

Send documentation comments to nexus5kdocs@cisco.com



Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Release Notes, Releases 4.2(1)N1(1), 4.2(1)N2(1), and 4.2(1)N2(1a)

Release: 4.2(1)N2(1a) -October 13, 2010
Part Number: OL-22747-01 E0

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 27.



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

Revision	Date	Description
A0	April 29, 2010	Created release notes for Release 4.2(1)N1(1)
B0	July 30, 2010	Updated release notes for Release 4.2(1)N2(1).
C0	October 13, 2010	Updated release notes for Release 4.2(1)N2(1a).
D0	March 17, 2011	Updated Limitations with Cisco Nexus 2148 Fabric Extender information.
E0	March 28, 2011	Updated Limitations with IGMP Snooping limitation.
F0	August 17, 2011	Added Enabling NPIV feature information.



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

Send documentation comments to nexus5kdocs@cisco.com

Contents

This document includes the following sections:

- [Introduction, page 2](#)
- [New and Changed Features, page 5](#)
- [Upgrade/Downgrade, page 7](#)
- [Installing Expansion Modules, page 9](#)
- [Limitations, page 10](#)
- [Caveats, page 14](#)
- [Related Documentation, page 27](#)
- [Obtaining Documentation and Submitting a Service Request, page 28](#)

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 2](#)
- [Cisco Nexus 5010 Switch, page 3](#)

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.

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

- 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. The N5K-M1060 8GFC GEM uses release 4.1(3)N2(1).

Cisco Nexus 2000 Series Fabric Extenders

The Cisco Nexus 2000 Series Fabric Extender was first released in Release 4.0(1a)N2(1). It 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.

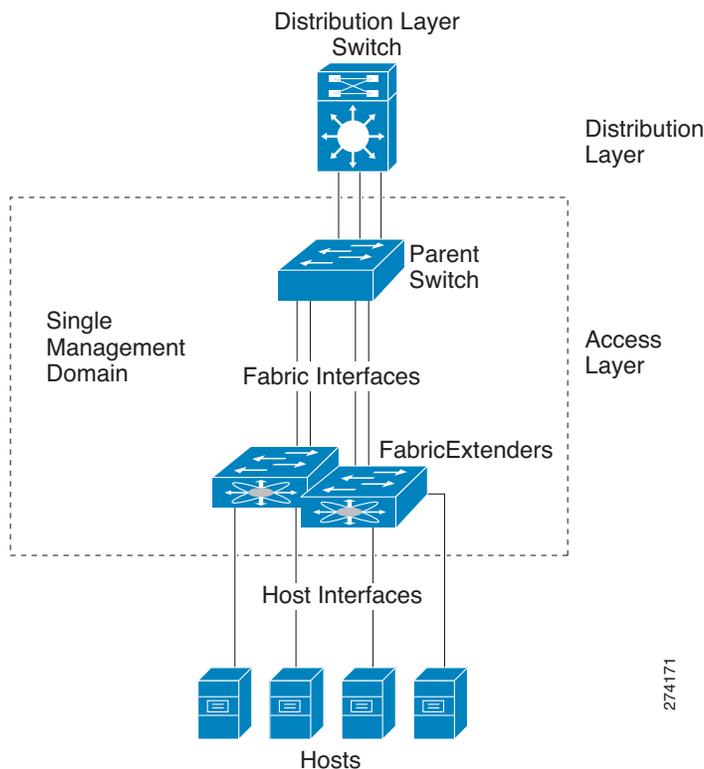
Scaling across a multitude of 1-Gigabit Ethernet, 10-Gigabit Ethernet, unified fabric, rack, and blade server environments, the Fabric Extender is designed to simplify data center architecture and operations.

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.

Send documentation comments to nexus5kdocs@cisco.com

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 2248TP Fabric Extender, page 4](#)
- [Cisco Nexus 2232PP Fabric Extender, page 5](#)
- [Cisco Nexus 2224TP Fabric Extender, page 5](#)

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.

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.

Cisco Nexus 2248TP Fabric Extender

The Cisco Nexus 2248TP is a stackable 1RU 450mm deep switch supporting 48 1000-TX host ports and 4 10G SFP+ network ports. Both 100Mbps and GE are supported on the 48 TX host ports.

Send documentation comments to nexus5kdocs@cisco.com

Host ports can be configured in a Etherchannel as well as part of a VPC. It is typically used in conjunction with the Cisco Nexus 5000 Series switch. The Cisco Nexus 2248TP is managed and configured by the upstream switch. The Fabric Extender software is shipped with the Cisco Nexus 5000 Series switch software. The Fabric Extender downloads the software image from the switch the same way that a module would download it from the supervisor in a modular chassis.

Cisco Nexus 2232PP Fabric Extender

The Cisco Nexus 2232PP is a stackable 1RU 450mm deep switch supporting 32 10G/1G SFP+ host ports and 8 10G SFP+ network ports. Host ports can be configured in a Etherchannel as well as part of a VPC. It is typically used in conjunction with the Cisco Nexus 5000 Series switch. The Cisco Nexus 2232PP is managed and configured by the upstream switch. The Fabric Extender software is shipped with the Cisco Nexus 5000 Series switch software. The Fabric Extender downloads the software image from the switch the same way that a module would download it from the supervisor in a modular chassis.

Cisco Nexus 2224TP Fabric Extender

This version of the Cisco Nexus 2000 product family is similar to the Cisco Nexus 2248T product but has 24 100/1000BaseT downlink ports, and 2 SFP+ uplink ports.

New and Changed Features

This section briefly describes the new features introduced in the Cisco NX-OS 4.2(1)N1(1) and Cisco NX-OS 4.2(1)N2(1) releases. This section includes the following topics:

- [Enabling NPIV, page 5](#)
- [Cisco NX-OS Release 4.2\(1\)N2\(1\), page 5](#)
- [Cisco NX-OS Release 4.2\(1\)N1\(1\), page 6](#)

Enabling NPIV

Beginning with Cisco NX-OS Release 4.1(3)N1(1), use the **feature npiv** command to enable NPIV. In previous releases, use the **npiv enable** command. Cisco NX-OS Release 4.2(1) documentation that refers to **npiv enable** will be corrected in future versions of the documentation.

Cisco NX-OS Release 4.2(1)N2(1)

Cisco NX-OS Release 4.2(1)N2(1) includes the following new or changed features:

- **Support for 1G on the Cisco Nexus 2000 2232PP**
- **Support for LACP Fast Timers.**
Changing the LACP timer rate allows you to modify the duration of the LACP timeout. Using the **lACP rate** command, you can set the rate at which LACP control packets are sent to an LACP-supported interface. You can change the timeout rate from the default rate (30 seconds) to a fast rate (1 second). This command is supported only on LACP-enabled interfaces.

Send documentation comments to nexus5kdocs@cisco.com

- **Support for the Cisco Nexus 2000 2224TP-1GE Fabric Extender**
(Cisco Nexus 2000 1GE FEX, 2 power supplies, 1 Fan Module, 24x100/1000-BaseT+2x10G-BaseT)
- **Support for queue-limit functionality for the Cisco Nexus 2000 host ports.**

Cisco NX-OS Release 4.2(1)N1(1)

Cisco NX-OS Release 4.2(1)N1(1) includes the following new or changed features:

- **ISSU**

In Service Software Upgrade (ISSU) support on the N5K provides the capability to perform transparent software upgrades, reducing downtime and allowing you to integrate the newest features and functions with little or no effect on network operation for Ethernet, storage, and converged network environments.



Note Non disruptive software upgrade or downgrade is supported between the Cisco NX-OS 4.2(1)N1(1) release and any release after the Cisco NX-OS 4.2(1)N1(1) release in specific topologies. Upgrade or downgrade between any release before 4.2(1)N1(1) to 4.2(1)N1(1) and beyond would be disruptive. For details, refer to the Upgrade Downgrade Configuration Guide.

- **F_Port Trunking and Channeling**

Trunking, also known as VSAN trunking, enables you to interconnect ports to transmit and receive frames in more than one VSAN over the same physical link. As of Cisco NX-OS Software Release 4.2(1)N1(1), trunking will be supported on F ports on the Nexus 5000 series. F_Port Channeling will allow the configuration of 16 member ports per port channel and a maximum of 4 port channels per Nexus 5000 switch. F_Port Trunking and Channeling is supported on the links between a Fiber Channel switch in the npv mode and another in full FC switch mode.

- **VTP Transparent**

In the Cisco Nexus 5000 Series switches, the VLAN Trunking Protocol (VTP) works in transparent mode, allowing you to extend a VTP domain across the device. Layer 2 trunk interfaces, Layer 2 trunk over physical interfaces, and Layer 2 port channels will support VTP transparent functionality. This feature relays all VTP protocol packets that the device receives on a trunk port onto all other trunk ports. When the VTP feature is disabled, VTP protocol packets are not relayed.

- **Support for the Cisco Nexus 2232 and the Cisco Nexus 2248**

- **Local Port Channels on FEX-10G and FEX-100/1000**

Cisco NX-OS Software Release 4.2(1)N1(1) will enable users to configure local port channels with or without LACP on the Nexus 2248PP and Nexus2232TP ports.

- **ACLs for SNMP Communities**

Cisco NX-OS Software Release 4.2(1)N1(1) will enable the assignment of an access list (ACL) to a community to filter incoming SNMP requests. If the assigned ACL allows the incoming request packet, SNMP processes the request. If the ACL denies the request, SNMP drops the request and sends a system message.

- **Error-Disable Recovery**

Cisco NX-OS Software Release 4.2(1)N1(1) will enable users to configure the automatic error-disable recovery timeout for a particular error-disable cause and configure the recovery period.

Send documentation comments to nexus5kdocs@cisco.com

The `errdisable recovery interval` command can be used to change the recovery period within a range of 30 to 65535 second and also change the recovery timeout for a particular err-disable cause.

- **AAA Command Authorization**

The AAA authorization and authentication cache feature allows users to cache authorization and authentication responses for a configured set of users or service profiles. When a TACACS+ server authorization method is configured, users can authorize every command that executed with the TACACS+ server. This includes all EXEC mode commands and all Configuration Mode commands.

- **Support for FCoE on 10GE Fabric Extender host ports on the Cisco Nexus 2232PP**
- **Support for 100Mbps/1000Mbps and auto negotiation on the 2248TP**
- **Support for DOM on SFP-DOM optics**
- **ACLs for SNMP community string**
- **Fabric Manager support**
- **Support for 12000 STP logical ports with a maximum of 4,000 non-edge STP logical ports**
- **Support for a maximum of 12 Fabric Extenders dual-homed to a vPC Cisco Nexus 5000 Series switch pair and a maximum of 576 hosts connected to Fabric Extenders connected to Cisco Nexus 5000 Series switches.**

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 27.

Upgrade/Downgrade

This section describes issues you may encounter when you upgrade to or downgrade from releases. This section includes the following topics:

- [Downgrading from Cisco NX-OS Release 4.2\(1\)N2\(1\), page 7](#)
- [Downgrading from Cisco NX-OS Release 4.2\(1\)N1\(1\), page 8](#)
- [Upgrading to Cisco NX-OS Release 4.2\(1\)N1\(1\), page 8](#)
- [Upgrade Downgrade Matrix, page 9](#)

Downgrading from Cisco NX-OS Release 4.2(1)N2(1)

Starting with the with Cisco NX-OS Release 4.2(1)N2(1) release, LACP rate fast is supported. If you would like to downgrade to previous releases, **install all** prints the following warning:

```
"Configuration not supported - LACP fast rate is enabled",
"Use \"lACP rate normal\" on those interfaces"
```

You should change the LACP rate to normal before downgrading. If you ignore the warning and force **install all** to proceed, then it is possible that the leftover LACP rate fast configuration would still be active with previous releases of software but the behavior would be unpredictable and link flap may occur. It is recommended that you follow the warning from **install all** and change the LACP rate setting to normal. For details see, CSCth93787

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

Downgrading from Cisco NX-OS Release 4.2(1)N1(1)

When downgrading from the Cisco NX-OS Release 4.2(1)N1(1) to earlier releases like 4.1(3)N1(1) and 4.1(3)N2(1), features like Trunking F-Port/F-Port-Channel, ACL for SNMP communities, and Disabling of HTTP Server are not shown as compatible. If you downgrade to an earlier release from Cisco NX-OS Release 4.2(1)N1(1) while these features are configured, then these features will not work as they are not supported in the previous release. For details see, CSCtd71387, CSCtd66949, CSCtg12017, CSCtf76649.

When downgrading from Cisco NX-OS Release 4.2(1)N1(1) to the 4.1(3)N1(1) release, channel-group configuration on FEX ports is lost after the downgrade. You need to reconfigure the channel-group on FEX ports after the downgrade. For details see CSCtb34477.

Upgrading to Cisco NX-OS Release 4.2(1)N1(1)

When upgrading from pre Cisco NX-OS 4.2(1)N1(1) releases, to any release, policy description is lost. This problem does not exist when upgrading from Cisco NX-OS Release 4.2(1)N1(1) and beyond and any future releases. You should re-configure policy description after an upgrade. For details see, CSCth14225.

When upgrading from earlier releases to Cisco NX-OS Release 4.2(1)N1(1), FC port security configuration may be lost. You should reconfigure FC Port Security after the upgrade to Cisco NX-OS Release 4.2(1)N1(1). For details, see CSCtg43468

When upgrading from earlier releases to Cisco NX-OS Release 4.2(1)N1(1), ASCII configuration conversion may fail if there is description configured under a class-map or policy map. You should remove description under class-map/policy-map before upgrading to Cisco NX-OS Release 4.2(1)N1(1), or reapply the QOS configuration again after the upgrade. For details see, CSCtg55716.

Upgrading Power Sequencer

Under certain conditions, a voltage spike exceeding the system voltage guard band and glitch filter settings may result in a power cycle of the system mezzanine board. This results in the failure of ports on the mezzanine board. To solve the issue, you need to upgrade to the Cisco NX-OS 4.2(1)N1(1) release and make sure that power sequencer has been upgraded to v1.2 version (show version). Follow the power-sequencer upgrade procedure to upgrade the power sequencer version to v1.2 to get around CSCsy21017 and CSCth33969.

If the switch is upgraded to the Cisco NX-OS 4.2(1)N1(1) release, but is not power-cycled as the procedure instructs, the switch will have instructions for power sequencer upgrade, but the power sequencer will not actually be upgraded. The show version will indicate a v1.2 power sequencer, but that only indicates the power sequencer upgrade instructions have been programmed. Therefore, if the switch admin cannot confirm the power off/on of the switch, it is advisable to perform a power off/on to ensure the power sequencer is actually upgraded.

The steps to upgrade the power-sequencer with the Cisco NX-OS release 4.2(1)N1(1) are as follows:

-
- Step 1 Download the Cisco NX-OS Release 4.2(1)N1(1) kickstart and system image to the system.
 - Step 2 Enter the **install all kickstart kickstart_url system system_url** command to start and upgrade to the Cisco NX-OS release 4.2(1)N1(1). When prompted to confirm the upgrade, review the upgrade table and select y to proceed.

Send documentation comments to nexus5kdocs@cisco.com

- Step 3** After completing the installation, the system reloads and displays a Cisco NX-OS Release 4.2(1)N1(1) image.
- Step 4** Repeat step 2 to re-install the Cisco NX-OS Release 4.2(1)N1(1) image. During this process, the upgrade table should display the upgrade action for the power-sequencer and then upgrade the power-sequencer.
- Step 5** After **install all** has completed installation, power-cycle the system. If you skip this step, the power-sequencer will not be updated till the next power-cycle.
- Step 6** After the system comes up, confirm that the power-sequencer has been upgraded by running **show version**. The **show version** only confirms if the power-sequencer has the updated instructions. The updated instructions do not take effect if the system was not power-cycled.

Upgrade Downgrade Matrix

Cisco NX-OS Release	Upgrade to Cisco NXOS Release 4.2(1)N1(1) and beyond	Downgrade from Cisco NXOS Release 4.2(1)N1(1) and beyond
4.0(0)N1(1)	No	No
4.0(1a)N1(1)		
4.0(1a)N2(1)	Yes	No
4.0(1a)N2(1a)		
4.1(3)N1(1)	Yes	Yes
4.1(3)N1(1a)		
4.1(3)N2(1)	Yes	Yes
4.1(3)N2(1a)		

Installing Expansion Modules

When you install an expansion module on a Cisco Nexus 5000 Series switch, check the status of the module installation in the system logs, as follows:

```
e7-dut-1# show module
Mod Ports  Module-Type                Model                Status
-----
1    40      40x10GE/Supervisor        N5K-C5020P-BF-XL-SU  active *
2     6       6x1/2/4/8G FC Module     N5K-M1060            ok
3     6       6x1/2/4/8G FC Module     N5K-M1060            ok

Mod  Sw                Hw      World-Wide-Name(s) (WWN)
---  -
1    4.1(3)N2(1)      1.2    --
2    4.1(3)N2(1)      0.0    20:41:00:0d:ec:b4:6a:80 to 20:46:00:0d:ec:b4:6a:80
3    4.1(3)N2(1)      0.0    20:81:00:0d:ec:b4:6a:80 to 20:86:00:0d:ec:b4:6a:80

Mod  MAC-Address(es)                Serial-Num
---  -
1    000d.ecb4.6a88 to 000d.ecb4.6aaf    JAF1314AQHR
2    000d.ecb4.6ab0 to 000d.ecb4.6ab7    JAF1325BBGE
3    000d.ecb4.6ab8 to 000d.ecb4.6abf    JAF1325BBJG
```

Send documentation comments to nexus5kdocs@cisco.com

If the module is not seated properly, an error message is displayed as follows:

```
2009 Aug 3 23:45:16 Edge-2 %PFMA-2-MOD_INSERTION_FAILED: Module 2 insertion failed.
Module might not be seated properly. Please try removing the module and the n
re-insert after five seconds or more.
```

For details see, the Expansion Modules section of the Cisco Nexus 5000 Series Hardware Installation Guide.

Limitations

This section describes the limitations in Nexus 5000 Series switches and the Cisco Nexus 2000 Series Fabric Extenders for the Cisco NX-OS 4.2(1)N2(1a), Cisco NX-OS 4.2(1)N2(1) and Cisco NX-OS 4.2(1)N1(1) releases.

- When upgrading from pre Cisco NX-OS 4.2(1)N1(1) releases, to any release, policy description is lost. This problem does not exist when upgrading from Cisco NX-OS Release 4.2(1)N1(1) and beyond and any future releases. You should re-configure policy description after an upgrade. For details see, CSCth14225.
- Starting with the with Cisco NX-OS Release 4.2(1)N2(1) release, LACP rate fast is supported. If you would like to downgrade to previous releases, **install all** prints the following warning:


```
"Configuration not supported - LACP fast rate is enabled",
  "Use \"lACP rate normal\" on those interfaces"
```

 You should change the LACP rate to normal before downgrading. If you ignore the warning and force **install all** to proceed, then it is possible that the leftover LACP rate fast configuration would still be active with previous releases of software but the behavior would be unpredictable and link flap may occur. It is recommended that you follow the warning from **install all** and change the LACP rate setting to normal. For details see, CSCth93787
- When an FC SPAN dest port is changed from SD to F mode and back to SD mode on a NPV switch, the port goes into error-disabled state. Perform a shut/no-shut after the mode change recovers the port. This problem is only with NPV mode. For details, see CSCtf87701
- If you configure a Nexus 2248TP port to 100Mbps speed instead of auto-negotiation, auto-negotiation does not occur. This is expected behavior. Both sides of the link should be configured to both hardwired speed or both autonegotiate.
 - no speed** autonegotiates, advertises all speeds (only full duplex)
 - speed 1000** autonegotiates only for a 802.3x pause
 - speed 100** does NOT autonegotiate, pause cannot be advertised. Peer must be set to NOT autonegotiate and fix at 100Mbps (similar to the N2248TP)
 For details, see CSCte81998
- Given the implementation of single CPU ISSU, STP root on PVST region with the switches on MST region is not supported. The PVST simulation on the boundary ports go into PVST sim inconsistent blocked state. So the STP active path is broken. In order to avoid this, to avoid the above, move all STP root on MST region. By doing so you cannot support non-disruptive ISSU since it will fail for Non-Edge Designated Forwarding Ports required for ISSU. For details see CSCtf51577. Refer to the configuration guide for scenarios and topology in which non-disruptive upgrade is supported.
- IGMP queries sent in CSCtf94558 are group-specific queries which are sent with destination IP/MAC address as the group's address.

Send documentation comments to nexus5kdocs@cisco.com

GS queries are sent for IP address : 224.1.14.1 to 224.1.14.100 [0100.5E01.0E01 to 0100.5E01.0E64]

These are not link-local addresses. So, by default they will not be flooded by hardware into the vlan. They will be sent only to the ports which have joined this group.

This is expected behavior during ISSU.

In another scenario, the IGMP global queries [dest IP 224.0.0.1] get flooded correctly in the vlan.

Group specific queries are not forwarded to ports other than the one that joined the group during ISSU. The reason to forward group specific queries towards hosts is to avoid them having to leave the group. However if no group has joined the group then this is not a problem, where as if there is an interface that has joined the group then the queries are expected to make it to the host. While the behavior is different from when ISSU is not occurring, it is sufficient and works as expected and there is no traffic impact. For details, see CSCtf94558.

- The meaning of MTU configuration has changed from pre-4.2(1)N1(1) releases to 4.2(1)N1(1) release. In pre-4.2(1)N1(1) releases, the configured MTU included the Ethernet payload and Ethernet headers. In 4.2(1)N1(1) release, the configured MTU includes only the Ethernet payload and not the Ethernet headers. When upgrading/downgrading between pre-4.2(1)N1(1) and 4.2(1)N1(1) releases, NX-OS will automatically convert the configuration to address this semantic change by adding/subtracting 38 to the MTU to address the Ethernet header size.

In a VPC configuration, the MTU per class needs to be consistent on both switches in the VPC domain for the VPC peer-link to come up. When upgrading/downgrading a working VPC setup between pre-4.2(1)N1(1) and 4.2(1)N1(1) releases, the MTU is adjusted appropriately to make sure that the MCT peer-link always comes up.

However if you add a peer-link between two switches in a VPC domain that are identically configured (MTU in particular) with one switch running 4.2(1)N1(1) and another switch running pre-4.2(1)N1(1), then the VPC peer-link will not come up as the MTU will be inconsistent between the two switches.

This is not an issue when upgrading or downgrading a working switch pair in a VPC domain, but is only a problem when adding a peer-link between two switches running 4.2(1)N1(1) and pre-4.2(1)N1(1) releases that were not in the same VPC domain before.

The solution is to upgrade/downgrade one switch to match the version on the other switch and reconfigure to the MTU to be consistent on both sides. For details see CSCtg27538

- The channel-group configuration is not applied to the Cisco Nexus 2000 downlink interface after downgrading to the Cisco NX-OS Release 4.1(3)N1(1) software. This happens if the **speed 1000** command is present under the context of the port-channel. As a workaround, reconfigure the **channel-group** command after the system comes up and reapply the config from the saved configuration in the bootflash. (For details see CSCtc06276)
- When a Private VLAN port is configured as a TX (egress)SPAN source, the traffic seen at the SPAN destination port is marked with the VLAN that the frame ingressed into the switch with. There is no workaround.
- 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.
- The Cisco Nexus 2148 Fabric Extender does not support frames with the dot1p vlan 0 tag.
- 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.
- 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.
- The class-fcoe cannot be removed even if Fibre Channel is not enabled on a switch.
- 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.

Send documentation comments to nexus5kdocs@cisco.com

IGMP Snooping Limitation

On the Cisco Nexus 5010 switch and the Cisco Nexus 5020 switch with a Cisco Nexus 2000 Series Fabric Extender (FEX) installed, unregistered IP multicast packets on one VLAN are forwarded to other VLANs where IGMP snooping is disabled. We recommend that you do not disable IGMP snooping on the Cisco Nexus 5010 switch and the Cisco Nexus 5020 switch. A static IGMP join can be configured for devices intended to receive IP multicast traffic but not to send IGMP join requests. This limitation applies to the Cisco Nexus 5010 switch and the Cisco Nexus 5020 switch only.

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 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.

Send documentation comments to nexus5kdocs@cisco.com

Caveats

This section includes the following topics:

- [Open Caveats, page 14](#)
- [Resolved Caveats—Cisco NX-OS Release 4.2\(1\)N2\(1a\), page 23](#)
- [Resolved Caveats—Cisco NX-OS Release 4.2\(1\)N2\(1\), page 24](#)
- [Resolved Caveats—Cisco NX-OS Release 4.2\(1\)N1\(1\), page 26](#)

Open Caveats

This section lists the open caveats for the Cisco NX-OS 4.2(1)N2(1) release.

- CSCth69160

Symptom: The SVI on secondary VLAN does not work.

Workaround: Assign SVI to non Private VLANs
- CSCti11823

Symptom: The Cisco NXOS version 4.2(1)N1() supports only 10GE on the Nexus 2232 host ports. If the administrator plugs in 1Gig SFPs on Nexus 2232 host ports with Cisco NXOS 4.2(1)N1(1) running, the SFP validation fails, the LED blinks amber and the interface shows LinkNotConnected. The administrator performs an ISSU from Cisco NXOS 4.2(1)N1(1) to Cisco NXOS 4.2(1)N2(1) and the Cisco NXOS 4.2(1)N2(1) version supports GE on Nexus 2232 host ports. After ISSU, if the administrator configures speed 1000 on the interface with GE SFP plugged in, the GE interface comes up (traffic flows) but the GE port LED continues to blink amber. The LED should be green after and ISSU to the Cisco NXOS version 4.2(1)N2(1) and configuration of speed 1000.

Workaround: Unplug and re-plug SFP and the LED will turn green.
- CSCtc62994

Symptom: With combining RBAC roles (multiple roles assigned to the same user account), interface policies in those roles do not work on per role basis. In the example shown in this bug, user lan-admin is assigned to the following 3 roles: LAN-admin, LAN-admin-ETH and LAN-admin-no. Roles LAN-admin and LAN-admin-no include eth1/1 in their interface policy permit list and both deny the **shutdown** command under their permitted interfaces.

Roles LAN-admin-ETH excludes eth1/1 from its interface policy permit-list and permits the **shutdown** command under its permitted interfaces.

When all 3 roles are assigned to the user, lan-admin, the user should not be allowed to **shutdown interface eth1/1** but it currently can.

Workaround: Do not configure multiple roles for the same user account
- CSCtd31131

Symptom: If you change PVLAN configuration as follows very rapidly (less than 5 seconds between any two commands), then the interface may end up in an err-disabled state.

```
DUT2(config-vlan)# vlan 421
DUT2(config-vlan)# private-vlan community
DUT2(config-vlan)# private-vlan isolated
DUT2(config-vlan)# private-vlan community
DUT2(config-vlan)# private-vlan isolated
```

Send documentation comments to nexus5kdocs@cisco.com

Workaround: **Shut** the interface, followed by **no shut** if the interface recovers properly.

- CSCtf32340

Symptom: When you change the VSAN or Interface Scope of an existing Role via DM, a **This entry already exists** error dialog is shown and the change is not applied

Workaround: Use the CLI to make the scope changes for the existing role.

- CSCtf79253

Symptom: This is the case of having an STP topology with parallel links (there is a loop) on the secondary switch with a path cost for the non-vpc ports is smaller than the vpc ports. Therefore, peer-link would end up being blocked by STP. In this case, transient traffic loop could be formed.

Workaround: Recommended best practice topology for deployment does not have parallel links to VPCs that have their cost tweaked to be higher than VPCs. By default all VPC links have preferred (smaller) cost over non VPC links.

- CSCtf98638

Symptom: The message "%SYSMGR-5-SUBPROC_KILLED "System Manager (core-client) (PID 7679) hasn't caught signal 9 (no core)." is printed when the system is rebooting after an install.

Freeing memory in the file system.

```
2010 Apr 1 08:39:09 SW2-5020-ANIL %$ VDC-1 %$ %CALLHOME-2-EVENT: SW_CRASH
cimxmlserver in slot 1 crashed with crash type : stateful crash
[#####] 100% -- SUCCESS
```

Loading images into memory.

```
[#####] 100% -- SUCCESS
```

Saving supervisor runtime state.

```
Apr 1 08:39:18 %SYSMGR-5-SUBPROC_KILLED "System Manager (core-client)" (PID 7679)
hasn't caught signal 9 (no core).
[#####] 100% -- SUCCESS
```

Workaround: There is no functional impact. Ignore the message

- CSCtg33706

Symptom: Debug LACP is not available on the Fabric Extender ports

Workaround: Use LACP event histories etc. on the Fabric Extender. Other show commands should also help with debug information.

- CSCtd15304

Symptom: When you perform a Cisco Nexus 5000 series switch software install using the Fabric Manager, the Success Reset status message is shown just before the switch reboots.

Workaround: To determine the status of the software install, do the following:

- Close the wizard and go to the main FM screen
- Click on the Rediscover...button in the toolbar and wait for the rediscovery to complete. Once the rediscovery is complete, you may encounter either of the scenarios:
 - If the topology shows the switch with a cross, the switch is rebooting/down. Wait for some time and repeat the second step again.

If the topology shows the switch as discovered/managed, select the corresponding fabric tree node from the Logical Domains tree. In the right hand side panel, under the Switches tab information about the switches along with their current versions is displayed. Use this to confirm the status of the software install.

Send documentation comments to nexus5kdocs@cisco.com

- CSCtd70554

Symptom: When you downgrade from the Cisco NX-OS Releases 4.1(3)N2(1) or 4.1(3)N2(1a) to the Cisco NX-OS Release 4.1(3)N1(1) or 4.1(3)N1(1a), the **feature fc-port-security** command does not get converted to **feature port-security**. As a result, the FC port security configuration gets lost and remains disabled.

Workaround: After the downgrade to Cisco NX-OS Releases 4.1(3)N1(1) or 4.1(3)N1(1a), re-enable the FC port security feature by executing the **command feature port-security**.
- CSCta77490

Symptom: When the type of a pvlan is toggled from being a regular vlan to a pvlan and back to regular vlan in very small interval of time, the type change fails.

Workaround: Issue the type change commands with a 5 seconds gap in between.
- CSCtb34546

Symptom: When a PACL with **deny ip any any** is applied on mgmt0, CFS discovery gets stuck.

```
ip access-list ip1
  10 deny ip any any
```

Applying such a ACL on the mgmt0

```
int mgmt 0
  ip access-group ip1 in
```

would cause this issue.

Workaround: Remove the **deny ip any any** rule from the PACL applied on mgmt0 interface.
- CSCtb61197

Symptom: When a port-channel provisioning fails because the system has reached the limit of number of port-channels supported, output of **show san-portchannel** will still display the port-channel as present but **down**. The port-channel will be seen as **down** even if the member interface is operationally up because it could not be provisioned correctly due to resource limitation.

Workaround: None.
- CSCtc44231

Symptom: LACP port-channel doesn't come up. A vlan is deleted from the switch which is also configured as native vlan for the lacp port-channel.

Workaround: Create the vlan or remove the native vlan config from the lacp port-channel.
- CSCtc77180

Symptom: Ports are error disabled with error **Ethernet interface not present** if **feature fcoe** is enabled immediately after the switch prompt comes up on bootup.

Workaround: Enable **feature fcoe** after confirming that the interfaces are displayed in the output of **show interface brief**.
- CSCtb58641

Symptom: If a mac-address moves from an isolated host port to a promiscuous trunk port, in certain conditions, the mac-address is never cleared from the system.

Workaround: None
- CSCtc36397

Symptom: Changing the role-priority and flapping the peer-link does not change the roles of the vPC peers. This happens when one of the switch has its role as Operational primary due to an earlier reload of the primary switch.

Send documentation comments to nexus5kdocs@cisco.com

Workaround: None

- CSCtb84512

Symptom: In mixed span mode where ethernet port-channel, vfc and FC ports are span sources and ethernet interface is a span destination, vfc flap causes the traffic coming in on ethernet port-channel to be not spanned.

Workaround: Remove and add span source command for the ethernet port-channel.

- CSCtb94310d

Symptom: With a san port-channel as the source and ethernet interface as the destination, removing the channel-group config from the san port-channel member causes monitor session to go to error state.

Workaround: Unconfigure and reconfigure the monitor session.

- CSCtb53820

Symptom: After **save** and **reload** with a monitor session configuration where source is a vsan and destination is an fc port, the monitor session goes to error state.

Workaround: Unconfigure and reconfigure the monitor session.

- CSCtc04213

Symptom: Vlan related configurations do not get applied on a range of interfaces. This issue may occur when you try to configure a vlan configuration on a range of interfaces in a way that the number of interfaces being configured at a time is greater than 192. As a result, even though the vlan configuration command returns with no errors on the console, the vlans do not get applied to the interfaces. You can confirm this by running the cli, **show system internal ethpm info interface *one of affected***

Workaround: Select a smaller range of interfaces to apply the vlan configuration to.

- CSCsz82199

Symptom: You cannot enable std.pause on a port-channel interface connected between two Cisco Nexus 5000 Series switches. Enable std.pause between two Cisco Nexus 5000 Series switches and configure std.pause in the hardware.

Workaround: None

- CSCsv39939

Symptom: Incorrect values are displayed for interface capabilities for ports for Cisco Nexus 2000 Series Fabric extenders connected to a Cisco Nexus 5000 Series switch. In particular, the number of input and output queues for the ports are displayed as zero instead of two.

Workaround: This is a display issue and does not impact functionality.

- CSCsz81365

Symptom: SPAN session should stop reflecting packets as soon as mapping is removed.

Workaround: None.

- CSCta04383

Symptom: When you install a new image from one of two vPC switches, the vPC switch gets upgraded. Also, each of the connected Fabric Extenders update their firmware with the new version, but continue to run the current version and stay connected to the other vPC switch. When you reload the upgraded switch, but revert back to the older version of software on the switch, both the vPC

Send documentation comments to nexus5kdocs@cisco.com

switches and all the Fabric Extenders run the older version of software. However, the Fabric Extenders have the more recent version of software in the firmware. When the Fabric Extenders reload there is loss of connectivity with the hosts.

Workaround: Re-install the older version of software from either of the two switches while the Fabric Extenders are connected to that switch.

- CSCsx35870

Symptom: The CLI times out when a large number of VLANs are created and deleted followed by PVLAN creation and deletion. The system indicates that the Ethernet Port Manager (ethpm) has timed out to communicate with the SPAN manager or PVLAN manager. As a result, some of the PVLAN interfaces will be error disabled.

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

- CSCsx59489

Symptom: Call home notifications for events generated when both a switch and a Fabric Extender are rebooting may contain a timestamp of January 1 1970 as shown in the following example:

```
System Notification from sample-system - environment:minor - 1970-01-01 00:00:00
GMT+0000
```

The most likely scenario where this would occur is after a power failure or after issuing the **reload all** command. The event is generated before the Fabric Extender connects with the switch and before the local time is updated for the Fabric Extender.

Workaround: None.

- CSCsx80279

Symptom: When traffic is sent at line rate as a single burst, all addresses are not learned when egress interfaces are FEX facing ports. This problem does not occur if sustained traffic is sent for more than 0.4 seconds.

Workaround: Resending the unlearned MAC addresses would render them relearned.

- CSCsy99816

Symptom: When a Cisco Nexus 2000 Fabric Extender is already online and a fabric interface that is not part of a port channel is configured with a serial number that does not match that of the FEX, the fabric interface will be brought down, and **show interface fex** does not display the reason for being down.

Workaround: None. This is a display issue.

- CSCsy02439

Symptom: Under some circumstances, the FC MAC driver displays the following error message:

```
%KERN-3-SYSTEM_MSG: FCP_ERRFCP_PORT: gat_fcp_isr_ip_fcmac_sync_intr@424, jiffies =
0x1c422:Unknown intr src_id 41 - kernel
```

The error message is when an unused interrupt in the MAC fires. The error message does not indicate any functional problem.

Workaround: None

- 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

- 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.
- 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 period when FEX interfaces are not yet created but interface configuration is applied and fails. (See also CSCsw21301)

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.
- 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. (See also CSCsx60187)

Workaround: Do not configure duplicated IPv4 or IPv6 addresses.
- CSCsv02866

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

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.
- 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.

Send documentation comments to nexus5kdocs@cisco.com

- 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.
- 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.
- 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
- 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

Send documentation comments to nexus5kdocs@cisco.com

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 error-disabled 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.

- 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:

Send documentation comments to nexus5kdocs@cisco.com

'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.

- 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.

- 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,

Send documentation comments to nexus5kdocs@cisco.com

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.

- CSCsv93922

Symptom: If the modulo(%) operator is used in a Cisco Nexus 2000 Series Fabric Extender description the **show fex <fex-id>** command brings up the following error message

ERROR: bad format: non escaped % not followed by 's'.

Workaround: Remove the modulo(%) operator from the Cisco Nexus 2000 Series Fabric Extender description

- CSCsv95478

Symptom: The Cisco Nexus 2000 Series Fabric Extender pinning redistribute command does not wait for user prompt with a yes or no operation.

Workaround: None.

- CSCsv15775

Symptom: When priority tagged frames are received on Cisco Nexus 2000 Series Fabric Extender ports, they are dropped and not forward on the native or default VLAN of the port. The MAC addresses are not learned.

Workaround: None.

- CSCth93531

Symptom: The show port-channel load-balance command shows the correct output only when source interface is specified.

Workaround: None.

Resolved Caveats—Cisco NX-OS Release 4.2(1)N2(1a)

This section lists the resolved caveats for this release.

- CSCti82166

Symptom: In a vPC set up with a Cisco Nexus 5000 switch running NX-OS 4.2(1)N2(1), when a vPC primary switch is reloaded with the **reload** command, the vPC secondary switch will not assume the vPC primary role and hence suspends all vPCs.

This is seen in a vPC set up which uses the Mgmt interface for **peer-keepalive**. The **reload** command is done on the vPC primary switch by a user logged into the switch using a non-admin username. In set ups using in-band peer-keepalives (SVI on the Cisco Nexus 5000), this issue is not seen.

This issue is seen because with a **reload** command, the vPC primary switch does not shut down the Mgmt interface first. Due to this, the vPC secondary switch only sees the peer-link go down but it still sees the vPC primary via the peer-keepalive link. Therefore, it does not assume vPC primary role and suspends all vPCs on the switch leading to traffic loss.

Workaround: If vPC primary switch needs to be reloaded for any reason, power cycle the box rather than using **reload** command. Alternately, from a console session on vPC primary switch, shut down the Mgmt interface first and then issue **reload** command without saving the configuration. If this is

Send documentation comments to nexus5kdocs@cisco.com

done by telnet/ssh to Mgmt interface, you will loose access to the switch. Then use inband peer-keepalives using SVIs on the Nexus 5000. and login using **admin** username. Downgrade NX-OS to 4.2(1)N1(1).

- CSCth69804

Symptom:

A Cisco Nexus 5000 switch running NX-OS 4.2(1)N2(1) configured for AAA, may prompt for a password when the **reload** command issued.

```
Lab-N5k# reload
WARNING: This command will reboot the system
Do you want to continue? (y/n) [n] y
We trust you have received the usual lecture from the local System Administrator. It
usually boils down to these two things:
    #1) Respect the privacy of others.
    #2) Think before you type.
Password:
Sorry, try again.
Password:
Sorry, try again.
Password:
Sorry, try again.
/usr/bin/sudo: 3 incorrect password attempts
```

Workaround: The switch does a reload after pressing Enter three times. Disable the AAA configuration and use local accounts defined on the switch for managing the switch. In a vPC peer set up, if **reload** is issued on the vPC primary, the vPC secondary switch does not take over smoothly. See related bug CSCti82166 for details of the vPC issue.

- CSCti18651

Symptom: A Cisco Nexus 5000 switch may crash when you run the following snmpwalk:

```
snmpwalk -v 2c -c <community> <ip> .1.3.6.1.4.1.9.9.225.1.3.1.1.1.369098852
SNMP has to be configured on the Cisco Nexus switch. The following is seen in the output of show logging nvram:
```

```

%$ VDC-1 %$ %SYSMGR-2-SERVICE_CRASHED: Service "eth_port_channel" (PID 3938) hasn't
caught signal 6 (core will be saved).
%$ VDC-1 %$ %SYSMGR-2-LAST_CORE_BASIC_TRACE: core_client_main: PID 6182 with message
filename = 0x101_eth_port_channel_log.3938.tar.gz .
```

And a core file will be created in 'show cores':

```
Nexus5000# show cores
Module-num      Instance-num    Process-name    PID      Core-create-time
-----
1               1              eth_port_channel 3938     Aug 3 16:10
```

Workaround: Avoid polling this specific object.

Resolved Caveats—Cisco NX-OS Release 4.2(1)N2(1)

This section lists the resolved caveats for this release.

- CSCtd54090

Symptom: A special NetAPP vendor specific an ELS frame that is sent after a FLOGI is not routed back. This puts the NetAPP port in a bad mode. This is triggered when the NetAPP CNA sends the special ELS frame after an FLOGI is complete.

Send documentation comments to nexus5kdocs@cisco.com

Workaround: None.

- CSCtd54245

Symptom: A Cisco Nexus 5000 Series switch that is configured with a san-port-channel to another Nexus or MDS product can see a gap in the switching of SAN traffic over the port-channel if a member is added to it. A gap of 300-500 msec can be seen on traffic going over a san-port-channel if a new member is added or a member that was down comes back operational.

Workaround: None.

- CSCte81499

Symptom: Basic System Configuration Dialog always rejects FC settings after write erase and reboot, as **feature fcoe** is not enabled yet.

Sample excerpt from basic configuration dialog with error

```
Enter basic FC configurations (yes/no) [n]: yes
Configure default physical FC switchport interface state (shut/noshut) [shut]:
Configure default physical FC switchport trunk mode (on/off/auto) [on]:
Configure default zone policy (permit/deny) [deny]: permit
Enable full zoneset distribution? (yes/no) [n]: yes
The following configuration will be applied:
...
system default switchport shutdown san
system default switchport trunk mode on
system default zone default-zone permit
system default zone distribute full
Would you like to edit the configuration? (yes/no) [n]:
Use this configuration and save it? (yes/no) [y]:
Error: There was an error executing atleast one of the command
Please verify the following log for the command execution errors.
Syntax error while parsing 'system default switchport shutdown san'
Syntax error while parsing 'system default switchport trunk mode on'
Syntax error while parsing 'system default zone default-zone permit'
Syntax error while parsing 'system default zone distribute full'

Would you like to save the running-config to startup-config? (yes/no) [n]: yes
[#####] 100%
```

This occurs if you select Yes to the option, **Enter basic FC configurations (yes/no) [n]:** during the basic configuration dialog.

Workaround: Enable **feature fcoe** in CLI after which you will be able to successfully configure the system default configuration as needed.

- CSCtg43468

Symptom: The **fc-port-security** configuration on a Nexus 5000 switch is missing when upgrading from the 4.1 release to the 4.2 release. If the admin has configured a port-security database for a vsan, the configuration will be lost after the Nexus 5000 switch reboots with the new release of software. The **show feature** will display the feature as enabled, but the configuration is not there.

Workaround: The admin has to re-configure FC port-security associations after upgrade from 4.1 to 4.2.

- CSCte99041

Symptom: When a Cisco Nexus 2000 Fabric Extender is connected to two Cisco Nexus 5000 Series Switches via a VPC, the server facing port configuration on the Cisco Nexus 2000 Fabric Extender must be configured the same on both the Cisco Nexus 5000 Series Switches. If the admin configures mismatching flowcontrol or speed on the server facing ports of the Cisco Nexus 2000 Fabric Extenders, then the port shows as link not connected. The error message is incorrect.

Send documentation comments to nexus5kdocs@cisco.com

Workaround: Configure matching configuration on both Cisco Nexus 5000 Series Switches for server facing ports on the Cisco Nexus 2000 Fabric Extender in this configuration.

- CSCtf07750

Symptom: Traps for bridge mib are not yet support on the Cisco Nexus 5000 Series Switches. Whereas in the Cisco Nexus 7000 Series, they are supported

```
agn11-5020(config)# snmp-server enable traps bridge?
```

```
% Invalid command at '^' marker.
```

```
n7k2(config)# snmp-server enable traps bridge ?
<CR>
newroot          Enable SNMP STP Bridge MIB newroot traps
topologychange  Enable SNMP STP Bridge MIB topologychange traps
```

Workaround: None.

- CSCtf40646

Symptom: The **show interface <> capabilities** shows incorrect speed for the Cisco Nexus 2248TP and the Cisco Nexus 2232PP server facing interfaces The Cisco Nexus 2248TP server facing ports only support 100Mbps and 1Gbps but the capability shows 10/100/1000Mbps. In 4.2(1)N1(1), the Cisco Nexus 2232PP only support 10Gbps on server facing interfaces but the capability shows both 1Gbps and 10Gbps.

Workaround: Ignore the non-applicable speeds

- CSCtf80675

Symptom: When the admin unbinds an interface from a VFC and binds it to a new VFC, the counter values are also copied over from the old VFC.

Workaround: Clear counters after rebinding

- CSCtb95741

Symptom: vPC configured on a port-channel with FEX host interface as a member does not come up. The **show vpc inconsistency** shows STP parameters as mismatched.

This happens when you configure port-channel on FEX host interface, configure spanning tree parameters (for example **port-type edge**) on the port-channel and configure vpc on the the port-channel.

Workaround: First, remove the vpc config on the port-channel and remove spanning tree parameters on the port-channel. Next, configure vpc on the port-channel and add spanning tree parameters, if needed.

Resolved Caveats—Cisco NX-OS Release 4.2(1)N1(1)

This section lists the resolved caveats for this release.

- CSCtd58753

Symptom: After adding & removing channel-group config on an mrouter port, multicast traffic flood to the port continues indefinitely, even after the PIM router or IGMP querier on the port has been removed (and the port is not an mrouter port anymore).

Workaround: Perform **shut, no shut** on the affected interface.

Send documentation comments to nexus5kdocs@cisco.com

- CSCtc09510

Symptom: On bootup of the switch, some of fex host interfaces go into BPDU errdisable state even when BPDU filter is enabled on that interface. This happens when BPDUs are received and processed on an interface, on bootup, before the BPDU filter configuration is applied.

Workaround: Bring the interface(s) administratively down and re-enable them.

```
switch# configure terminal
switch(config)# interface ethernet 109/1/1
switch(config-if)# shut
switch(config-if)# no shut
```
- 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.
- CSCsu48008

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

Workaround: Get the OperStatus from the CLI using the **show interface vfc x** command.
- CSCsu01188

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

Workaround: None.
- CSCta13997

Symptom: When vpc peer-link is down on Cisco Nexus 5000 Series switches, ports fail to come up on a Cisco Nexus 2000 Fabric Extender that is dual homed to the Nexus 5000 Series switches with vPC.

Workaround: Restore the peer-link before making new connections to the Cisco Nexus 2000 Fabric Extender.
- CSCsl62545

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

Workaround: None

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

The following are related Cisco Nexus 5000 Series documents:

- *The Cisco Nexus 5000 Series Switch CLI Software Configuration Guides*
- *Cisco Nexus 5000 Series Command Reference*
- *Cisco Nexus 5000 Series Hardware Installation Guide*

Send documentation comments to nexus5kdocs@cisco.com

- *Cisco MDS 9000 and Nexus 5000 Series Fabric Manager Software Configuration Guide, Cisco Fabric Manager*
- *Cisco Nexus 5000 Series and CiscoNexus 2000 Series MIB Quick Reference*

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 Fabric Extender Software Configuration Guide1*
- *Cisco Nexus 2000 Series Fabric Extender Hardware Installation 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.

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

©2010 Cisco Systems, Inc. All rights reserved.

♻️ Printed in the USA on recycled paper containing 10% postconsumer waste.