monitor session 2
swor28(config-monitor)# source interface ethernet 100/1/2 tx
ERROR: Eth100/1/2: Ports on a fex can be tx source in one session only
swor28(config-monitor)#
```

- When a FEX port is configured as a tx-source, the multicast traffic on all VLANs for which the tx-source port is a member, will be SPAN-ed. The FEX port sends out only multicast packets that are not filtered by IGMP snooping. For example, if FEX ports 100/1/1-12 are configured on VLAN

Send documentation comments to nexus5kdocs@cisco.com

11 and the switch port 1/5 sends multicast traffic on VLAN 11 in a multicast group, and hosts connected to FEX ports 100/1/3-12 are interested in receiving that multicast traffic (through IGMP), then that multicast traffic goes out on FEX ports 100/1/3-12, but not on 100/1/1-2.

If you configure SPAN Tx on port 100/1/1, although the multicast traffic does not egress out of port 100/1/1, the SPAN destination does receive that multicast traffic. This is due to a design limitation.

- When a FEX port is configured as both SPAN rx-source and tx-source, the broadcast, non-IGMP layer-2 multicast, and unknown unicast frames originating from that port may be seen twice on the SPAN destination, once on the ingress and once on the egress path. On the egress path, the frames are filtered by the FEX to prevent them from going out on the same port on which they were received. For example, if FEX Port 100/1/1 is configured on VLAN 11, and is also configured as SPAN rx-source and tx-source and a broadcast frame is received on that port, the SPAN destination recognizes two copies of the frame, even though the frame is not sent back on port 100/1/1.
- A FEX port can not be configured as a SPAN destination. Only a switch port can be configured and used as a SPAN destination.

Caveats

This section includes the following topics:

- [Open Caveats, page 14](#)
- [Resolved Caveats—Cisco NX-OS Release 4.0\(1a\)N2\(1\), page 21](#)
- [Resolved Caveats—Cisco NX-OS Release 4.0\(1a\)N1\(1\), page 22](#)
- [Resolved Caveats—Cisco NX-OS Release 4.0\(0\)N1\(2a\), page 24](#)
- [Resolved Caveats—Cisco NX-OS Release 4.0\(0\)N1\(2\), page 25](#)
- [Resolved Caveats—Cisco NX-OS Release 4.0\(0\)N1\(1a\), page 26](#)
- [Resolved Caveats—Cisco NX-OS Release 4.0\(0\)N1\(1\), page 27](#)

Open Caveats

This section lists the open caveats for this release.

- CSCsx67695

Symptom: When a Gatos expansion module (GEM) port has an interface policy configuration that is consistent with the system policy, then when the GEM is removed from the switch, the system policy is changed to a new one. When the GEM is inserted into the switch, there can be an inconsistency between the GEM port interface policy and the new system policy (the new system policy may not have a system class that was in the old system policy). No error is reported for this inconsistency.

Workaround: Remove the GEM port interface policy, fix the interface policy, then reapply to the GEM port.

- CSCsx68778

Symptom: You cannot configure commands under the interface range

Workaround: Configure command under HIF ports of one FEX at a time.

Send documentation comments to nexus5kdocs@cisco.com

- CSCsx60187

Symptom: Duplicate IP addresses are configured for multiple SVI interfaces.

Workaround: Do not configure two more SVIs with the same IP address.

- CSCsx35870

Symptom: The CLI times out when large number of VLANs are created and deleted followed by PVLAN creation and deletion. The system will syslog indicating that Ethernet Port Manager (ethpm) timing out to communicate with SPAN manager or PVLAN manager. As a result, some of the PVLAN interfaces will be error disabled. FEX ports are configured as a member of PVLAN and some of them are member of regular VLAN.

Workaround: Perform **shut** and **no shut** on the error disabled interfaces.

- CSCsx39481

Title: After **vlan intf delete**, its ipv6 address cannot be added to a diff.

Symptom: After deleting an up vlan interface which contains an ipv6 address, you cannot add the same ipv6 address to another vlan interface. This is only seen with ipv6 addresses. IPV4 addresses do not have this problem.

Workaround: Bring back the deleted VLAN interface, delete the IPv6 address and then delete the VLAN interface again. Followed by the addition of the same IPv6 address on a different VLAN interface.

- CSCsx54086

Symptom: If source-vlan is configured for a monitor session and a downgrade is done from 4.0(1a)N2(1), the source-vlan configuration will be lost. The impact is the src-vlan traffic may not get spanned as desired after downgrade.

Workaround: After downgrade, reconfigure the source-vlan configuration for a monitor session

- CSCsx54270

Symptom: When upgrading to 4.0(1a)N2(1), if all cos values map to no-drop classes in the **system qos service-policy** configuration, the service-policy application will fail. The impact is the traffic may not conform to the expected behavior after upgrade.

Workaround: After upgrade, reconfigure the policymap to include atleast one cos in a non no-drop class and reapply policymap to **system qos service-policy** configuration.

- CSCsx40562

Symptom: ACL drop traffic with 802.1p cos values greater than 3 may not get spanned if all four user qos classes are not configured in **system qos service-policy** configuration.

Workaround: Configure all four user qos classes (except **class-default** and **class-fcoe**) under **system qos service-policy** to span all ACL drop traffic.

- CSCsw21301

Symptom: Cisco NX-OS does not provide offline configuration support. The creation of a FEX interface depends on the FEX coming online after a fabric interface configuration. When you restore configuration from a file, there periodwhen FEX interfaces are not yet created but interface configuration is applied and fails.

Workaround: If system configuration is to be restored from a configuration file (copied locally or through tftp), you can separate the FEX interface part of the configuration (if any) into a different file. Copy the main file first, then wait for FEX to come online, and then copy the separate FEX interface configuration file. Alternately, you can copy twice.

Send documentation comments to nexus5kdocs@cisco.com

- CSCsv52871

Symptom: When multiple FEX host ports congest FEX uplinks with no-drop class of traffic, pause asserted on the FEX host ports is uneven. As a result, traffic from some FEX host ports observe a better throughput than others.

Workaround: None
- CSCsx24526

Symptom: When a FEX host port is congested, an uncongested port in the same block of 8 ports will experience some traffic loss. The loss varies based on how many sources are trying to congest the congested port.

Workaround: None
- CSCsv93263

Symptom: Cisco NX-OS does not provide offline configuration support. The creation of a FEX interface depends on the FEX coming online after a fabric interface configuration. When you restore configuration from a file, there periodwhen FEX interfaces are not yet created but interface configuration is applied and fails.

Workaround: If system configuration is to be restored from a configuration file (copied locally or through tftp), you can separate the FEX interface part of the configuration (if any) into a different file. Copy the main file first, then wait for FEX to come online, and then copy the separate FEX interface configuration file. Alternately, you can copy twice.
- CSCsv81694

Symptom: The auto learn static MAC entry is removed if the port on which the same MAC address is dynamically learned is flapped. The static MAC address is removed from the software as well as the hardware.

Workaround: Re-add the static MAC entry through the CLI.
- CSCsu01188

Symptom: No traps are sent when SFPs for Gigabit Ethernet and 10-Gigabit Ethernet are removed or inserted.

Workaround: None.
- CSCsv56881

Symptom: Each Switched Virtual Interface (SVI) for inband management must be configured with a different IP address. IPv6 has an error check feature. When an administrator enters a duplicate IPv6 address across two SVIs, the software fails the command due to the duplicate address. A similar error check should exist for IPv4 address configuration on SVIs.

Workaround: Do not configure duplicated IPv4 or IPv6 addresses.
- CSCsv24214

Symptom: When downgrading a Cisco Nexus 5000 Series switch from running a 4.0(1a)N1 image to a 4.0(0)N1(1) or 4.0(0)N1(1a) image, the startup configuration is not restored. This is caused by the defect CSCsq74395, which was resolved in 4.0(0)N1(2) and 4.0(0)N1(2a).

Workaround: After a downgrade, manually copy the startup configuration to the running configuration and reboot the system.
- CSCsv02866

Symptom: The command **show interface ethernet transceiver details** may show invalid calibration for DOM-supported 1 G SFP.

Send documentation comments to nexus5kdocs@cisco.com

Workaround: None.

- CSCsv10783

Symptom: The **show startup-config** command does not show the correct mode for a channel group. It shows the current mode for the channel group. The correct mode is stored. On a subsequent reload, the correct mode is configured on the channel group. The **show startup-config** always shows the current mode instead of the mode that was saved to startup. On a subsequent reload, the correct mode is configured. This is only a show command issue.

Workaround: None.

- CSCsv00989

Symptom: The **show interface ethernet transceiver details** command may show all zero values for a DOM-capable 1 G SFP.

Workaround: None.

- CSCsu48008

Symptom: When a Virtual Fibre Channel (VFC) interface is down, the **fclIfOperStatusCause** MIB object does not report the correct reason.

Workaround: Get the OperStatus from the CLI using the **show interface vfc x** command.

- CSCsu77946

Symptom: Within a configuration session, when you enable statistics on the PACL add more than 252 ACES to the ACL, and apply it to an interface, an error message is generated as the statistics counter is exhausted. Even if you try to remove the statistics keyword, it does not get removed. The result is that the ACL cannot be applied to the interface. This problem occurs only with a configuration session, and only after a configuration failure.

Workaround: Reduce the size of the ACL (fewer than 252 ACES) and re-apply the ACL to an interface. The statistics keyword will still remain and consume hardware resources.

- CSCsu93313

Symptom: Within a configuration session, 125 unique VACLs are created with a total of 1023 TCAM entries. The verify command fails with the following message:

```
d2-switch-2(config)# configure session 30
Config Session started, Session ID is 1
d2-switch-2(config-s)# verify
Failed to start Verification: Message Timed Out
d2-switch-2(config-s)# commit
Failed to complete Verification: no free label
```

This problem occurs with large VACL configurations. Once the Cisco Nexus 5000 Series switch is in this state, subsequent VACL configurations fail.

Workaround: Reload to recover the configuration.

- CSCsv19979

Symptom: Any FC port set to SD mode does not come up until the speed is configured manually. The port goes into the error disabled state and the only way to bring it online as SD is to manually set the speed 2 G or 4 G.

Workaround: Configure the speed manually to 2 G or 4 G.

Send documentation comments to nexus5kdocs@cisco.com

- CSCsv00402

Symptom: When you downgrade from the Cisco NX-OS 4.0(1a)N1 release to a previous software release, any static IGMP entries that have been configured over an Ether channel are lost after the downgrade.

Workaround: After you downgrade to the previous release and reload the switch, reconfigure any static IGMP groups configured over an etherchannel. Alternately, you can also do a **copy startup running** to reload the startup configuration. After that do a **copy running startup** to make sure the static IGMP entries are re added properly to the startup configuration.

- CSCsr20499

Symptom: When you restore a configuration to running-config from a configuration file, ACL manager may leak memory. The size of the leak is related to the size of ACL configurations and the number of times the restoration occurs. The switch may reboot if the ACL configuration is very large and the restoration occurs too many times.

Workaround: None.

- CSCsq64251

Symptom: TACACS+ fails if the user name input at login initiates a directed request authentication. The syntax to authenticate a directed request to a switch is **username@(IP address or name of TACACS+ server)**.

Workaround: Use RADIUS for directed request authentication.

- CSCsq76688

Symptom: The neighboring device for the Cisco Discovery Protocol (CDP) is not removed after shutting down the port for CDP hold time interval.

Workaround: None.

- CSCsl62545

Symptom: The fan LED on the Device Manager displays in amber color even though the fan is operating properly.

Workaround: None.

- CSCsr28868

Symptom: When the Fibre Channel over Ethernet (FCoE) feature is disabled, any untagged Ethernet packet with 00 00 in the Ethertype/length field is treated as an invalid packet and is forwarded out with a bad Ethernet CRC.

Workaround: None.

- CSCsr35452

Symptom: When the **ntp peer** command is configured on the MDS fabric and is distributed using CFS, the Nexus 5000 Series switch appends an incorrect VRF name **AC** to the command instead of **VRF management**.

Workaround: Use the **ntp server** command to synchronize time across the fabric.

- CSCsr36661

Symptom: When IGMP group membership is statically configured with private VLAN (PVLAN) host ports, the hardware gets programmed correctly. However, the membership information is not programmed for PVLAN host ports after the switch is reloaded.

Workaround: Delete and add the private VLAN association once again.

Send documentation comments to nexus5kdocs@cisco.com

- CSCsr64291
Symptom: When a port is congested due to traffic from multiple inputs, the traffic mapped to a user-defined QoS class can get discarded at a very slow rate (fewer frames per second) in spite of being configured with strict priority scheduling.
Workaround: None.
- CSCsr68690
Symptom: When an egress SPAN is configured on a port transmitting jumbo or large frames, the spanned frames are truncated to 2384 bytes.
Workaround: None.
- CSCsl21529
Symptom: An incorrect MTU value is displayed in the **show interface** command output. The Cisco Nexus 5000 Series switch only supports class-based MTU. Per-interface level MTU configuration is not supported. The switch supports jumbo frames by default. However, the **show interface** command output currently displays an incorrect MTU value of 1500 bytes.
Workaround: None.
- CSCsm03765
Symptom: The Set operation on the CISCO-IP-IF-MIB is not supported. You cannot set the mgmt0 IP address using SNMP.
Workaround: Use the CLI to set the mgmt0 IP address.
- CSCsm16222
Symptom: CFS does not support roles configuration distribution. Enter the **show cfs application** command to see the registered applications.
Workaround: Any features not registered with CFS need to be configured locally on the switch.
- CSCsl73766
Symptom: CFS does not support RADIUS configuration distribution. Enter the **show cfs application** command to see the registered applications.
Workaround: Any features not registered with CFS need to be configured locally on the switch.
- CSCso25966
Symptom: When an LACP port channel is configured between Catalyst 6500 and Cisco Nexus 5000 Series switches, and the configurations on both sides of the port channel do not match, the Catalyst 6500 LACP ports may change to the errordisable state.
Workaround: Fix the configuration to make it consistent on both peer switches of the port channel, and perform a **shut** and **no shut** operation on the Catalyst 6500 port channel interface.
- CSCso27446
Symptom: When a **shutdown** command is issued to the mgmt0 interface on a Cisco Nexus 5000 Series switch, the link never goes down and the remote end does not indicate that the link is down.
Workaround: None.
- CSCso46345
Symptom: The current version of NX-OS software running on the Cisco Nexus 5000 Series switches does not support Brocade i10K interop mode 4. The i10k v9.2.0.8 is supported by MDS in SAN-OS 3.2(2c), and 3.2(3) with interop mode 1 and 4.
Workaround: None.

Send documentation comments to nexus5kdocs@cisco.com

- CSCso74872

Symptom: When two SNMP walks are started simultaneously, one of them may fail with the following error:

```
OID not increasing
```

This problem does not occur with a single SNMP walk.

Workaround: This is not a permanent failure. Restart the walk and the problem will not occur as long as there is no other SNMP walk in progress.

- CSCso84269

Symptom: Occasionally, when reload is executed after bootup, and there has been no configuration change, the switch will display the following warning:

```
'WARNING: There is unsaved configuration!!!'
```

Workaround: Enter the **copy running startup** command. The problem will disappear.

- CSCsq10026

Symptom: When the small form-factor pluggable (SFP) is not in the Ethernet port, the **show interface** command output displays a bandwidth of 1 Gbps. When the SFP is plugged in, the bandwidth is displayed correctly (10 Gbps).

Workaround: None.

- CSCsq17571

Symptom: When an SNMP user creates or deletes virtual interface group (VIG), virtual Ethernet (VEth) or virtual FC (VFC) interfaces, the accounting log displayed by the **show accounting log** command does not get updated.

Workaround: Use the CLI for the configuration, which will update the accounting log.

- CSCsq35527

Symptom: When IGMP snooping is enabled on a switch, and the switch is the STP root and an STP topology change occurs, the IP multicast traffic may take a long time to converge. During this time, the IP multicast traffic may get affected.

Workaround: Configure a shorter query interval on the IGMP router to reduce the time it takes for ip-multicast traffic to converge in this topology.

- CSCsq35728

Symptom: When a SAN port channel is created, the following syslog message is displayed:

```
2008 May 20 06:09:13 switch %PORT_CHANNEL-3-MSG_SEND_FAILURE: failed to send
MAP_PARAM_FROM_CHANNEL to sap 45: Broken pipe"
```

There is no functionality loss and this message can be ignored.

Workaround: None.

- CSCso01268

Symptom: The following error message is displayed when a module is hot-swapped out:

```
2005 Jan 1 00:08:23 switch %KERN-4-SYSTEM_MSG: SI-VDC map entry <0, 0x0> does not
exist! - kernel"
```

There is no functionality loss and the message can be ignored.

Workaround: None.

Send documentation comments to nexus5kdocs@cisco.com

- CSCsq57558

Symptom: Enhanced Inter Switch Link (EISL) encapsulation is not supported on a Fibre Channel SPAN destination port. When EISL encapsulation is configured on the SPAN destination (SD) port, a VFT header is added to the packets going out of the SD port. The VFT header has VSAN information that helps distinguish traffic for various VSANs. This applies to any Fibre Channel SPAN source. By default, EISL encap is not added and packets are sent out without a VFT header, and independent of the type of SPAN source port. For a Cisco Nexus 5000 Series switch, the **switchport encap EISL** command does not have any effect.

Workaround: Use the Ethernet SPAN destination port to SPAN Fibre Channel traffic. FCoE packets going out of the Ethernet SPAN destination port will contain VSAN information in the Ethernet VLAN tag.

- CSCsq90423

Symptom: EISL encapsulation is not supported on Fibre Channel SPAN destination port in NPV mode. When EISL encapsulation is configured on the SPAN destination (SD) port, a VFT header is added to the packets going out of the SD port. The VFT header has VSAN information, which helps distinguish traffic for various VSANs. This applies to any Fibre Channel SPAN source. By default, EISL encapsulation is not added and packets are sent out without a VFT header, and independent of the type of SPAN source port. For a Cisco Nexus 5000 Series switch, the **switchport encap EISL** command does not have any effect.

Workaround: Use the Ethernet SPAN destination port to SPAN Fibre Channel traffic. FCoE packet going out of the Ethernet SPAN destination port will contain VSAN information in the Ethernet VLAN tag.

Resolved Caveats—Cisco NX-OS Release 4.0(1a)N2(1)

This section lists the resolved caveats for this release.

- CSCsv70815

Symptom: The default VRF is the system default for a VRF setting. Ideally, applications using VRF (such as TACACS+) assume a default VRF value if a VRF configuration is not specified by the administrator. However, TACACS+ is not set up properly unless the default VRF is configured.

Workaround:

The Cisco Nexus 5000 Series switch supports two possible VRFs, the default and the management. Configure the desired VRF when using TACACS+ service. You can configure the desired VRF using either of these configurations:

```
aaa group server tacacs+ t1
server 10.193.149.54
use-vrf management
```

```
aaa group server tacacs+ t2
server 20.1.1.2
use-vrf default
```

- CSCsv55655

Symptom: When the Cisco Nexus 5000 Ethernet port is configured in the 1 G mode of operation using the **speed 1000** command, it does not advertise and auto-negotiate the flow-control configuration. As a result, the link peer does not learn about the capabilities of the Cisco Nexus 5000 Series switches and does not enable flow-control on its end.

Send documentation comments to nexus5kdocs@cisco.com

Workaround: Disable auto-negotiation on the link peer and enable flow control for flow control to work over the link.

- CSCsv52513

Symptom: The VLAN interface 1 is internally created by the SVI daemon by default, so this VLAN interface cannot be deleted from the CLI.

Workaround: None.

- CSCsu50589

Symptom: If an invalid IP address is configured within a RADIUS configuration (such as a RADIUS server), any manipulation of the RADIUS configuration causes timeouts to the console such as **show running-configuration** or **copy running to startup**.

Workaround: Configure a valid IP address for a RADIUS configuration.

- CSCsu66201

Symptom: If the name server IP address is unreachable, the TACACS daemon gets stuck trying to resolve the server names. While it is stuck, TACACS commands, including common commands such as **show running-config** and **copy running-config startup-config**, are not processed.

Workaround: Fix the network connectivity to the name server IP address.

- CSCso65934

Symptom: Virtual interfaces are created by specifying the interface name in configuration mode. If the interface does not exist, the system creates the interface, and then enters interface configuration mode. If the user role prohibits access to that interface, this CLI is rejected and the user does not enter interface configuration mode, although the interface is created. Similarly, if a user does not have access to a virtual interface, the interface can still be deleted with the **no interface** command.

Workaround: None.

- CSCso82992

Symptom: Delete and insert role scopes for a role may fail and display the following error:

```
entry already exists
```

The problem occurs when you execute the following steps:

- Delete all role scopes for a role.
- Insert new role scopes for the same role.

Workaround: Repeat the above steps once again.

Resolved Caveats—Cisco NX-OS Release 4.0(1a)N1(1)

This section lists the resolved caveats for this release.

- CSCso91286

Symptom: When TACACS+ authentication is used to authenticate AAA users using ACS, the Cisco Cisco Nexus 5000 Series switch ignores the user to role binding information specified in the ACS. Users are logged in with their default roles. The default role for a new user is network-operator and for a user who is an administrator is network-admin.

Workaround: The user-to-role binding needs to be configured locally on the Cisco Nexus 5000 Series switch for the role binding to take effect.

Send documentation comments to nexus5kdocs@cisco.com

- CSCsu32247

Symptom: The Cisco Nexus 5000 Series switch executes the power on self test (POST) at bootup to validate hardware integrity of the ASICs. When a Fibre Channel port is connected to an HBA, the HBA driver could assert a LOS that causes a false failure to be reported by POST for one of the ASICs. As a result, all the ports serviced by the ASIC on the GEM are marked for hwFailure.

Workaround: Bypass POST at bootup by performing the following steps:

```
switch(config)# diagnostic bootup level bypass
switch(config)# copy running-config startup-config
switch(config)# reload
```

- CSCsv05115

Symptom: The Cisco Nexus 5000 Series switch crashes if CFS callhome is enabled on a it after a CFS callhome commit is performed on an attached MDS.

Workaround: None.

- CSCsv30392

Symptom: The Cisco Nexus 5020 switch has a Pktmgr memory leak in version 4.0(0)N1(2). This causes STP to stop functioning after awhile causing a Layer 2 Loop. After breaking the redundant connections, the switch is unable to be managed, due to a No buffer space available message.

Condition: The Cisco Nexus 5020 switch is setup in a triangle topology with two 6500 switches. Code version 4.0(0)N1(2) is loaded on the the Cisco Nexus 5020 switch. The redundant link had to be shut down in order to stop the loop.

Workaround: To fix the broken state, do not configure SVI.

- CSCso99821

Symptom: If PVLANS are created and deleted continuously and without pausing, the Ethernet interface may not be configurable and you have to reboot.

Workaround: Pause between the creation and deletion of PVLANS and do not perform multiple PVLAN operations at the same time. Alternately, you can create a PVLAN before any PVLAN interface is created and remove the switch port PVLAN from the interface before the PVLAN is deleted.

- CSCsr52118

Symptom: When you perform delete, add, shutdown or no shutdown operations on a VLAN, the port channel interface may lose VLAN membership in the forwarding plane. As a result, ports will not participate in any of the forwarding operations on that VLAN. This behavior applies to access port channels where the switch port access VLAN configuration matches the deleted and re-added VLAN. This behavior can occur for trunk port channels, if the deleted or re-added VLAN matches the native VLAN of the port channel.

Workaround: Enter the **shutdown** command or the **no shutdown** command on the port channel.

- CSCsr39670

Symptom: Although SNMP notification is enabled on the switch, traps for the power supply and fan modules are not received.

Workaround: None.

- CSCsr47531

Symptom: When a VSAN is configured as SPAN source, traffic from all the member ports is spanned to the SPAN destination port. When a switch is rebooted, the VSAN SPAN source remains in the down state.

Send documentation comments to nexus5kdocs@cisco.com

Workaround: Delete and add the VSAN sources for the SPAN session.

Resolved Caveats—Cisco NX-OS Release 4.0(0)N1(2a)

- CSCsu08988

Symptom: Telnet access is available only after reloading the Nexus 5020 switch with the **no telnet server enable** command running.

Workaround: Execute the **no telnet server enable** command once again after reload even if the command is saved in the startup-config. Additionally, you can also apply filters to the SVI to allow only trusted hosts to communicate with the system.

- CSCsu32247

Symptom: The Nexus 5000 Series switch executes the power-on self-test (POST) at bootup to validate hardware integrity of the ASICs. When a Fibre Channel port is connected to a host bus adapter (HBA), the HBA driver can trigger a signal loss that causes a false failure to be reported by POST for one of the ASICs. As a result, all the ports serviced by the ASIC on the GEM are marked for hardware failure.

Workaround: Do not connect HBAs to the FC expansion modules.

- CSCsu40126

Symptom: When the Cisco Nexus 5000 Series switch is configured to operate in the N-port virtualization (NPV) mode, the LOGOs received over the server port are not processed properly. This results in stale Fibre Channel IDs at the switch.

Workaround: Perform **shut** and **no shut** command operations on all the server ports to clear the stale state on the NPV switch.

- CSCsm66194 and CSCsr66209

Symptom: Logins may not be uniformly balanced across all available border ports when the Cisco Nexus 5000 Series switch operates in the NPV mode under conditions such as the following:

- The switch is reloaded.
- NPIV is disabled and re-enabled on the NPV core-switch.
- The NPV core-switch is reloaded.
- A new NP link is added.

Workaround: Perform **shut** and **no shut** command operations on all the server ports to rebalance the logins.

- CSCsu25775

Symptom: When the Cisco Nexus 5020 Series switch is connected to a 110 V power supply, the following problems occur:

- The syslog displays warnings at bootup. Although the hardware supports the 110 V input, the operating system incorrectly logs the warning message.
- The show environment power command incorrectly displays the available power as a negative value.

Workaround: The problem is not significant and the system operates normally with 110 V power supply input. There is no workaround to avoid the syslog message.

Send documentation comments to nexus5kdocs@cisco.com

Resolved Caveats—Cisco NX-OS Release 4.0(0)N1(2)

This section lists the resolved caveats for this release.

- CSCsq61505

Symptom: Problems are encountered while upgrading to Cisco NX-OS Release 4.0(0)N1(2) from Cisco NX-OS Releases 4.0(0)N1(1) or 4.0(0)N1(1a).

Workaround: None.

- CSCsq67305

Symptom: When the switch is operating in the N-port virtualization (NPV) mode, traffic may be disrupted on active Fibre Channel links if you enter the **shutdown** or **no shutdown** command on any of the NP mode uplink interfaces. The disruption also occurs if you change the interface mode of a port from F to NP or from NP to F. Traffic is disrupted for all F mode ports in the same VSAN as the NP-mode port.

Workaround: Change only the interface state or interface mode on NP mode Fibre Channel interfaces during a maintenance window.

- CSCso56749

Symptom: The current software does not have the ability to tag supervisor sourced frames separately depending on control or data. The frame always goes out with CoS 0.

Workaround: None.

- CSCso91286

Binding information specified in the Access Control System (ACS) is ignored when the Terminal Access Controller Access Control Plus (TACACS+) authentication is used to authenticate the AAA user using ACS.

Symptom: When TACACS+ authentication is used to authenticate the AAA user using ACS, the Cisco Nexus 5000 Series switch ignores the user-to-role binding information specified in the ACS. You are logged in with the default role of network-operator (for new users) and network-admin.

Workaround: Configure the user-to-role binding locally on the Cisco Nexus 5000 Series switch for the role binding to take effect.

- CSCsq23027

Symptom: Occasional Phy loopback failure is reported in power-on self-test (POST) routines. This is a sporadic issue with front port POST routines. Occasionally, when the system comes up, ports fail the loopback test.

Workaround: Reload the switch to confirm that it is a hardware defect.

- CSCsq27576

Symptom: FC-SP authentication is supported only with a switch over the E/TE port. Authentication with native Fibre Channel (FC) and Fibre Channel over Ethernet (FCoE) initiator or target is not supported.

Workaround: None.

- CSCsq32710

Symptom: SNMP users configured with a user-defined roles cannot retrieve any Fibre Channel interfaces.

Workaround: Use one of the predefined roles on the switch such as network-admin or network-operator.

Send documentation comments to nexus5kdocs@cisco.com

- CSCsq37899

Symptom: Deleting class-fcoe or class-default from an output policy map will cause **show policy-map** and **show running-config terminal** commands to be inconsistent if the **priority** keyword was configured for either of the classes.

Workaround: The default bandwidth percentage for **class-fcoe** and **class-default** is 50 percent. Before you remove these classes from an output policy, make sure that the remaining classes in the policy map do not exceed 50 percent. Alternatively, if you want to allocate minimum bandwidth to class-fcoe or class-default, configure the bandwidth for these classes to 0 percent.

- CSCsq39683

Symptom: Fibre Channel links with traffic running on it might generate syslog errors such as the following:

```
2008 May 21 15:08:55 switch %KERN-3-SYSTEM_MSG: FCP_ERRFCP_PORT:
gat_fcp_isr_process_blk_intr@1441, jiffies = 0xb1cec:ISR threshold reached, reg_block
= 0x8, num_regs = 6, idx = 0, src_bit = 11 - kernel
```

There is no loss in functionality and these messages can be ignored.

Workaround: None.

- CSCso83662

Symptom: One global MAC address is used in all STP and PVRST frames originating on various ports. This can result in inconsistent MAC moves on peer switches connected with multiple links.

Workaround: None.

Resolved Caveats—Cisco NX-OS Release 4.0(0)N1(1a)

This section lists the resolved caveats for this release.

- CSCsq53614

Symptom: Ethernet port channels configured between two Cisco Nexus 5000 Series switches cause a memory leak in the DCBX process. This leak eventually leads to the DCBX process crashing and a system reboot (in a few days). This issue occurs only with Ethernet port channels between Cisco Nexus 5000 Series switches. Port channels using the Catalyst 6500 Series switches do not have this problem.

Workaround: Disable receive and transmit Link Layer Discovery Protocol (LLDP) on the port channel members by entering the following configuration under each Ethernet interface that is a member of the port channel:

```
Interface Ethernet 1/1
  no lldp receive
  no lldp transmit
```

- CSCsq36609

Symptom: For a virtual Ethernet or virtual Fibre Channel interface configured as a SPAN source, changing the interface administrative state results in a memory leak. Deleting and adding SPAN sessions also causes memory leaks. The memory leaks eventually cause a switch reboot.

Workaround: Avoid using virtual Ethernet and virtual Fibre Channel interfaces as SPAN sources. Avoid deleting and adding SPAN sessions.

Send documentation comments to nexus5kdocs@cisco.com

Resolved Caveats—Cisco NX-OS Release 4.0(0)N1(1)

This was the first release of Cisco NX-OS for Nexus 5000 Series switches. There are no resolved caveats for this release.

Cisco Fabric Manager

Beginning with Cisco Fabric Manager release 3.4(1a), Nexus 5000 Series switches are supported by Fabric Manager.

If you are deploying Cisco Nexus 5000 Series switches with FCoE, you should operate Fabric Manager in Display FCoE mode. Display FCoE mode displays additional tree nodes, menu items, toolbar buttons, and topology nodes and links related to FCoE. To convert to Display FCoE mode, edit the `server.properties` file to set the `display FCoE` property to true.

For additional information about Fabric Manager, see *Cisco Nexus 5000 Series Switch Fabric Manager Software Configuration Guide, Release 4.0*

The following sections apply to Fabric Manager support for Cisco Nexus 5000 Series switches:

- [Limitations, page 27](#)
- [Caveats, page 28](#)

Limitations

This section lists the Cisco Fabric Manager limitations related to managing Cisco Nexus 5000 Series switches.

Ethernet Configuration

You cannot configure physical Ethernet interfaces using Fabric Manager or Device Manager. You must configure physical Ethernet interfaces using CLI commands.

SPAN

You cannot use Device Manager to configure Ethernet or virtual Ethernet interfaces as SPAN source ports or to configure Ethernet interfaces as destination ports. The workaround is to configure SPAN using CLI commands.

Zoning

In the Edit Local Full Zone Database tool, virtual Fibre Channel interfaces must be specified using the Switch Port WWN method of adding members to a zone. In the Add Members to Zone dialog box, the Switch & Port method and the Domain & Port method are not supported for virtual Fibre Channel interfaces.

Send documentation comments to nexus5kdocs@cisco.com

Caveats

This section lists the Cisco Fabric Manager caveats related to managing Cisco Nexus 5000 series switches.

Open Caveats

- CSCq57019
Symptom: The Fabric Manager N-Port Virtualization (NPV) Wizard does not list the Cisco Nexus 5000 Series switches that are operating in NPV mode.
Workaround: None.
- CSCsq06170
Symptom: In Fabric Manager, only one power supply is displayed for a Cisco Nexus 5000 Series switch with two power supplies.
Workaround: Use Device Manager, which displays the correct information when you use the Power Supplies menu item from the Physical menu.
- CSCso82992
Symptom: When using the Roles dialog box in Device Manager, if you edit the role scopes for a role that already has role scopes defined, and then click the **Apply** button, an error message dialog box states that the entry already exists.
Workaround: The changes are saved if you click the **Apply** button again.
- CSCsq23436
Symptom: The load-balancing and traffic-engineering features are not supported on Cisco Nexus 5000 Series switches in NPV mode, but Fabric Manager does not disable the configuration of these features. In the Switches > NPV information pane, the Load Balancing tab provides a check box to enable the feature, and the Traffic Engineering tab provides a dialog box to create a traffic engineering session.
Workaround: Ignore these configuration options for Cisco Nexus 5000 Series switches.
- CSCsq14828
Symptom: The web client does not display flow statistics or the Ethernet interface statistics for initiators using Fibre Channel over Ethernet (FCoE) connections to the Cisco Nexus 5000 Series switch.
Workaround: Use Flow Statistics in Fabric Manager, which displays the flow statistics correctly. The Device Manager displays the Ethernet interface statistics correctly.
- CSCsq32710
Symptom: If you logged in to the Cisco Nexus 5000 Series switch using a user-defined role, you cannot retrieve the Fibre Channel interfaces using SNMP.
Workaround: You can retrieve the information using CLI commands.

Related Documentation

The Nexus 5000 Series documentation is available at the following URL:

http://www.cisco.com/en/US/products/ps9670/tsd_products_support_series_home.html

Send documentation comments to nexus5kdocs@cisco.com

The following are related Cisco Nexus 5000 Series documents:

- *Cisco Nexus 5000 Series Switch CLI Software Configuration Guide*
- *Cisco Nexus 5000 Series Switch Fabric Manager Software Configuration Guide*
- *Cisco Nexus 5000 Series System Messages Reference*
- *Cisco Nexus 5000 Series MIB Quick Reference*
- *Cisco Nexus 5000 Series Command Reference*
- *Cisco Nexus 5000 Series Hardware Installation Guide*

Cisco Nexus 2000 Series documentation is available at the following URL:

http://www.cisco.com/en/US/products/ps10110/tsd_products_support_series_home.html

The following are related Cisco Nexus 2000 Series documents:

- *Cisco Nexus 2000 Series Hardware Installation Guide*
- *Cisco Nexus 2000 Series CLI Software Configuration Guide*

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, 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>

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.

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0812R)

© 2008 Cisco Systems, Inc. All rights reserved.

Send documentation comments to nexus5kdocs@cisco.com