Cisco Nexus 7000 Series NX-OS Release Notes, Release 6.2

Date: May 25, 2015
Current Release: 6.2(12)

This document describes the features, caveats, and limitations for Cisco NX-OS software for use on the Cisco Nexus 7000 Series Switches. Use this document in combination with documents listed in the "Upgrade and Downgrade" section on page 253.

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 7000 Series NX-OS Release Notes, Release 6.x*:
http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-release-notes-list.html

Cisco Systems, Inc.
www.cisco.com

Table 1 shows the online change history for this document.

Table 1 Online History Change

Date	Description
August 22, 2013	Created release notes for Release 6.2(2).
August 26, 2013	• Added CSCuh18007to the "Open Caveats—Cisco NX-OS Release 6.2" section. • Added the "ECMP Support in Hardware" limitation. • Added the "CLI Command for Breakout Capabilities" limitation. • Added a caveat about FEX queuing to the "Upgrade or Downgrade Caveats" section. • Added IPv6 Inter-AS Option B lite as a supported MPLS feature.
September 5, 2013	• Added CSCug37851 to the "Resolved Caveats—Cisco NX-OS Release 6.2(2)" section. • Added CSCui13170 to the "Open Caveats—Cisco NX-OS Release 6.2" section. • Updated the supported FEX modules for the N77-F248XP-23E I/O module in Table 3. • Added vPC and vPC+ to the "New Software Features" section. • Updated OTV in the "New Software Features" section.
September 6, 2013	• Updated the "ISSU Upgrade Steps" section. • Updated "VACL Configuration Should Be Removed Before ISSU" in the "Upgrade or Downgrade Caveats" section. • Added CSCtf36357 to the "Resolved Caveats—Cisco NX-OS Release 6.2(2)" section. • Added Security Features to the "New Software Features" section.
September 10, 2013	• Modified the description of the "Behavior of Control Plane Packets on an F2e Series Module" limitation. • Added the "WCCP Support in a Mixed Mode VDC" limitation.
September 16, 2013	Updated vPC and vPC+ in the "New Software Features" section.
September 18, 2013	Added "Upgrade With an M2 Series Module Installed" to the "Upgrade/Downgrade Paths and Caveats" section.
September 30, 2013	Updated FabricPath in the "New Software Features" section.
October 16, 2013	Created release notes for Release 6.2(2a).
October 21, 2013	Added the Cisco Nexus B22 Fabric Extender for HP (N2K-B22HP) to Table 2 and Table 3.
October 22, 2013	Updated the "Security Features" section of the "Cisco NX-OS Release 6.2(2) Software Features" section.
October 30, 2013	Updated Table 3 to add support for Cisco Nexus B22 Fabric Extender for HP (N2K-B22HP) on the 48-port 1/10 Gigabit Ethernet SFP+ I/O F2 Series module (N7K-F248XP-25).
November 21, 2013	Updated the "Limitations" section to add "Behavior of Control Plane Packets on an F2e Series Module".

Table 1 Online History Change

Date	Description
December 4, 2013	• Updated Table 4 to list correct Product ID for the Cisco Nexus 7000 Series Network Analysis Module. • Updated the "Open Caveats—Cisco NX-OS Release 6.2" section to add CSCul81224.
December 10, 2013	Updated Table 9 to add support for Fe and F2e Series modules hardware as ERSPAN destinations.
December 20, 2013	Created the release notes for Release 6.2(6).
January 6, 2014	• Updated Table 3 to add FEX module support for the F3 Series modules in Release 6.2(6). • Updated Table 5 to add the F3 Series modules and supported transceivers for Release 6.2(6)
January 7, 2014	Added Physical Port vPC to the new features for Release 6.2(6).
January 7. 2014	Updated the "OTV Support"section and the "OTV" section under New Software Features for Release 6.2(6).
January 8, 2014	Added the Cisco Nexus B22HP Fabric Extender for HP (N2K-B22HP) to the "Cisco Nexus Fabric Extenders" section.
January 4, 2014	Updated the ISSU and ISSD information in the "Upgrade/Downgrade Paths and Caveats" section.
January 22, 2014	Updated the "Open Caveats—Cisco NX-OS Release 6.2" section to add CSCul91443.
January 23, 2014	Updated the "Interoperability Between Modules" information.
January 24, 2014	Updated the Updated the "Limitations" section to add "The no hardware ejector enable Command is Not Recommended for Long-Term Use"
February 5, 2014	• Updated the Cisco NX-OS Release 6.2(6) Software Features section to remove inaccurate FIPs certification feature from. • Updated Table 3 to add the F3-Series 12-port 40-Gigabit Ethernet SFP+ I/O module (N7K-F312FQ-25) for Cisco Nexus 7000 switches.
February 17, 2014	• Updated the "Features Available on F2, F2e, and F3 Series Modules" section to add MACSec support for the F2e Series modules. • Updated "Interoperability Between Modules" section to add that you cannot interoperate the F3 Series plus the F2 and/or F2e Series plus M2 Series in the same VDC.
February 20, 2014	• Updated "Non-ISSU Upgrade Steps" section.
February 26, 2014	Created the release notes for Release 6.2(6a).
March 10, 2014	Revised Table 5.
April 2, 2014	Removed support for ERSPAN destination for F3 Series modules.
April 3, 2014	Revised Tables 2 and 3.
April 8, 2014	Correct Cisco NX-OS Release to 6.2(6) for CLI command for breakout capabilities.
April 25, 2014	Created the release notes for Release 6.2(8).

Table 1 Online History Change

Date	Description
April 26, 2014	Added two additional caveats to the "Open Caveats—Cisco NX-OS Release 6.2" section.
April 28, 2014	Added additional caveats to the "Caveats" section.
May 1, 2014	Added additional caveats to the "Caveats" section.
May 6, 2014	Added additional caveat to the "Open Caveats—Cisco NX-OS Release 6.2"section.
May 12, 2014	Updated information for LISP in "Cisco NX-OS Release 6.2(8) Software Features" section.
May 13, 2014	Added list of modules and required level of software to "CLI Command for Breakout Capabilities" section.
May 22, 2014	Added additional caveat to the "Caveats" section.
June 3, 2014	Updated Table 9.
June 10, 2014	Updated Table 5.
June 12, 2014	Added additional caveat to the "Caveats" section.
June 25, 2014	Added the "Resolved Caveats—Cisco NX-OS Release 6.2(8a)" section.
July 7, 2014	Added the "Resolved Caveats—Cisco NX-OS Release 6.2(6b)" section. Updated the "Supported Upgrade and Downgrade Paths" to add Release 6.2(6b) and Release 6.2(8a).
July 25, 2014	Updated the "Supported Upgrade and Downgrade Paths" tables to add 6.1(4a). Updated the "VDC Ports Can Become Unallocated After a Downgrade" caveat in the "Upgrade or Downgrade Caveats" section. Added an additional caveat to the"Upgrade or Downgrade Caveats" section. Added CSCup65230 to the "Resolved Caveats—Cisco NX-OS Release 6.2(2)"section.
August 8, 2014	Added an additional caveat to the"Upgrade or Downgrade Caveats" section. Updated the "Limitations" section to add "Native VLAN Change Causes Link Flap". Added CSCug71801 to the "Resolved Caveats—Cisco NX-OS Release 6.2(2)"section.
September 11, 2014	Added CSCui72592 and CSCuo93631 to the "Open Caveats—Cisco NX-OS Release 6.2" section.
October 20, 2014	Created the release notes for Release 6.2(10).
November 5, 2014	Added CSCul05775 to "Open Caveats—Cisco NX-OS Release 6.2" section.
November 11, 2014	Added "Resolved Caveats—Cisco NX-OS Release 6.2(8b)" section.
November 14, 2014	Added CSCur57579 to "Open Caveats—Cisco NX-OS Release 6.2" section.
November 17, 2014	Updated the "Cisco NX-OS Release 6.2(10) Software Features" section to add "License Requirement MP-BGP RR for FabricPath (aka DFA)."
February 3, 2015	Created the release notes for Release 6.2(12).
February 4, 2015	Added CSCus76724 to "Open Caveats—Cisco NX-OS Release 6.2"section.

Table 1 Online History Change

Date	Description
February 12, 2015	Updates to "Open Caveats—Cisco NX-OS Release 6.2" and "Resolved Caveats—Cisco NX-OS Release 6.2(12)" sections. Updated the "Supported Upgrade and Downgrade Paths".
March 2, 2015	Removed "Not Customer Visible" Bugs from the "Caveats" section.
April 23, 2015	Added CSCub54436 to the "Resolved Caveats—Cisco NX-OS Release 6.2(2)"section.
May 20, 2015	Added CSCuq39448 to the "Open Caveats—Cisco NX-OS Release 6.2"section.
May 25, 2015	Added CSCuo82450 to the "Resolved Caveats—Cisco NX-OS Release 6.2(12)" sections.

Contents

This document includes the following sections:

Introduction

The Cisco NX-OS software for the Cisco Nexus 7000 Series fulfills the routing, switching, and storage networking requirements of data centers and provides an Extensible Markup Language (XML) interface and a command-line interface (CLI) similar to Cisco IOS software.

System Requirements

This section includes the following topics:

- Memory Requirements, page 6
- Supported Device Hardware, page 6

Memory Requirements

Cisco NX-OS Release 6.2 software requires 8 GB of memory. If you have a Cisco Nexus 7000 Series system with a Supervisor 1 module with 4 GB of memory, you must upgrade to 8 GB of memory using the memory upgrade kit, N7K-SUP1-8GBUPG=, before you install Cisco NX-OS Release 6.2.

Instructions for upgrading to the new memory are available in the “Upgrading Memory for Supervisor Modules” section of the *Cisco Nexus 7000 Series Hardware Installation and Reference Guide.*

Supported Device Hardware

The Cisco NX-OS software supports the Cisco Nexus 7000 Series that includes Cisco Nexus 7000 switches and Cisco Nexus 7700 switches. You can find detailed information about supported hardware in the *Cisco Nexus 7000 Series Hardware Installation and Reference Guide.*

Table 2 shows the Cisco Nexus 7000 Series hardware supported by Cisco NX-OS Release 6.x, Release 5.x, and Release 4.x software.

Table 3 shows the FEX modules supported by the Cisco Nexus 7000 Series I/O modules.

Table 4 shows the service modules supported by the Cisco Nexus 7000 Series switches.

Table 5 shows the transceiver devices supported by each release.

For a list of minimum recommended Cisco NX-OS software releases for use with Cisco Nexus 7000 Series switches, see the document *Minimum Recommended Cisco NX-OS Releases for Cisco Nexus 7000 Series Switches.*

Table 2 Cisco Nexus 7000 Series Hardware Supported by Cisco NX-OS Software

Product ID	Hardware	Minimum Software Release
N7K-C7004	Cisco Nexus 7004 chassis	6.1(2)
N7K-C7009	Cisco Nexus 7009 chassis	5.2(1)
N7K-C7010	Cisco Nexus 7010 chassis	4.0(1)
N7K-C7018	Cisco Nexus 7018 chassis	4.1(2)
N7K-C7004-FAN=	Replacement fan for the Cisco Nexus 7004 chassis	6.1(2)
N7K-C7009-FAN	Replacement fan for the Cisco Nexus 7009 chassis	5.2(1)
N7K-C7010-FAN-S	System fan tray for the Cisco Nexus 7010 chassis	4.0(1)
N7K-C7010-FAN-F	Fabric fan tray for the Cisco Nexus 7010 chassis	4.0(1)

Table 2 Cisco Nexus 7000 Series Hardware Supported by Cisco NX-OS Software (continued)

Product ID	Hardware	Minimum Software Release
N7K-C7018-FAN	Fan tray for the Cisco Nexus 7018 chassis	4.1(2)
N7K-AC-3KW	3.0-kW AC power supply unit	6.1(2)
N7K-DC-3KW	3.0-kW DC power supply unit	6.1(2)
N7K-AC-6.0KW	6.0-kW AC power supply unit	4.0(1)
N7K-AC-7.5KW-INT N7K-AC-7.5KW-US	7.5-kW AC power supply unit	4.1(2) 4.1(2)
N7K-DC-6.0KW N7K-DC-PIU N7K-DC-CAB=	6.0-kW DC power supply unit (cable included) DC power interface unit DC 48 V, -48 V cable (spare)	5.0(2) 5.0(2) 5.0(2)
N7K-SUP2E	Supervisor 2 Enhanced module	6.1(1)
N7K-SUP2	Supervisor 2 module	6.1(1)
N7K-SUP1	Supervisor 1 module	4.0(1)
N7K-SUP1-8GBUPG	Supervisor module memory kit upgrade	5.1(1)
N7K-C7009-FAB-2	Fabric module, Cisco Nexus 7000 Series 9-slot	5.2(1)
N7K-C7010-FAB-2	Fabric module, Cisco Nexus 7000 Series 10-slot	6.0(1)
N7K-C7010-FAB-1	Fabric module, Cisco Nexus 7000 Series 10-slot	4.0(1)
N7K-C7018-FAB-2	Fabric module, Cisco Nexus 7000 Series 18-slot	6.0(1)
N7K-C7018-FAB-1	Fabric module, Cisco Nexus 7000 Series 18-slot	4.1(2)
N7K-F312FQ-25	Cisco Nexus 7000 12-port 40-Gigabit Ethernet QSFP+ I/O module (F3 Series)	6.2(6)
N7K-F306CK-25	Cisco Nexus 7000 6-port 100-Gigabit Ethernet CPAK I/O module (F3 Series)	6.2(10)
N7K-F348XP-25	Cisco Nexus 7000 48-port 1/10-Gigabit Ethernet SFP+ I/O module (F3 Series)	6.2(12)
N7K-F248XT-25E	Enhanced 48-port 1/10 GBASE-T RJ45 module (F2e Series)	6.1(2)
N7K-F248XP-25E	Enhanced 48-port 1/10 Gigabit Ethernet SFP+ I/O module (F2e Series)	6.1(2)

Table 2 Cisco Nexus 7000 Series Hardware Supported by Cisco NX-OS Software (continued)

Product ID	Hardware	Minimum Software Release
N7K-F248XP-25	48-port 1/10 Gigabit Ethernet SFP+ I/O module (F2 Series)	6.0(1)
N7K-F132XP-15	32-port 1/10 Gigabit Ethernet module (F1 Series)	5.1(1)
N7K-M202CF-22L	2-port 100-Gigabit Ethernet I/O module XL (M2 Series)	6.1(1)
N7K-M206FQ-23L	6-port 40-Gigabit Ethernet I/O module XL (M2 Series)	6.1(1)
N7K-M224XP-23L	24-port 10-Gigabit Ethernet I/O module XL (M2 Series)	6.1(1)
N7K-M108X2-12L	8-port 10-Gigabit Ethernet I/O module XL[1]	5.0(2)
N7K-M132XP-12	32-port 10-Gigabit Ethernet SFP+ I/O module	4.0(1)
N7K-M132XP-12L	32-port 10-Gigabit Ethernet SFP+ I/O module XL[1]	5.1(1)
N7K-M148GS-11	48-port 1-Gigabit Ethernet SFP I/O module	4.1(2)
N7K-M148GS-11L	48-port 1-Gigabit Ethernet I/O module XL[1]	5.0(2)
N7K-M148GT-11	48-port 10/100/1000 Ethernet I/O module	4.0(1)
N7K-M148GT-11L	48-port 10/100/1000 Ethernet I/O module XL[1]	5.1(2)
N77-C7706	Cisco Nexus 7706 chassis	6.2(6)
N77-C7718	Cisco Nexus 7718 chassis	6.2(2)
N77-C7710	Cisco Nexus 7710 chassis	6.2(2)
N77-C7718-FAN	Fan, Cisco Nexus 7718 chassis	6.2(2)
N77-C7710-FAN	Fan, Cisco Nexus 7710 chassis	6.2(2)
N77-AC-3KW	Cisco Nexus 7700 AC power supply	6.2(2)
N77-DC-3KW	Cisco Nexus 7700 DC power supply	6.2(2)
N77-C7706-FAB-2	Fabric Module, Cisco Nexus 7706 chassis	6.2(6)
N77-C7718-FAB-2	Fabric Module, Cisco Nexus 7718 chassis	6.2(2)
N77-C7710-FAB-2	Fabric Module, Cisco Nexus 7710 chassis	6.2(2)
N77-SUP2E	Cisco Nexus 7700 Supervisor 2 Enhanced module	6.2(2)

Table 2 Cisco Nexus 7000 Series Hardware Supported by Cisco NX-OS Software (continued)

Product ID	Hardware	Minimum Software Release
N77-F348XP-23	Cisco Nexus 7700 48-port 1/10-Gigabit Ethernet SFP+ I/O module (F3 Series)	6.2(6)
N77-F324FQ-25	Cisco Nexus 7700 24-port 40-Gigabit Ethernet QSFP+ I/O module (F3 Series)	6.2(6)
N77-F312CK-26	Cisco Nexus 7700 12-port 100-Gigabit Ethernet CPAK I/O module (F3 Series)	6.2(6)
N77-F248XP-23E	Cisco Nexus 7700 Enhanced 48-port 1/10 Gigabit Ethernet SFP+ I/O module (F2e Series)	6.2(2)
N2K-C2248TP-1GE	Cisco Nexus 2248TP Fabric Extender	5.2(1)
N2K-C2224TP-1GE	Cisco Nexus 2224TP Fabric Extender	5.2(1)
N2K-C2248TP-E	Cisco Nexus 2224TP Fabric Extender	6.1(1)
N2K-C2232PP-10GE	Cisco Nexus 2232PP Fabric Extender	5.2(1)
N2K-C2232TM	Cisco Nexus 2232TM Fabric Extender	6.1(1)
N2K-C2232TM-E	Cisco Nexus 2232TM Fabric Extender	6.2(2)
N2K-C2248PQ	Cisco Nexus 2248PQ Fabric Extender	6.2(2)
N2K-B22HP	Cisco Nexus B22 Fabric Extender for HP	6.2(2)

1. Requires the Cisco Nexus 7010 Scalable Feature Package license (N7K-C7010-XL) or the Cisco Nexus 7018 Scalable Feature Package license (N7K-C7018-XL), depending on the chassis, to enable all XL-capable I/O modules to operate in XL mode.

Table 3 FEX Modules Supported by Cisco Nexus 7000 Series Modules

Cisco Nexus 7000 Series Module	FEX Module	Minimum Software Release
12-port 40-Gigabit Ethernet F3-Series QSFP I/O module (N7K-F312FQ-25) for Cisco Nexus 7000 Series switches	N2K-C2224TP-1GE N2K-C2248TP-1GE N2K-C2232PP-10GE N2K-C2232TM N2K-C2248TP-E N2K-C2232TM-E N2K-C2248PQ N2K-B22HP[1]	6.2(6)
24-port Cisco Nexus 7700 F3 Series 40-Gigabit Ethernet QSFP I/O module (N77-F324FQ-25)	N2K-C2224TP-1GE N2K-C2248TP-1GE N2K-C2232PP-10GE N2K-C2232TM N2K-C2248TP-E N2K-C2232TM-E N2K-C2248PQ N2K-B22HP[1]	6.2(8)
48-port Cisco Nexus 7700 F3 Series 1/10-Gigabit Ethernet SFP+ I/O module (N77-F348XP-23)	N2K-C2224TP-1GE N2K-C2248TP-1GE N2K-C2232PP-10GE N2K-C2232TM N2K-C2248TP-E N2K-C2232TM-E N2K-C2248PQ N2K-B22HP	6.2(6)
32-port 10-Gigabit Ethernet SFP+ I/O module (N7K-M132XP-12) 32-port 10-Gigabit Ethernet SFP+ I/O module (N7K-M132XP-12L)	N2K-C2248TP-1GE	5.2(1)
	N2K-C2224TP-1GE N2K-C2232PP-10GE	5.2(1)
	N2K-C2232TM N2K-C2248TP-E	6.1(1)
	N2K-2232TM-E N2K-C2248PQ N2K-B22HP	6.2(2)

Table 3 FEX Modules Supported by Cisco Nexus 7000 Series Modules (continued)

Cisco Nexus 7000 Series Module	FEX Module	Minimum Software Release
24-port 10-Gigabit Ethernet I/O M2 Series module XL (N7K-M224XP-23L)	N2K-C2224TP-1GE N2K-C2248TP-1GE N2K-C2232PP-10GE N2K-C2232TM N2K-C2248TP-E	6.1(1)
	N2K-C2232TM-E N2K-C2248PQ N2K-B22HP	6.2(2)
48-port 1/10 Gigabit Ethernet SFP+ I/O F2 Series module (N7K-F248XP-25)	N2K-C2224TP-1GE N2K-C2248TP-1GE N2K-C2232PP-10GE	6.0(1)
	N2K-C2232TM N2K-C2248TP-E	6.1(1)
	N2K-2232TM-E N2K-2248PQ N2K-B22HP	6.2(2)
Enhanced 48-port 1/10 Gigabit Ethernet SFP+ I/O module (F2e Series) (N7K-F248XP-25E)	N2K-C2224TP-1GE N2K-C2248TP-1GE N2K-C2232PP-10GE N2K-C2232TM N2K-C2248TP-E	6.1(2)
	N2K-2232TM-E N2K-C2248PQ N2K-B22HP	6.2(2)
48-port 1/10 Gigabit Ethernet SFP+ I/O module (F2e Series) (N77-F248XP-23E)	N2K-C2224TP-1GE N2K-C2248TP-1GE N2K-C2232PP-10GE N2K-C2232TM N2K-C2232TM-E N2K-C2248PQ N2K-C2248TP-E N2K-B22HP	6.2(2)

[1] FEX server-facing interfaces should be configured in autonegotiate mode. Do not force a specific data rate. See DDTS CSCuj84520 for additional information.

Note The Cisco Nexus 7000 Enhanced F2 Series 48-port 1/10 GBASE-T RJ-45 Module (N7K-F248XT-25E) does not support Cisco Nexus 2000 Fabric Extenders.

Table 4 Service Modules Supported by Cisco Nexus 7000 Series Switches

Service Module	Product ID	Minimum Software Release
Cisco Nexus 7000 Series Network Analysis Module	NAM-NX1	6.2(2)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases

I/O Module	Product ID	Transceiver Type	Minimum Software Version
N77-F312CK-26	CPAK-100G-LR4	Cisco 100GBASE-LR4 CPAK	6.2(6)
	CPAK-100G-SR10	Cisco 100GBASE-SR10 CPAK	6.2(6)
N77-F324FQ-25	QSFP-40G-SR-BD	Cisco 40G BiDi QSFP+	6.2(6)
	QSFP-40G-SR4 QSFP-40G-SR4-S	40GBASE-SR4 QSFP+	6.2(6)
	QSFP-40G-CSR4	40GBASE-CSR4 QSFP+	6.2(6)
	QSFP-40GE-LR4 QSFP-40G-LR4-S	40GBASE-LR4 QSFP+	6.2(6)
	FET-40G	Cisco 40G Fabric Extender Transceiver (FET)	6.2(8)
	QSFP-H40G-ACUxM	40GBASE-CR4 QSFP+ Direct Attach Copper Cable Active (7 m, 10 m)	6.2(8)
	QSFP-4X10G-ACxM	40GBASE-CR4 QSFP+ to 4x SFP+ Twinax Direct Attach Copper Breakout Cable Active (7 m, 10 m)	6.2(8)
	QSFP-H40G-AOCxM	40GBASE-AOC (Active Optical Cable) QSFP Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m, 15 m)	6.2(8)
	QSFP-4X10G-AOCxM	40GBASE-AOC (Active Optical Cable) QSFP to 4x10G SFP+ Breakout Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(8)
	WSP-Q40GLR4L	40GBASE-LR4 lite (2km SMF) QSFP+	6.2(10)
	QSFP-40G-LR4	40GBASE-LR4 QSFP+ (Ethernet and OTU3 capable)	6.2(12)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	QSFP-40G-ER4	40GBASE-ER4 QSFP+ (40km)	6.2(12)
N77-F348XP-23	CWDM-SFP-xxxx[1]	1000BASE-CWDM	6.2(8)
	DWDM-SFP-xxxx[1]	1000BASE-DWDM	6.2(8)
	GLC-TE	1000BASE-T SFP	6.2(10)
	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.2(6)
	SFP-10G-AOCxM	110GBASE-AOC (Active Optical Cable) SFP+ Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(10)
	SFP-10G-SR SFP-10G-SR-S	10GBASE-SR SFP+	6.2(6)
	SFP-10G-LR SFP-10G-LR-S	10GBASE-LR SFP+	6.2(6)
	SFP-10G-ER SFP-10G-ER-S	10GBASE-ER SFP+	6.2(6)
	SFP-10G-ZR SFP-10G-ZR-S	10GBASE-ZR SFP+	6.2(6)
	DWDM-SFP10G-xx.xx	10GBASE-DWDM SFP+	6.2(6)
	SFP-10G-LRM[2]	10GBASE-LRM SFP+	6.2(8)
	SFP-H10GB-CUxM	SFP-H10GB-CUxM Twinax Cable Passive (1 m, 3 m, 5 m)	6.2(8)
	SFP-H10GB-CUxM	SFP-H10GC-CUxM Twinax Cable Passive (1.5 m, 2 m, 2.5 m)	6.2(8)
	SFP-H10GB-ACUxM	SFP-H10GB-ACUxM Twinax Cable Active (7 m, 10 m)	6.2(8)
	SFP-GE-T	1000BASE-T SFP	6.2(8)
	SFP-GE-S	1000BASE-SX SFP (DOM)	6.2(8)
	SFP-GE-L	1000BASE-LX/LH SFP (DOM)	6.2(8)
	SFP-GE-Z	1000BASE-ZX SFP (DOM)	6.2(8)
	GLC-LH-SM	1000BASE-LX/LH SFP	6.2(8)
	GLC-LH-SMD	1000BASE-LX/LH SFP	6.2(8)
	GLC-SX-MM	1000BASE-SX SFP	6.2(8)
	GLC-SX-MMD	1000BASE-SX SFP	6.2(8)
	GLC-ZX-SM	1000BASE-ZX SFP	6.2(8)
	GLC-ZX-SMD	1000BASE-ZX SFP	6.2(8)
	GLC-T	1000BASE-T SFP	6.2(8)
	GLC-BX-D	1000BASE-BX10-D	6.2(8)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	GLC-BX-U	1000BASE-BX10-U	6.2(8)
	GLC-EX-SMD	1000BASE-EX SFP	6.2(8)
	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.2(8)
N7K-F348XP-25	CWDM-SFP-xxxx[1]	1000BASE-CWDM	6.2(12)
	DWDM-SFP-xxxx[1]	1000BASE-DWDM	6.2(12)
	GLC-TE	1000BASE-T SFP	6.2(12)
	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.2(12)
	SFP-10G-AOCxM	110GBASE-AOC (Active Optical Cable) SFP+ Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(12)
	SFP-10G-SR SFP-10G-SR-S	10GBASE-SR SFP+	6.2(12)
	SFP-10G-LR SFP-10G-LR-S	10GBASE-LR SFP+	6.2(12)
	SFP-10G-ER SFP-10G-ER-S	10GBASE-ER SFP+	6.2(12)
	SFP-10G-ZR SFP-10G-ZR-S	10GBASE-ZR SFP+	6.2(12)
	DWDM-SFP10G-xx.xx	10GBASE-DWDM SFP+	6.2(12)
	SFP-10G-LRM[2]	10GBASE-LRM SFP+	6.2(12)
	SFP-H10GB-CUxM	SFP-H10GB-CUxM Twinax Cable Passive (1 m, 3 m, 5 m)	6.2(12)
	SFP-H10GB-CUxM	SFP-H10GC-CUxM Twinax Cable Passive (1.5 m, 2 m, 2.5 m)	6.2(12)
	SFP-H10GB-ACUxM	SFP-H10GB-ACUxM Twinax Cable Active (7 m, 10 m)	6.2(12)
	SFP-GE-T	1000BASE-T SFP	6.2(12)
	SFP-GE-S	1000BASE-SX SFP (DOM)	6.2(12)
	SFP-GE-L	1000BASE-LX/LH SFP (DOM)	6.2(12)
	SFP-GE-Z	1000BASE-ZX SFP (DOM)	6.2(12)
	GLC-LH-SM	1000BASE-LX/LH SFP	6.2(12)
	GLC-LH-SMD	1000BASE-LX/LH SFP	6.2(12)
	GLC-SX-MM	1000BASE-SX SFP	6.2(12)
	GLC-SX-MMD	1000BASE-SX SFP	6.2(12)
	GLC-ZX-SM	1000BASE-ZX SFP	6.2(12)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	GLC-ZX-SMD	1000BASE-ZX SFP	6.2(12)
	GLC-T	1000BASE-T SFP	6.2(12)
	GLC-BX-D	1000BASE-BX10-D	6.2(12)
	GLC-BX-U	1000BASE-BX10-U	6.2(12)
	GLC-EX-SMD	1000BASE-EX SFP	6.2(12)
	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.2(12)
N7K-F312FQ-25	QSFP-40G-SR-BD	Cisco 40G BiDi QSFP+	6.2(6)
	QSFP-40G-SR4 QSFP-40G-SR4-S	40GBASE-SR4 QSFP+	6.2(6)
	QSFP-40G-CSR4	40GBASE-CSR4 QSFP+	6.2(6)
	QSFP-40GE-LR4 QSFP-40G-LR4-S	40GBASE-LR4 QSFP+	6.2(6)
	FET-40G	Cisco 40G Fabric Extender Transceiver (FET)	6.2(6)
	QSFP-H40G-ACUxM	40GBASE-CR4 QSFP+ Direct Attach Copper Cable Active (7 m, 10 m)	6.2(8)
	QSFP-4X10G-ACxM	40GBASE-CR4 QSFP+ to 4x SFP+ Twinax Direct Attach Copper Breakout Cable Active (7 m, 10 m)	6.2(8)
	QSFP-H40G-AOCxM	40GBASE-AOC (Active Optical Cable) QSFP Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m, 15 m)	6.2(8)
	QSFP-4X10G-AOCxM	40GBASE-AOC (Active Optical Cable) QSFP to 4x10G SFP+ Breakout Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(8)
	WSP-Q40GLR4L	40GBASE-LR4 lite (2km SMF) QSFP+	62(10)
	QSFP-40G-LR4	40GBASE-LR4 QSFP+ (Ethernet and OTU3 capable)	6.2(12)
	QSFP-40G-ER4	40GBASE-ER4 QSFP+ (40km)	6.2(12)
N7K-F306CK-25	CPAK-100G-LR4	Cisco 100GBASE-LR4 CPAK	6.2(10)
	CPAK-100G-SR10	Cisco 100GBASE-SR10 CPAK	6.2(10)
N77-F248XP-23E	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.2(2)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	SFP-10G-AOCxM	10GBASE-AOC (Active Optical Cable) SFP+ Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(2)
	SFP-10G-SR SFP-10G-SR-S	10GBASE-SR SFP+	6.2(2)
	SFP-10G-LR SFP-10G-LR-S	10GBASE-LR SFP+	6.2(2)
	SFP-10G-ER SFP-10G-ER-S	10GBASE-ER SFP+	6.2(2)
	SFP-10G-LRM	10GBASE-LRM SFP+	6.2(2)
	SFP-10G-ZR[3] SFP-10G-ZR-S	10GBASE-ZR SFP+	6.2(2)
	SFP-H10GB-CUxM	SFP-H10GB-CUxM Twinax Cable Passive (1 m, 3 m, 5 m)	6.2(2)
	SFP-H10GB-CUxM	SFP-H10GC-CUxM Twinax Cable Passive (1.5 m, 2 m, 2.5 m)	6.2(2)
	SFP-H10GB-ACUxM	SFP-H10GB-ACUxM Twinax Cable Active (7 m, 10 m)	6.2(2)
	SFP-GE-T	1000BASE-T SFP	6.2(2)
	SFP-GE-S	1000BASE-SX SFP (DOM)	6.2(2)
	SFP-GE-L	1000BASE-LX/LH SFP (DOM)	6.2(2)
	SFP-GE-Z	1000BASE-ZX SFP (DOM)	6.2(2)
	GLC-LH-SM	1000BASE-LX/LH SFP	6.2(2)
	GLC-LH-SMD	1000BASE-LX/LH SFP	6.2(2)
	GLC-SX-MM	1000BASE-SX SFP	6.2(2)
	GLC-SX-MMD	1000BASE-SX SFP	6.2(2)
	GLC-ZX-SM	1000BASE-ZX SFP	6.2(2)
	GLC-ZX-SMD	1000BASE-ZX SFP	6.2(2)
	GLC-T	1000BASE-T SFP	6.2(2)
	GLC-TE	1000BASE-T SFP	6.2(10)
	GLC-BX-D	1000BASE-BX10-D	6.2(2)
	GLC-BX-U	1000BASE-BX10-U	6.2(2)
	GLC-EX-SMD	1000BASE-EX SFP	6.2(2)
	CWDM-SFP-xxxx[1]	1000BASE-CWDM	6.2(2)
	DWDM-SFP10G-xx.xx[1]	10GBASE-DWDM SFP+	6.2(2)
	DWDM-SFP-xxxx[1]	1000BASE-DWDM	6.2(2)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
N7K-F248XP-25	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.0(1)
	SFP-10G-SR SFP-10G-SR-S	10GBASE-SR SFP+	6.0(1)
	SFP-10G-LR SFP-10G-LR-S	10GBASE-LR SFP+	6.0(1)
	SFP-10G-ER SFP-10G-ER-S	10GBASE-ER SFP+	6.0(1)
	SFP-10G-LRM	10GBASE-LRM SFP+	6.0(1)
	SFP-10G-ZR[3] SFP-10G-ZR-S	10GBASE-ZR SFP+	6.1(1)
	SFP-10G-AOCxM	10GBASE-AOC (Active Optical Cable) SFP+ Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(2)
	SFP-H10GB-CUxM	SFP-H10GB-CUxM Twinax Cable Passive (1 m, 3 m, 5 m)	6.0(1)
	SFP-H10GB-CUxM	SFP-H10GC-CUxM Twinax Cable Passive (1.5 m, 2 m, 2.5 m)	6.2(2)
	SFP-H10GB-ACUxM	SFP-H10GB-ACUxM Twinax Cable Active (7 m, 10 m)	6.0(1)
	SFP-GE-T	1000BASE-T SFP	6.0(1)
	SFP-GE-S	1000BASE-SX SFP (DOM)	6.0(1)
	SFP-GE-L	1000BASE-LX/LH SFP (DOM)	6.0(1)
	SFP-GE-Z	1000BASE-ZX SFP (DOM)	6.0(1)
	GLC-TE	1000BASE-T SFP	6.2(10)
	GLC-LH-SM	1000BASE-LX/LH SFP	6.0(1)
	GLC-LH-SMD	1000BASE-LX/LH SFP	6.0(1)
	GLC-SX-MM	1000BASE-SX SFP	6.0(1)
	GLC-SX-MMD	1000BASE-SX SFP	6.0(1)
	GLC-ZX-SM	1000BASE-ZX SFP	6.0(1)
	GLC-ZX-SMD	1000BASE-ZX SFP	6.2(2)
	GLC-T	1000BASE-T SFP	6.0(1)
	GLC-BX-D	1000BASE-BX10-D	6.0(1)
	GLC-BX-U	1000BASE-BX10-U	6.0(1)
	GLC-EX-SMD	1000BASE-EX SFP	6.1(1)
	CWDM-SFP-xxxx[1]	1000BASE-CWDM	6.0(1)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	DWDM-SFP10G-xx.xx[1]	10GBASE-DWDM SFP+	6.1(1)
	DWDM-SFP-xxxx[1]	1000BASE-DWDM	6.0(1)
N7K-F248XP-25E	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.1(2)
	SFP-10G-SR SFP-10G-SR-S	10GBASE-SR SFP+	6.1(2)
	SFP-10G-LR SFP-10G-LR-S	10GBASE-LR SFP+	6.1(2)
	SFP-10G-ER SFP-10G-ER-S	10GBASE-ER SFP+	6.1(2)
	SFP-10G-LRM	10GBASE-LRM SFP+	6.1(2)
	SFP-10G-ZR[3] SFP-10G-ZR-S	10GBASE-ZR SFP+	6.1(2)
	SFP-10G-AOCxM	10GBASE-AOC (Active Optical Cable) SFP+ Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(2)
	SFP-H10GB-CUxM	SFP-H10GB-CUxM Twinax Cable Passive (1 m, 3 m, 5 m)	6.1(2)
	SFP-H10GB-CUxM	SFP-H10GC-CUxM Twinax Cable Passive (1.5 m, 2 m, 2.5 m)	6.2(2)
	SFP-H10GB-ACUxM	SFP-H10GB-ACUxM Twinax Cable Active (7 m, 10 m)	6.1(2)
	SFP-GE-T	1000BASE-T SFP	6.1(2)
	SFP-GE-S	1000BASE-SX SFP (DOM)	6.1(2)
	SFP-GE-L	1000BASE-LX/LH SFP (DOM)	6.1(2)
	SFP-GE-Z	1000BASE-ZX SFP (DOM)	6.1(2)
	GLC-LH-SM	1000BASE-LX/LH SFP	6.1(2)
	GLC-LH-SMD	1000BASE-LX/LH SFP	6.1(2)
	GLC-SX-MM	1000BASE-SX SFP	6.1(2)
	GLC-SX-MMD	1000BASE-SX SFP	6.1(2)
	GLC-ZX-SM	1000BASE-ZX SFP	6.1(2)
	GLC-ZX-SMD	1000BASE-ZX SFP	6.1(2)
	GLC-T	1000BASE-T SFP	6.1(2)
	GLC-TE	1000BASE-T SFP	6.2(10)
	GLC-BX-D	1000BASE-BX10-D	6.1(2)
	GLC-BX-U	1000BASE-BX10-U	6.1(2)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	GLC-EX-SMD	1000BASE-EX SFP	6.1(2)
	CWDM-SFP-xxxx[1]	1000BASE-CWDM	6.1(2)
	DWDM-SFP10G-xx.xx[1]	10GBASE-DWDM SFP+	6.1(2)
	DWDM-SFP-xxxx[1]	1000BASE-DWDM	6.1(2)
N7K-F132XP-15	SFP-10G-SR SFP-10G-SR-S	10GBASE-SR SFP+	5.2(1)
	GLC-TE	1000BASE-T SFP	6.2(10)
	SFP-10G-LR SFP-10G-LR-S	10GBASE-LR SFP+	5.1(1)
	SFP-10G-ER SFP-10G-ER-S	10GBASE-ER SFP+	5.1(1)
	SFP-10G-LRM	10GBASE-LRM SFP+	5.1(1)
	SFP-10G-ZR[3] SFP-10G-ZR-S	10GBASE-ZR SFP+	6.1(1)
	SFP-10G-AOCxM	10GBASE-AOC (Active Optical Cable) SFP+ Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(2)
	SFP-H10GB-CUxM	SFP-H10GB-CUxM Twinax Cable Passive (1 m, 3 m, 5 m)	5.1(1)
	SFP-H10GB-CUxM	SFP-H10GC-CUxM Twinax Cable Passive (1.5 m, 2 m, 2.5 m)	6.2(2)
	SFP-H10GB-ACUxM	SFP-H10GB-ACUxM Twinax Cable Active (7 m, 10 m)	5.1(1)
	SFP-GE-T	1000BASE-T SFP	5.1(1)
	SFP-GE-S	1000BASE-SX SFP (DOM)	5.1(1)
	SFP-GE-L	1000BASE-LX/LH SFP (DOM)	5.1(1)
	SFP-GE-Z	1000BASE-ZX SFP (DOM)	5.1(1)
	GLC-LH-SM	1000BASE-LX/LH SFP	5.1(1)
	GLC-SX-MM	1000BASE-SX SFP	5.1(1)
	GLC-ZX-SM	1000BASE-ZX SFP	5.1(1)
	GLC-ZX-SMD	1000BASE-ZX SFP	6.2(2)
	GLC-T	1000BASE-T SFP	5.1(1)
	GLC-LH-SMD	1000BASE-LX/LH SFP	5.2(1)
	GLC-SX-MMD	1000BASE-SX SFP	5.2(1)
	GLC-EX-SMD	1000BASE-EX-SFP	6.1(1)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	DWDM-SFP10G-xx.xx [1]	10-GBASE-DWDM SFP+	6.1(1)
N7K-M108X2-12L	SFP-10G-SR[4] SFP-10G-SR-S	10GBASE-SR SFP+	5.2(3a)
	SFP-10G-LR[4] SFP-10G-LR-S	10GBASE-LR SFP+	5.2(3a)
	SFP-10G-LRM[4]	10GBASE-LRM SFP+	5.2(1)
	SFP-H10GB-CUxM[4]	SFP-H10GB-CUxM Twinax Cable Passive (1 m, 3 m, 5 m)	5.2(1)
	CVR-X2-SFP10G	OneX Converter Module - X2 to SFP+ Adapter	5.2(1)
	X2-10GB-CX4	10GBASE-CX4 X2	5.1(1)
	X2-10GB-ZR	10GBASE-ZR X2	5.1(1)
	X2-10GB-LX4	10GBASE-LX4 X2	5.1(1)
	X2-10GB-SR	10GBASE-SR X2	5.0(2a)
	X2-10GB-LR	10GBASE-LRX2	5.0(2a)
	X2-10GB-LRM	10GBASE-LRM X2	5.0(2a)
	X2-10GB-ER	10GBASE-ERX2	5.0(2a)
	DWDM-X2-xx.xx=[1]	10GBASE-DWDM X2	5.0(2a)
N7K-M148GS-11	SFP-GE-S	1000BASE-SX	4.1(2)
	GLC-SX-MM		4.1(2)
	SFP-GE-L	1000BASE-LX	4.1(2)
	GLC-LH-SM		4.1(2)
	SFP-GE-Z	1000BASE-ZX	4.1(2)
	GLC-ZX-SM		4.1(2)
	GLC-EX-SMD	1000BASE-EX SFP	6.2(2)
	GLC-ZX-SMD	1000BASE-ZX SFP	6.2(2)
	GLC-T	1000BASE-T	4.2(1)
	SFP-GE-T		4.2(1)
	GLC-BX-D	1000BASE-BX10-D	5.2(1)
	GLC-BX-U	1000BASE-BX10-U	5.2(1)
	GLC-SX-MMD	1000BASE-SX	5.2(1)
	GLC-LH-SMD	1000BASE-LX	5.2(1)
	CWDM-SFP-xxxx[1]	1000BASE-CWDM	4.2(1)
	DWDM-SFP-xxxx[1]	1000BASE-DWDM	4.2(1)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
N7K-M148GS-11L	SFP-GE-S	1000BASE-SX	5.0(2a)
	GLC-SX-MM		5.0(2a)
	SFP-GE-L	1000BASE-LX	5.0(2a)
	GLC-LH-SM		5.0(2a)
	SFP-GE-Z	1000BASE-ZX	5.0(2a)
	GLC-ZX-SM		5.0(2a)
	GLC-EX-SMD	1000BASE-EX SFP	6.2(2)
	GLC-ZX-SMD	1000BASE-ZX SFP	6.2(2)
	GLC-T	1000BASE-T	5.0(2a)
	SFP-GE-T		5.0(2a)
	GLC-BX-D	1000BASE-BX10-D	5.2(1)
	GLC-BX-U	1000BASE-BX10-U	5.2(1)
	GLC-SX-MMD	1000BASE-SX	5.2(1)
	GLC-LH-SMD	1000BASE-LX	5.2(1)
	GLC-TE	1000BASE-T SFP	6.2(10)
	DWDM-SFP-xxxx[1]	1000BASE-DWDM	5.0(2a)
	CWDM-SFP-xxxx[1]	1000BASE-CWDM	5.0(2a)
N7K-M132XP-12	FET-10G	Cisco Fabric Extender Transceiver (FET)	5.1(1)
	SFP-10G-SR SFP-10G-SR-S	10GBASE-SR SFP+	4.2(6)
	SFP-10G-LR SFP-10G-LR-S	10GBASE-LR SFP+	4.0(3)
	SFP-10G-ER SFP-10G-ER-S	10GBASE-ER SFP+	4.0(1)
	SFP-H10GB-ACUxM[3]	SFP-H10GB-ACUxM Twinax Cable Active (7 m, 10 m)	5.1(2)
N7K-M132XP-12L	FET-10G	Cisco Fabric Extender Transceiver (FET)	5.1(1)
	SFP-10G-SR SFP-10G-SR-S	10GBASE-SR SFP+	5.1(1)
	SFP-10G-LR SFP-10G-LR-S	10GBASE-LR SFP+	5.1(1)
	SFP-10G-ER SFP-10G-ER-S	10GBASE-ER SFP+	5.1(1)
	SFP-10G-LRM	10GBASE-LRM SFP+	5.1(1)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	SFP-10G-ZR[3] SFP-10G-ZR-S	10GBASE-ZR SFP+	6.1(1)
	SFP-10G-AOCxM	10GBASE-AOC (Active Optical Cable) SFP+ Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(2)
	SFP-H10GB-ACUxM	SFP-H10GB-ACUxM Twinax Cable Active (7 m, 10 m)	5.1(1)
	SFP-H10GB-CUxM[3]	SFP-H10GB-CUxM Twinax Cable Passive (1 m, 3 m, 5 m)	5.1(2)
	SFP-H10GB-CUxM	SFP-H10GC-CUxM Twinax Cable Passive (1.5 m, 2 m, 2.5 m)	6.2(2)
	DWDM-SFP10G-xx.xx[1]	10GBASE-DWDM SFP+	6.1(1)
N7K-M224XP-25L	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.1(1)
	SFP-10G-SR SFP-10G-SR-S	10GBASE-SR SFP+	6.1(1)
	SFP-10G-LR SFP-10G-LR-S	10GBASE-LR SFP+	6.1(1)
	SFP-10G-ER SFP-10G-ER-S	10GBASE-ER SFP+	6.1(1)
	SFP-10G-ZR[3] SFP-10G-ZR-S	10GBASE-ZR SFP+	6.1(1)
	SFP-10G-LRM	10GBASE-LRM SFP+	6.1(1)
	SFP-10G-AOCxM	10GBASE-AOC (Active Optical Cable) SFP+ Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(2)
	SFP-H10GB-ACUxM	SFP-H10GB-ACUxM Twinax Cable Active (7 m, 10 m)	6.1(1)
	SFP-H10GB-CUxM[3]	SFP-H10GB-CUxM Twinax Cable Passive (1m, 3m, 5m)	6.1(1)
	SFP-H10GB-CUxM	SFP-H10GC-CUxM Twinax Cable Passive (1.5 m, 2 m, 2.5 m)	6.2(2)
	DWDM-SFP10G-xx.xx[1]	10GBASE-DWDM SFP+	6.1(1)
N7K-M206FQ-23L	QSFP-40G-SR-BD	Cisco 40G BiDi QSFP+	6.2(6)
	FET-40G	Cisco 40G Fabric Extender Transceiver (FET)	6.2(6)

Table 5 Transceivers Supported by Cisco NX-OS Software Releases (continued)

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	QSFP-40G-SR4 QSFP-40G-SR4-S	40GBASE-SR4 QSFP+	6.1(1)
	QSFP-40G-CSR4	40GBASE-CSR4 QSFP+	6.2(2)
	QSFP-40GE-LR4 QSFP-40G-LR4-S	40GBASE-LR4 QSFP+	6.1(4)
	QSFP-H40G-ACUxM	40GBASE-CR4 QSFP+ Direct Attach Copper Cable Active (7 m, 10 m)	6.2(2)
	QSFP-4X10G-ACxM	40GBASE-CR4 QSFP+ to 4x SFP+ Twinax Direct Attach Copper Breakout Cable Active (7 m, 10 m)	6.2(8)
	QSFP-H40G-AOCxM	40GBASE-AOC (Active Optical Cable) QSFP Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m, 15 m)	6.2(8)
	QSFP-4X10G-AOCxM	40GBASE-AOC (Active Optical Cable) QSFP to 4x10G SFP+ Breakout Cable (1 m, 2 m, 3 m, 5 m, 7 m, 10 m)	6.2(8)
	WSP-Q40GLR4L	40GBASE-LR4 lite (2km SMF) QSFP+	62(10)
	QSFP-40G-LR4	40GBASE-LR4 QSFP+ (Ethernet and OTU3 capable)	6.2(12)
	QSFP-40G-ER4	40GBASE-ER4 QSFP+ (40km)	6.2(12)
N7K-M202CF-22L	CFP-40G-SR4	40GBASE-SR4 CFP	6.1(2)
	CFP-40G-LR4	40GBASE-LR4 CFP	6.1(2)
	CFP-100G-SR10	100GBASE-SR10 CFP	6.1(3)
	CFP-100G-LR4	100GBASE-LR4 CFP	6.1(1)
	CFP-100G-ER4	100GBASE-ER4 CFP	6.2(10)

1. For a complete list of supported optical transceivers of this type, go to the Cisco Transceiver Module Compatibility Information page.
2. Multimode fiber supported on ports 41 to 48 only. Single-mode fiber support if applicable to all ports.
3. Minimum version supported is -02.
4. Requires CVR-X2-SFP10G, OneX Converter Module (X2 to SFP+ Adapter)

Upgrade/Downgrade Paths and Caveats

Note The XSD and POAP software files are exactly the same for Cisco NX-OS Release 6.2(6a) as for Release 6.2(6). Use those files.

This section includes information about upgrading or downgrading Cisco NX-OS software on Cisco Nexus 7000 Series devices. It includes the following sections:

- Supported Upgrade and Downgrade Paths, page 24
- ISSU Upgrade Steps, page 26
- Non-ISSU Upgrade Steps, page 26
- Upgrade or Downgrade Caveats, page 27

Supported Upgrade and Downgrade Paths

Note Before you upgrade or downgrade your Cisco NX-OS software, we recommend that you read the complete list of caveats in this section to understand how an upgrade or downgrade might affect your network, depending on the features that you have configured.

Do not change any configuration settings or network settings during a software upgrade. Any changes in the network settings might cause a disruptive upgrade.

See Table 6 for the in-service software upgrade (ISSU) path to Cisco NX-OS Release 6.2(x). Releases that are not listed for a particular release train do not support a direct ISSU or ISSD to the current release.

Note Only the ISSU combinations in the following table have been tested and are supported.

Table 6 Supported ISSU and ISSD Paths for the Cisco Nexus 7000 Series Chassis

Current Release	Release Train	Releases That Support ISSU to the Current Release	Releases That Support ISSD from the Current Release
NX-OS Release 6.2(12)	6.2	6.2(8a), 6.2(8b), 6.2(10)	No support
NX-OS Release 6.2(10)	6.2	6.2(8), 6.2(8a), 6.2(8b)	No support
NX-OS Release 6.2(8b)	6.2	6.2(8a)	No support
	6.1	6.1(5a)	No support
	5.2	5.2(9a)	No support
NX-OS Release 6.2(8a)	6.2	6.2(2), 6.2(2a), 6.2(6), 6.2(6a), 6.2(6b), 6.2(8)	6.2(2), 6.2(2a), 6.2(6), 6.2(6a), 6.2(6b), 6.2(8)
	6.1	6.1(3), 6.1(4), 6.1(4a), 6.1(5)	No support
	5.2	5.2(7), 5.2(9)	No support
NX-OS Release 6.2(8)	6.2	6.2(2), 6.2(2a), 6.2(6), 6.2(6a)	6.2(2), 6.2(2a), 6.2(6), 6.2(6a)
	6.1	6.1(3), 6.1(4), 6.1(4a), 6.1(5)	No support
	5.2	5.2(7), 5.2(9)	No support
NX-OS Release 6.2(6b)	6.2	6.2(2), 6.2(2a), 6.2(6), 6.2(6a)	6.2(2), 6.2(2a), 6.2(6), 6.2(6a)

Table 6 Supported ISSU and ISSD Paths for the Cisco Nexus 7000 Series Chassis

Current Release	Release Train	Releases That Support ISSU to the Current Release	Releases That Support ISSD from the Current Release
NX-OS Release 6.2(6a)	6.2	6.2(2), 6.2(2a), 6.2(6)	6.2(2), 6.2(2a), 6.2(6)
	6.1	6.1(3), 6.1(4), 6.1(4a)	No support
	6.0	6.0(4)	No support
	5.2	5.2(7), 5.2(9)	No support
NX-OS Release 6.2(6)	6.2	6.2(2), 6.2(2a)	6.2(2), 6.2(2a)
	6.1	6.1(3), 6.1(4), 6.1(4a)	No support
	6.0	6.0(4)	No support
	5.2	5.2(7), 5.2(9)	No support
NX-OS Release 6.2(2a)	6.2	6.2(2)	6.2(2)
	6.1	6.1(2), 6.1(3), 6.1(4), 6.1(4a)	No support
	6.0	6.0(4)	No support
	5.2	5.2(4), 5.2(5), 5.2(7), 5.2(9)	No support
NX-OS Release 6.2(2)	6.1	6.1(2), 6.1(3), 6.1(4), 6.1(4a)	No support
	6.0	6.0(4)	No support
	5.2	5.2(4), 5.2(7), 5.2(9)	No support

See Table 7 for the in-service software upgrade (ISSU) path to Cisco NX-OS Release 6.2(x) for the Cisco Nexus 7700 Series chassis. Releases that are not listed for a particular release train do not support a direct ISSU or ISSD to the current release.

Table 7 Supported ISSU and ISSD Paths for the Cisco Nexus 7700 Series Chassis

Current Release	Release Train	Releases That Support ISSU to the Current Release	Releases That Support ISSD from the Current Release
NX-OS Release 6.2(12)	6.2	6.2(8a), 6.2(8b), 6.2(10)	No support
NX-OS Release 6.2(10)	6.2	6.2(8), 6.2(8a), 6.2(8b)	No support
NX-OS Release 6.2(8b)	6.2	6.2(8a)	No support
NX-OS Release 6.2(8a)	6.2	6.2(6), 6.2(6a), 6.2(6b), 6.2(8)	6.2(6), 6.2(6a), 6.2(6b), 6.2(8)
NX-OS Release 6.2(8)	6.2	6.2(6), 6.2(6a)	6.2(6), 6.2(6a)
NX-OS Release 6.2(6b)	6.2	6.2(6a)	6.2(6a)
NX-OS Release 6.2(6a)	6.2	6.2(6)	6.2(6)
NX-OS Release 6.2(6)	6.2	No support	No support

Unless otherwise noted, all releases within the same release train are ISSU and ISSD compatible to releases within the same train.

If you are running a Cisco NX-OS release earlier than Release 5.2, you can perform an ISSU in multiple steps. Table 8 lists the supported multistep ISSU paths.

Table 8 Multistep ISSU Paths

Starting Release	Intermediate Release	Destination Release
6.1(1)	6.1(4)	6.2(8a)
5.2(3)	5.2(9)	6.2(8a)
6.1(1)	6.1(4)	6.2(8a)
5.2(3)	5.2(9)	6.2(8a)
6.1(1)	6.1(4)	6.2(6b)
5.2(3a)	5.2(9)	6.2(6b)
6.1(1)	6.1(4)	6.2(2a)
5.2(3a)	5.2(9)	6.2(2a)
4.2(6), 5.0(3), 5.1(3), or 5.1(5)	5.2(7)	6.2(6b)
4.2(6), 5.0(3), 5.1(3), or 5.1(5)	5.2(7)	6.2(2a)

ISSU Upgrade Steps

To perform an ISSU upgrade to Release 6.2(x) from one of the ISSU supported releases listed in Table 6, follow these steps:

1. Enter the **show running-config aclmgr inactive-if-config** command for all VDCs.
2. Enter the **clear inactive-config acl** command for all VDCs.
3. If the configuration has any **mac packet-classify** configurations on any interfaces, remove all of the configurations by entering the **no mac packet-classify** command.
4. If there is an VACL configuration and the ISSU to Release 6.2(2a) is from a release earlier than Release 6.1(3), perform the ISSU in two steps:
   a. Upgrade to Release 6.1(3) or Release 6.1(4).
   b. Upgrade from Release 6.1(3) or Release 6.1(4) to Release 6.2(2a).
5. Start the ISSU procedure.

Non-ISSU Upgrade Steps

To perform a non-ISSU upgrade to Release 6.2(x) from any release earlier than Release 6.2(x), follow these steps:

1. Change the boot variable.
2. Enter the **copy running-config startup-config vdc-all** command.
3. Enter the **reload** command to reload the switch.

Note Allow time after the reload for the configuration to be applied.

To perform a non-ISSU upgrade from any PRIOR supported Cisco NX-OS release to Release 6.2(2) or 6.2(2a), follow these steps:

1. Enter the **show running-config aclmgr inactive-if-config** command for all VDCs.
2. Enter the **clear inactive-config acl** command for all VDCs.
3. Change the boot variable to boot Release 6.2(2) or 6.2(2a).
4. Enter the **copy running-config startup-config vdc-all** command.
5. Enter the **reload** command to reload the switch.

Note Allow time after the reload for the configuration to be applied.

6. Once all line cards come up, enter the **show running-config aclmgr** command in all VDCs.
7. Enter the **clear inactive-config acl** command for all VDCs.
8. Enter the **copy boot flash:/vdc_x/aclmgr-inactive-config.cfg running-config** command for all VDCs.

For complete instructions on upgrading your software, see the *Cisco Nexus 7000 Series NX-OS Upgrade Downgrade Guide.*

Non-ISSU upgrades are also referred to as a cold boot.

Upgrade or Downgrade Caveats

A software upgrade or downgrade can be impacted by the following features or hardware:

- The online diagnostics tests PortLoopbackTest, SnakeLoopbackTest, RewriteEngineText, and BootupPortLoopbackTest are not supported on the N77-F348XP-23 module.
- Feature Support

  Any features introduced in a release must be disabled before downgrading to a release that does not support those features.
- Unsupported Modules

  When manually downgrading from Cisco NX-OS Release 6.2(2) to an earlier release, first power down all modules that are unsupported in the downgrade image. Then, purge the configuration of the unsupported modules using the **purge module** *module_number* **running-config** command.
- The Intelligent Traffic Director (ITD) Configuration is Lost After an Upgrade or Downgrade

  The ITD configuration is not saved during an in-service software upgrade (ISSU) from Release 6.2(8) or Release 6.2(8a) to Release 6.2(10) or an in-service software downgrade (ISSD) to Release 6.2(8a) or Release 6.2(8). Before performing an ISSU or ISSD, you must remove the ITD configuration by using the **no feature itd** command. After the upgrade or downgrade, you must manually reapply the configuration.
- OTV AED Displays a Version Mismatch After an Upgrade

  If your OTV setup includes dual Authoritative Edge Devices (AEDs), the output of the **show otv** command displays a Version Mismatch after you perform an upgrade to from Release 6.1.X to Release 6.2.x because one AED is upgraded to 6.2.x and second AED is still running Release 6.1.x or an earlier release. This behavior is expected. For example:

  ```
  otv# show otv
  ```

```
OTV Overlay Information
Site Identifier 0000.0000.0751

Overlay interface Overlay1

VPN name            : Overlay1
VPN state           : UP
Extended vlans      : 1405-1406 2405-2406 3401-3406 3419-3422 3427-3428
(Total:16)
Control group       : 239.194.0.1
Data group range(s) : 232.19.58.0/24
 Broadcast group     : 239.194.0.1
Join interface(s)   : Eth7/9 (10.13.244.146)
 Site vlan           : 2905 (up)
 AED-Capable         : No (Version Mismatch)
Capability          : Multicast-Reachable
```

- The Auto-Recovery Disable and Auto-Recovery Timeout Configurations Might not be Persistent Across an Upgrade, Downgrade, or Reload.

  Before Cisco NX-OS 6.2(2), the auto-recovery default was disabled. In Cisco NX-OS Release 6.2(2) and later releases, auto recovery is enabled by default. Because the default behavior of auto recovery is changed, the following caveats apply:

  - After you perform an ISSU to Release 6.2(2) or a later release, auto recovery is enabled by default after the upgrade. If you want to disable auto recovery in Release 6.2(2) or a later release, you must use the **no auto-recovery** command to explicitly disable auto recovery.
  - After you perform a downgrade from Release 6.2(6) or 6.2(4a) to Release 6.2(2), auto recovery is enabled by default after the downgrade regardless of the auto-recovery configuration before the downgrade. If you want to disable auto recovery in Release 6.2(2) or a later release, you must use the **no auto-recovery** command to explicitly disable auto recovery.
  - After you reload a device that is running Release 6.2(2), 6.2(4a) or 6.2(6), auto recovery is enabled by default after the reload. If you want to disable auto recovery in Release 6.2(2) or a later release, you must use the **no auto-recovery** command to explicitly disable auto recovery after the reload.
  - If you changed the auto-recovery reload-delay value from the default value in Release 6.2(2), 6.2(4a), or 6.2(6), the value reverts to the default (240sec) after you perform a downgrade to any of the three specified releases and after you reload a device that is running any of the three specified releases. If you want a value other than the default for the auto-recovery reload-delay, you must reconfigure the auto-recovery reload-delay value after the downgrade or the reload.

- Auto Recovery Does not Restore the vPC.

  After a disruptive upgrade or sequential reloads, the Auto Recovery feature does not restore the vPC when the keep-alive and peer-link are down. By design, a disruptive upgrade causes the start-up configuration to run through a different process when it is installed. This issue is seen on sequential reloads until the binary configuration is rebuilt.

  To ensure that auto recovery is successful and the vPC is restored after a disruptive upgrade, use the **save running-config startup-config vdc-all** command after the disruptive upgrade and then reload the device.

- Supervisor CoS2q and DSCP2Q Mappings are Restored to Default Mappings

  After you perform an ISSU to Cisco NX-OS Release 6.2(6a) or a later release, one of the following warning messages is displayed:

  - `Supervisor cos2q, dscp2q mappings are restored to default mappings. Please apply cos2q, dscp2q mappings again. For dscp2q mappings, turn off and on dscp global knob.`

  - `Supervisor cos2q mappings are restored to default mappings. Please apply cos2q mappings again.`
  - `Supervisor dscp2q mappings are restored to default mappings. Please turn off and on dscp global knob.`

  If you are using the 8e4q4q template in Cisco NX-OS Release 6.2(6) or an earlier release and you perform an ISSU to Release 6.2(6a) or a later release, the Class of Service (Cos2q) and differentiated services code point (DSCP2q) maps are restored to the defaults in the supervisor engine. However, because the ISSU is non-disruptive, the mappings are not programmed in the hardware. The next time you save a new configuration, the default mappings are pushed to the module, which is disruptive.

  To avoid this issue, you must explicitly push the COS2q and DSCP2q mappings to the module after the ISSU.

  If you are not using the 8e4q4q template and you perform an ISSU from Cisco NX-OS Release 6.2(6) or an earlier release to Release 6.2(6a) or a later release, you can ignore the warning message; no action is required.

- VDC Ports Can Become Unallocated After a Downgrade

  If your Cisco Nexus 7000 Series device that is running Release 6.2(x) has **f2e** in its configuration as a part of the **limit-resource module-type** command, some interfaces might become unallocated after a downgrade from Release 6.2(x) to Release 6.1(x) or an earlier release.

  Because there is no **f2e** keyword for the **limit-resource module-type f2 f2e** command in releases before Release 6.2(x), the command fails and the VDC is converted to a limit-resource module-type m1 f1 m1xl m2xl VDC. This issue can occur even if you do not have any F2e Series modules in the chassis.

  To avoid this issue, do one of the following procedures:

Note Before any upgrade, save the startup configuration to bootflash.

  - If you have a saved backup of the switch configuration from the release to which you want to downgrade, do the following:

    1. Before downgrading to Release 6.1(x) or an earlier release, change the boot variable to 6.1(x).
    2. Enter the **copy running-config startup-config** command to save the boot variable to 6.1(x).
    3. Enter the **write erase** command on the switch.
    4. Reload the switch with the downgrade image.
    5. After all I/O modules and supervisor modules come up, enter the **copy** *saved-config* **running-config** command. For example, if your saved configuration is in the file bootflash:saved_config, enter the **copy bootflash:saved_config running-config** command.

  - If you do not have a backup switch configuration from the release to which you want to downgrade, follow these steps to change instances of **f2e** to **f2** in the **limit-resource module-type** command of the Release 6.2(x) configuration, before reloading the switch:

    1. Copy your release 6.2(x) configuration to a file on a remote server by entering the **copy running-config scp://** command.
    2. In the copied file, find all instances of **f2e** in the **limit-resource module-type f2 f2e** command and replace it with **f2**. Remove any redundant instances of **f2**. For example, if you replace **f2e** in the **limit-resource module-type f2 f2e** command, make sure that you change it to **limit-resource module-type f2.** Save the file.
    3. Enter the **copy running-config startup-config** command to save the boot variable to 6.2(x).

4. Enter the **write erase** command on the switch.
5. Reload the switch with the downgrade image.
6. After all I/O modules and supervisor modules come up, copy the edited file to your running configuration by entering the **copy** *saved-config* **running-config** command.

- F2e Interfaces Might Become Unallocated After an Upgrade

  All F2e interfaces might become unallocated immediately after an upgrade from Cisco NX-OS Release 6.1(x) to Release 6.2(x). Generally, this situation should not occur if you follow recommended upgrade procedures. However, if you find that all F2e interfaces are unallocated after the upgrade, before doing any further configuration, complete the following steps to fix the problem:

  1. Retrieve an ASCII version of the complete switch configuration and edit it as follows:
     - Replace every instance of **f2** with **f2 f2e** in any **limit-resource module-type** command and **any system module-type** command.
  2. When all the modules are online and all the VDCs are created and in active status, in the context of the Default VDC or the Admin VDC, enter the **copy modified-ascii-config running-config echo** command to apply the modified ASCII configuration file.
  3. After verifying the sanity of the current configuration in the context of the Default VDC or the Admin VDC, enter the **copy running-config startup-config vdc-all** command to save the configuration.

- Storage VDC

  Starting with Cisco NX-OS Release 6.2(x), shared interfaces must come from the module types that the storage VDC supports, and having F1 and F2 interfaces shared in the same VDC is not supported. In Release 6.2(x), the **vdc** command will fail if you attempt to share F2 interfaces with a VDC that supports only F1 (or vice versa).

  In Cisco NX-OS Release 6.1(x), you are allowed to share F2 interfaces with a storage VDC that supports only F1, and you can share F1 interfaces with a storage VDC that supports only F2. But having F1 and F2 shared interfaces in the same storage VDC is somewhat problematic and can potentially cause issues.

  If you have a Cisco NX-OS Release 6.1(x) configuration that has shared F1 and F2 interfaces in the same storage VDC, this configuration will be allowed temporarily in Release 6.2(x) if you are performing an upgrade through a binary configuration, which is the typical upgrade path, rather than through an ASCII configuration. This temporary situation provides a transition period to update the configuration; however, the functionality of the interfaces is not guaranteed during this time. The configuration might be lost after a switch reboot. In addition, the configuration will fail in Release 6.2(x), if you remove the interface and attempt to add it back to a shared storage VDC. Because of these constraints, we recommend that you remove the combination of shared interfaces before the upgrade to Release 6.2(x), or after it.

  Before or after an upgrade from Cisco NX-OS Release 6.1(x) to Release 6.2(x), apply either the **limit-resource module-type f1** command or the **limit-resource module-type f2** command to the storage VDC, and check that the following storage VDC configurations are removed:

  - Shared F2 interfaces with a storage VDC that support only F1, or shared F1 interfaces with a storage VDC that supports only F2
  - F1 and F2 interfaces in the same storage VDC

- Upgrade With an M2 Series Module Installed

  In rare instances, if you perform an ISSU from Cisco NX-OS Release 6.1(2) or an earlier release to Release 6.2(x) and you have an M2 or F2 Series module installed in your Cisco Nexus 7000 Series system, the upgrade might fail with the following error:

```
Return code 0x40710027 (BIOS flash-type verify failed)
```

To work around this issue, ISSU to Release 6.1(3) before you upgrade to Release 6.2(x) or a later release, or upgrade via a traditional reload. For additional details, see CSCud63092.

- FEX Host Interface

  When you upgrade Cisco NX-OS software by changing boot variables and reloading the device, make sure to save the FEX HIF configuration to the startup configuration, as well as another location (such as bootflash or an external server). Once the upgrade to a new release is complete, and the FEX is fully online and associated, reapply the FEX HIF configuration.

Note During the process of Cisco Fabric Extender (FEX) modules getting connected to a Cisco Nexus 7000 Series switch, if the switch is manually upgraded or downgraded, FEX host interfaces (HIFs) lose the configuration. To avoid it, if you are manually upgrading the vPC system, you must save the FEX HIF (FEX host interfaces connected to hosts) configurations to both the startup configuration file and to an external device before starting the reload, and reapply the configuration once the FEX module is fully online.

- FEX Queuing

  The FEX queuing feature is enabled by default if you perform an ISSU to a release that supports this feature from an earlier release that supports this feature. The feature is not enabled if you perform an ISSU to a release that supports this feature from an earlier release that does not support this feature.

  For an ISSU to Cisco NX-OS Release 6.2(2a) from Release 6.1(x), you must reload the FEX to enable FEX queuing after the ISSU. Enter the **show queuing interface** *fex-hif-port* command to verify if FEX queuing is enabled on a given FEX.

- VACL Configuration Should Be Removed Before ISSU

  If an active or inactive VACL configuration is present in the running configuration, an ISSU from any release earlier than Cisco NX-OS Release 6.1(3) to Release 6.2(2a) will not succeed.

  Enter the **show running-config aclmgr inactive-if-config** command to check for inactive policies. You can remove these policies by entering the **clear inactive config acl/qos** command.

  To work around this issue, remove all VACLs from the system before the ISSU, or perform a two-step upgrade to Release 6.1(3) and then to Release 6.2(2a).

- The MAC Packet Classify Configuration Should Be Removed Before ISSU

  If the **mac packet-classify** command is configured for any interface, the ACLMGR process might fail during an ISSU to Cisco NX-OS Release 6.2(2a).

  To work around this issue, remove all **mac packet-classify** commands from the configuration for all interfaces before the ISSU.

- OTV

  Any upgrade from an image that is earlier than Cisco NX-OS Release 6.2(2) to an image that is Cisco NX-OS Release 6.2(2) or later in an OTV network is disruptive. When you upgrade from any previous release, the OTV overlay needs to be shut down for ISSU to operate.

  For more details, see the "Preparing OTV for ISSU to Cisco NX-OS 5.2(1) or Later Releases in a Dual-Homed Site" section in the *Cisco Nexus 7000 Series NX-OS OTV Configuration Guide*.

- LISP

  If you have LISP configured on a Cisco Nexus 7000 Series device, you must remove the configuration before an ISSU.

- OSPF

  Cisco NX-OS Release 6.1 supports an increased number of Open Shortest Path First (OSPF) process instances per VDC. See the *Cisco Nexus 7000 Series NX-OS Verified Scalability Guide* for the latest verified number.

  If you have more than four OSPF v2 or more than four OSPF v3 process instances configured and you manually downgrade to an earlier release, you must remove instances 5 and higher. Use the following command to match an OSPF v2 process tag with an OSPF process instance:

  ```
  switch# show system internal sysmgr service name ospf
  Service "__inst_005__ospf" ("ospf", 13):  <= OSPF process instance
          UUID = 0x41000119, PID = 3402, SAP = 320
          State: SRV_STATE_HANDSHAKED (entered at time Mon Jul 23 05:11:33 2012).
          Restart count: 1
          Time of last restart: Mon Jul 23 05:11:33 2012.
          The service never crashed since the last reboot.
          Tag = 1                <= configured process tag
          Plugin ID: 1
  ```

  Use the **show system internal sysmgr service name ospfv3** command to match an OSPF v3 process tag with an OSPF v3 process instance.

  This enhancement was added for EIGRP and ISIS with Cisco NX-OS Release 6.2.

- ACL

  During an ISSU from Cisco NX-OS Release 5.2(x) or Release 6.1(x) to Release 6.2(2a), the system prompts you to clear inactive access control list (ACL) configurations. Enter the **clear inactive-config acl** command to clear any inactive ACL configurations.

- CoPP

  The default Control Plane Policing (CoPP) policy does not change when you upgrade the Cisco NX-OS software.

  If you manually downgrade without using ISSD to a release earlier than NX-OS Release 5.2(1), the CoPP configuration is lost, and a CoPP policy is no longer attached to the control plane.

- Aggressive Failure Detection Timers

  ISSU, stateful switchover (SSO), and graceful restart are not supported when aggressive failure detection timers are used for any Layer 3 protocols. Starting in Cisco NX-OS Release 5.2(3a), the First Hop Redundancy Protocol (FHRP) with aggressive timers has been validated for SSO or ISSU using the extended hold timer feature. Other protocols such as OSPF have been validated with aggressive timers without SSO or ISSU support. For additional information on aggressive timer support and extended hold timers for FHRP, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide*.

- IPFIB Errors

  During an upgrade to Cisco NX-OS Release 5.2(7) or a later release, the following error messages might appear:

  ```
  %IPFIB-SLOT2-2-FIB_TCAM_HA_ERROR: FIB recovery errors, please capture 'show
  tech forwarding l3 unicast' and 'show tech forwarding l3 multicast'
  ```

  In addition, the ipfib process might fail.

  This issue can be triggered when the following sequence of events occur:

  - You perform an ISSU to Cisco NX-OS Release 5.2(1), Release 5.2(3a), Release 5.2(4), or Release 5.2(5) release from an earlier 5.0(x) or 5.1(x) release and you have not reloaded the switch.

- You make configuration changes in the 5.2(x) release running on the Cisco Nexus 7000 Series system.
- You perform an ISSU to NX-OS Release 5.2(7) or a later release.

To work around this issue, follow these steps:

1. Prior to the upgrade, enter the following commands to avoid the issue:

a. Enter the **feature lisp** command.

b. Enter the **ip lisp etr** command for all virtual routing and forwarding (VRF) instances, followed by the **no ip lisp etr** command.

c. Enter the **no feature lisp** command.

2. If you experience this issue, reload the affected modules on your Cisco Nexus 7000 Series system.

Note The Transport Services Package license is required to enable the Locator/ID Separation Protocol (LISP). If you do not have this license, you can enable the grace period for it. If you cannot enable the grace period, perform an ISSU and reload the affected modules.

Perform these steps even if you are not using LISP because the issue can occur even if LISP is not running.

- BGP

  If both **send-community** and **send-community extended** are in the configuration for Cisco NX-OS 6.1 or an earlier release and an ISSU is performed, only **send-community extended** will be present in the configuration for a Cisco NX-OS 6.2 or later release after the ISSU. You will have to manually reconfigure **send-community**. The running configuration will show **send-community both** instead of both commands.

CMP Images

Cisco NX-OS Release 6.2(2a) does not include a new connectivity management processor (CMP) image.

Cisco NX-OS Release 6.2(2) includes a new CMP image for the Cisco Nexus 7000 Supervisor 1 module. The CMP is upgraded to Release 6.2(2) on a successful ISSU to Cisco NX-OS to Release 6.2(2). When the ISSU completes, reload the CMP image on the active and standby Supervisor 1 modules.

Cisco NX-OS Release 6.2(2) does not include a CMP image for the Cisco Nexus 7000 Supervisor 2 or Supervisor 2 Enhanced module because neither module has a CMP.

Cisco Nexus 7700 Series switches do not have a CMP.

For additional information about the CMP, see the *Cisco Nexus 7000 Series Connectivity Management Processor Configuration Guide.*

EPLD Images

Cisco NX-OS Release 6.2(8) includes new EPLD images for the M1 Series I/O modules with XL and M2 forwarding engines, the Network Analysis Module (NAM) Azuma FPGA, and the N77-F248XP-23E I/O module power manager and IO devices. For more information about upgrading to a new EPLD image, see the *Cisco Nexus 7000 Series FPGA/EPLD Upgrade Release Notes, Release 6.2.*

Cisco NX-OS Release 6.2(2a) does not include new EPLD images.

Cisco NX-OS Release 6.2(2) includes a new EPLD image for the Cisco Nexus 7000 Series Supervisor 1 and the Supervisor 2 module.

Cisco Nexus 7700 Series switches have an EPLD image that is programmed on the switches. This EPLD image is different than the EPLD image for the Cisco Nexus 7000 switches.

The Cisco Nexus 7000 Series Network Analysis Module (Cisco NAM-NX1) also includes an EPLD image that is programmed on the device.

New Hardware

This section briefly describes the new hardware introduced in Cisco NX-OS Release 6.2. For detailed information about the new hardware, see the *Cisco Nexus 7000 Series Hardware Installation and Reference Guide*.

This section includes the following topics:

New Hardware in Cisco NX-OS Release 6.2(12)

Cisco NX-OS Release 6.2(12) introduces new hardware:

- N7K-F348XP-25

The Cisco Nexus 7000 F3-Series module is a low-latency, high-performance, high-density 1/10 Gigabit Ethernet module. It is operationally consistent with the Cisco Nexus 7700 F3-Series modules and shares a common system architecture and the same application-specific integrated circuit (ASIC) technology.

New Hardware in Cisco NX-OS Release 6.2(10)

Cisco NX-OS Release 6.2(10) introduces new hardware:

- N7K-F306CK-25

Cisco NX-OS 6.2(10) supports the Cisco Nexus 7000 F3-Series 6-Port 100 Gigabit Ethernet Module. The Cisco Nexus 7000 F3-Series module is a low-latency, high-performance, high-density 100 Gigabit Ethernet module. It is operationally consistent with the Cisco Nexus 7700 F3-Series modules and shares a common system architecture and the same application-specific integrated circuit (ASIC) technology.

New Hardware in Cisco NX-OS Release 6.2(8)

Diagnostic Optical Monitoring (DOM) is supported on the F3 Series modules in Cisco NX-OS Release 6.2(8).

Breakout cables and 40-Gigabit copper cables are available for the following modules:

- N77-F324FQ-25
- N7K-F312FQ-25
- N7K-M206FQ-23L

The Cisco Nexus 2248TP-E FEX supports 10 Mbps. transmission speeds.

New Hardware in Cisco NX-OS Release 6.2(6)

Cisco NX-OS Release 6.2(6) introduces new hardware that is described in the following sections:

- Cisco Nexus 7706 Switch, page 35
- Cisco Nexus 7000 Series I/O Module, page 35
- Cisco Nexus 7700 Series I/O Modules, page 35

See Table 2 for the PIDs associated with the hardware components described in this section.

In addition, Release 6.2(6) introduces new cables that allow you to expand the number of ports on selected modules. These cables allow you to connect to one of the 40-Gigabit Ethernet ports on the following modules and connect the other end to up to four 10-Gigabit Ethernet ports. Two cables are available and offer the following:

- Support for the Cisco Nexus 7000 F3 Series 12-port 40-Gigabit Ethernet (QSFP+) Module to support up to 48 10-Gigabit Ethernet ports per slot for 18-slot, 10-slot, 9-slot, and 4-slot chassis.
- Support for the Cisco Nexus 7000 M2-Series 6-port 40-Gigabit Ethernet I/O Module XL to support up to 24 10-Gigabit Ethernet ports per slot for 18-slot, 10-slot, 9-slot, and 4-slot chassis.

Cisco Nexus 7706 Switch

The Cisco Nexus 7706 switch is a 6-slot chassis that delivers 1.32 Gbps per slot. It has two half-width slots for Supervisor 2E modules and four full-width slots for I/O modules. With the F3 Series modules, the switch provides 40-Gigabit Ethernet and 100-Gigabit Ethernet ports. The supervisor modules and I/O modules are interchangeable between the Cisco Nexus 7710 switch and the Cisco Nexus 7706 switch.The switch supports up to six fabric modules and uses three fan tray that provide front-to-back airflow cooling. To power the switch, you can use one or two AC or DC 3-kW power supply units, and you can use up to two more power supplies to provide power-supply or grid power redundancy.

Cisco Nexus 7000 Series I/O Module

Cisco NX-OS Release 6.2(6) supports the Cisco Nexus 7000 F3 Series 12-port 40-Gigabit Ethernet (QSFP+) Module (F3 Series).

This multiprotocol module supports Ethernet, FabricPath, FCoE, OTV, LISP, and VXLAN, MPLS, VPLS, DFA. This module supports Cisco Nexus 2000 Series Fabric Extenders.

Cisco NX-OS Release 6.2(8) supports F3 line cards as first hop in a LISP multi-hop setup.

Cisco Nexus 7700 Series I/O Modules

Cisco NX-OS Release 6.2(6) supports the following modules for the Cisco Nexus 7700 Series switch:

- Cisco Nexus 7700 F3 Series 48-port 1/10-Gigabit Ethernet (SFP+) Module (F3 Series)
- Cisco Nexus 7700 F3 Series 24-port 40-Gigabit Ethernet (SFP+) Module (F3 Series)

- Cisco Nexus 7700 F3 Series 12-port 100-Gigabit Ethernet (SFP+) Module (F3 Series)

These multiprotocol modules support Ethernet, FabricPath, FCoE, OTV, LISP, VXLAN, MPLS, VPLS, and DFA. These modules, except the 12-port 100-Gigabit Ethernet module, support Cisco Nexus 2000 Series Fabric Extenders.

Cisco NX-OS Release 6.2(8) supports F3 line cards as first hop in a LISP multi-hop setup.

New Hardware in Cisco NX-OS Release 6.2(2)

Cisco NX-OS Release 6.2 introduces new hardware that is described in the following sections:

- Cisco Nexus 7710 Switch, page 36
- Cisco Nexus 7718 Switch, page 36
- Cisco Nexus 7700 Series I/O Module, page 36
- Cisco Nexus Fabric Extenders, page 36
- Cisco Nexus 7000 Series Network Analysis Module, page 37

See Table 2 for the PIDs associated with the hardware components described in this section.

Cisco Nexus 7710 Switch

The Cisco Nexus 7710 switch is a 10-slot chassis that delivers 1.32 Gbps per slot. It has two half-width slots for two Supervisor 2E modules and eight full-width slots for I/O modules. The switch supports up to six fabric modules and up to eight 3000 W power supply units. There are three fan trays. Airflow is front to back in the Cisco Nexus 7710 switch.

Cisco Nexus 7718 Switch

The Cisco Nexus 7718 switch is a 18-slot chassis that delivers 1.32 Gbps per slot. It has two half-width slots for two Supervisor 2E modules, and 16 full-width slots for I/O modules. The switch supports up to 6 fabric modules and up to 16 3000 W power supply units. There are three fan trays. Airflow is front to back in the Cisco Nexus 7718 switch.

Cisco Nexus 7700 Series I/O Module

Cisco NX-OS Release 6.2(2) supports the Cisco Nexus 7700 Enhanced F2-Series 48-port 1/10-Gigabit Ethernet (SFP+) Module (F2e Series).

This multiprotocol module supports Ethernet, FabricPath, and FCoE, and it supports Cisco Nexus 2000 Series Fabric Extenders.

Cisco Nexus Fabric Extenders

Cisco NX-OS Release 6.2(2) supports the following Cisco Nexus Fabric Extenders:

- Cisco Nexus 2232TM-E 10GE Fabric Extender
- Cisco Nexus 2248PQ-10GE Fabric Extender

For additional information, see the *Cisco Nexus 2000 Series Hardware Installation Guide.*

- Cisco Nexus B22HP Fabric Extender (blade fabric extender for HP)

For additional information, see the *Cisco Nexus B22 Fabric Extender for HP Getting Started Guide*.

Cisco Nexus 7000 Series Network Analysis Module

The Cisco Nexus 7000 Series Network Analysis Module (Cisco NAM-NX1) is the first service module for the Cisco Nexus 7000 Series platform. With the Cisco NAM-NX1, you can implement network analysis and monitoring in the data center. Cisco NAM-NX1 can be installed into any one of the network module slots on a Cisco Nexus 7000 Series switch.

For additional information, see the *Cisco Prime Network Analysis Module User Guide 6.0*.

New Hardware in Cisco NX-OS Release 6.2(2a)

Cisco NX-OS Release 6.2(2a) does not include new hardware.

Changed Software Features

This section describes software features that are changed in Cisco NX-OS Release 6.2.

This section includes the following topics:

- VDC Changes in Cisco NX-OS Release 6.2, page 37
- Features Available on F2, F2e, and F3 Series Modules, page 38

VDC Changes in Cisco NX-OS Release 6.2

Cisco NX-OS Release 6.2(2) provides several changes to VDCs that apply to Cisco Nexus 7000 Series switches and Cisco Nexus 7700 Series switches.

- For a chassis with only F2e Series modules, the default VDC will be created using an F2e Series module as a supported module, unless you apply your own configuration.
- In Release 6.2(2), F2 Series modules can only be a part of an F2 VDC or an F2-F2E VDC.
- The F2e and F2 Series modules cannot exist with the F1 Series module in a VDC.
- A new VDC type, F2E, supports only F2e Series modules, but can be added to other VDC types to allow F2E to be part of the same VDC with F2 modules, or M1 or M2 Series modules. An F2e Series module and any M Series module can be configured in the same VDC.
- During an upgrade to Cisco NX-OS Release 6.2(2), if you have F2 and F2e Series modules, the VDC type automatically changes to the F2-F2E VDC.
- In an F2-F2E VDC, only the features supported on the F2 Series module are supported.

Review the "Upgrade or Downgrade Caveats" section on page 27 for information related to VDCs that you should be aware of before upgrading to Cisco NX-OS Release 6.2(2).

For additional information, see the *Cisco Nexus 7000 Series NX-OS VDC Configuration Guide*.

Features Available on F2, F2e, and F3 Series Modules

Some software features are not available on the F2 or F3 Series modules in Cisco NX-OS Release 6.x. See Table 9 for a list of features that have hardware and software support on the F2, F2e, and F3 Series modules.

Table 9 Feature Support on F2 Series, F2e Series, and F3 Series Modules

	Hardware Support			Software Support		
Feature	**F2 Series Module**	**F2e Series Module**	**F3 Series Module**	**F2 Series Module**	**F2e Series Module**	**F3 Series Module**
ACL Capture	Yes	Yes	Yes	No	No	6.2(6)

ERSPAN source
ERSPAN destination
	Yes
	Yes
		Yes
		Yes
			Yes
			Yes
				6.1(1)
				6.2(2)
					6.1(2)
					6.2(2)
						6.2(6)
						No

FCoE	Yes	Yes	Yes	6.1(1)	6.1(2) SFP+ only	No
GRE tunnels	No	No	Yes	N/A	N/A	No
LISP	No	No	Yes	N/A	N/A	No
MACSec	No	Yes	Yes	N/A	Yes	No
MPLS	No	No	Yes	N/A	N/A	No
NetFlow	Yes	Yes	Yes	6.1(2)	6.1(2)	6.2(6)
OTV	No	No	Yes	N/A	Internal interface, 6.2(2)	6.2(6)
PIM-Bidir	No	Yes	Yes	N/A	No	No
VLAN counters	No	Yes	Yes	N/A	6.2(2)	6.2(6)
Interoperability with M Series modules	No	Yes	Yes (M2)	N/A	6.2(2)	6.2(6)

Note The same features that are supported on an F2 Series module on a Cisco Nexus 7000 Series switch are also supported on a Cisco Nexus 7700 Series switch.

Note On the F3 Series, only the 10-Gigabit I/O module offers MACSec capabilities. The 40-Gigabit and 100-Gigabit F3 Series modules do not support MACSec.

New Software Features

This section briefly describes the new features introduced in Cisco NX-OS Release 6.2 software. For detailed information about the features listed, see the documents listed in the “Related Documentation” section. The “New and Changed Information” section in each of these books provides a detailed list of all new features and includes links to the feature description or new command.

This section includes the following topics:

For additional information about these features, see the *Cisco Nexus 7000 Series Switches Configuration Guides.*

Cisco NX-OS Release 6.2(2) Software Features

Cisco NX-OS Release 6.2(2) includes the new features described in the following sections:

MPLS

New MPLS features in Release 6.2(2) include the following:

- Any Transport over MPLS (AToM), which accommodates different types of Layer 2 packets, including Ethernet and VLAN, to enable the service provider to transport different types of traffic over the backbone.
- Pseudowire provisioning for AToM, which enables you to configure static pseudowires in cases where you cannot use directed control protocols, such as the Label Distribution Protocol or Resource Reservation Protocol over traffic-engineered tunnels (RSVP-TE).
- Ethernet over Multiprotocol Label Switching (EoMPLS), which is a Virtual Private Wire Service (VPWS) that is used to carry Layer 2 Ethernet frames over an MPLS network. EoMPLS enables service providers to offer emulated Ethernet services over existing MPLS networks.

- EoMPLS Graceful Restart, which adds support for a switch that is configured with the Label Distribution Protocol (LDP) Graceful Restart (GR) to assist its neighboring switches to recover gracefully from an interruption in service.
- Layer 2 and Layer 3 load balancing coexistence, which supports Layer 3 VPN and Layer 2 VPN forwarding that is performed independently on the switch using two different types of adjacencies. The forwarding is not impacted by having a different method of load balancing for the Layer 2 VPN.
- Virtual Private LAN Service, which supports a point-to-multipoint service between multiple customer sites using a mesh of point-to-point pseudowires over the provider core to emulate a LAN between the sites.
- Inter-AS Option B lite is supported.

FabricPath

FabricPath has been enhanced to include the following:

- Anycast Hot Standby Router Protocol (HSRP)
- An overload bit
- Support for multiple topologies (ISIS-MT)
- Layer 2 proxy learning
- Scale improvements

VDC

New VDC features in Release 6.2(2) are as follows:

- A new VDC type, F2e
- An Administrator VDC on the Supervisor 1 module
- Support for an F2e Series module and an M Series module in the same VDC

vPC and vPC+

Cisco NX-OS Release 6.2(2) includes the following new vPC and vPC+ features:

- A best practice macro for vPCs can be enabled in vPC domain configuration mode with the **mode auto** command. This feature enables the following vPC best practice features:
  - peer gateway
  - auto recovery
  - ip arp synchronize
  - ipv6 nd synchronize
- Release 6.2(2) supports Source-Specific Multicast (SSM) in a vPC+ domain.

OTV

The following Overlay Transport Virtualization (OTV) features are available in Release 6.2(2):

- The VLAN translation feature allows you to connect applications that reside in separate Layer 2 domains between data centers.
- Selective unknown Unicast flooding is a per MAC address configuration that allows OTV to flood across the DCI for the specified MAC address. This feature is particularly helpful for applications that go silent and timeout from the ARP tables.
- Dedicated broadcast group allows you to configure a separate multicast address for broadcast traffic. This feature is useful for organizations that need separate QoS policies for broadcast traffic.
- OTV has built-in BFD support that does not require any additional configuration on the OTV side, which helps with any reconvergence that OTV might have to handle.
- The scale of OTV and how fast it converges are improved in this release.
- F1 Series and F2e Series modules can be used as internal interfaces with the OTV VDC.

Routing Capabilities

Release 6.2(2) provides improved routing capabilities through these features:

- The Dynamic Host Configuration Protocol (DHCP) relay now supports IPv6.
- Bidirectional Forwarding Detection (BFD) has been enhanced to include a client for ISISv6, PIMv6, BGPv6, and OSPFv3.
- IPv6 logo phase 2 certification is available in this release.
- The Border Gateway Protocol (BGP) has been enhanced to include flexible distance manipulation and an injection map.
- Virtual Router Redundancy Protocol (VRRP) v3 is supported.
- VRF scale enhancements are in this release.
- Private VLAN (PVLAN) support is provided for vPCs and port channels.
- Layer 2 scale improvements are a part of this release.

Security Features

Release 6.2(2) includes enhancements for numerous security features, including these:

- ACL ternary content address memory (TCAM) bank mapping allows TCAM banks to accommodate more feature combinations in a more predictable manner. By using this feature, you can optimize space and maximize the utilization of TCAM banks.
- Added support for the DHCPv6 relay agent. You can configure a device to run a DHCPv6 relay agent, which forwards DHCPv6 packets between clients and servers. This feature is useful when clients and servers are not on the same physical subnet.
- For Control Plane Policing (CoPP)
  - Changed the behavior of multicast traffic from being policed at different rates in different classes to being grouped into three classes (multicast-host, multicast-router, and normal) and policed at consistent rates.
  - Added the ability to monitor CoPP with SNMP using the CISCO-CLASS-BASED-QOS-MIB.

SPAN

SPAN enhancements include the following:

- (2+6) bidirectional or 2 bidirectional and 12 unidirectional SPAN sessions for F1, F2, F2e, and M2 Series I/O modules.
- A rule-based SPAN filter.
- An F2 Series module or an F2e Series module provides Encapsulated Remote Switched Port Analyzer (ERSPAN) termination.

LISP

LISP includes these new features:

- The LISP Instance ID supports virtualization and provides a means of maintaining unique address spaces (or address space segmentation) in the control and data plane.
- The LISP Delegated Database Tree (DDT) defines a large-scale distributed database of LISP Endpoint Identifier (EID) space using a DDT node.
- Support for multicast.

IPSLA

IP service level agreement (IPSLA) features enhancements include PBR object tracking, ICMP, ECHO, and DNS.

FEX

Cisco Fabric Extenders support DSCP to queue mapping, and Layer 3 protocol adjacencies on host interfaces (HIFs). Queuing support is not available in this release for the Cisco Nexus 2248PQ-E Fabric Extender.

ISIS MT

This release introduces support RFC 5120 - M-ISIS, Multi-Topology (MT) Routing in NXOS IS-IS. Multitopology (MT) ISIS allows you to define a set of independent topologies for different protocols within a single IS-IS domain. Cisco NX-OS supports one topology for IPv4 and one topology for IPv6.

Confirm Commit Device Alias

The confirm commit feature is enabled by default. In the previous version of code prior to 6.2(10), the dialogue information is not shown, but in the later versions, it’s shown by default.

Cisco NX-OS Release 6.2(2a) Software Features

Cisco NX-OS Release 6.2(2a) includes the new features described in the following sections:

- RISE, page 43
- BFD for IPv6 Static Routes, page 43

- Static Route to a VLAN, page 43
- USGv6 and IPv6 Phase 2 Ready Logo, page 43
- FIPS, page 43
- Dynamic Fabric Automation (DFA), page 43

RISE

The Remote Integration Service Engine (RISE) architecture logically integrates the Citrix NetScaler appliance and the Cisco Nexus 7000 Series switch so that the Citrix NetScaler service appliance appears as a service module within the Cisco Nexus 7000 Series switch. RISE provides streamlined deployment, simplified configuration, and reduced operational costs for service appliances. RISE enables service integration with the Cisco Nexus 7000 Series switch virtual device context (VDC) architecture.

BFD for IPv6 Static Routes

In Cisco NX-OS Release 6.2(2a), you can configure BFD for all IPv6 static routes on an interface.

Static Route to a VLAN

A switch virtual interface (SVI) includes a static route to a VLAN in Cisco NX-OS Release 6.2(2a).

USGv6 and IPv6 Phase 2 Ready Logo

Cisco NX-OS Release 6.2(2a) includes features for USGv6 and IPv6 Phase 2 Ready Logo.

FIPS

Cisco NX-OS Release 6.2(2a) includes features that qualify this release for Federal Information Processing Standards (FIPS) 140-2 Level 1 certification for the Cisco Nexus 7000 Series. The release is planned to be submitted for certification in December 2013.

Dynamic Fabric Automation (DFA)

Cisco NX-OS Release 6.2(6a) is the first release to support the Cisco data center fabric solution called Dynamic Fabric Automation (DFA). DFA is evolutionary and is based on the industry leading Unified Fabric solution.

DFA focuses on simplifying, optimizing and automating data center fabric environments by offering an architecture based on four major pillars namely fabric management, workload automation, optimized networking and virtual fabrics. Each of these pillars provide a set of modular functions which can be used together or independently for easiness of adoption of new technologies in the data center environment.

For the Cisco Nexus 7000 Series devices, the F2, F2e, and F3 Series modules support this functionality. On the Cisco Nexus 7700 Series devices, the F2e and F3 Series modules are support this functionality.

Complete details on the DFA architecture can be found at http://www.cisco.com/c/en/us/solutions/data-center-virtualization/unified-fabric/dynamic_fabric_automation.html.

Cisco NX-OS Release 6.2(6) Software Features

Cisco NX-OS Release 6.2(6) includes the new features described in the following sections:

MVRP

The Multiple VLAN Registration Protocol (MVRP) is a VLAN membership protocol, in which end stations and bridges can issue or withdraw declarations related to the membership of VLANs. The attribute type is the 12-bit VLAN ID. MVRP is a pruning protocol that forbids the propagation of unknown broadcast and unknown multicast frames in the regions of the network that do not need to receive those frames.

Cisco NX-OS Release 6.2(6) provides support for this feature.

MAC Security

The MAC Security (MACSec) feature is used for encryption and decryption,

MACSec support is available on F2e Series modules in Cisco NX-OS Release 6.2(6), with the following caveats:

- F2 Series modules with copper interfaces—All ports support MACSec (N7K-F248XT-25E and N77-F248XT-25E).
- F2 Series modules with fiber interfaces—The last eight ports (41 to 48) support MACSec (N7K-F248XP-25E and N77-F248XP-25E).

VLAN Translation

VLAN translation provides flexibility in managing VLANs and data center-related services by allowing you to merge the two Layer 2 domains without actually changing the original VLAN number. For example, when two data centers are connected using some form of DCI such as OTV and reconfiguration is not worth the collateral damage it can cause.

Cisco NX-OS Release 6.2(6) provides support for VLAN translation. Per-port VLAN translation is supported on all hardware and software protocols.

OTV Support

Cisco NX-OS Release 6.2(6) provides support for Overlay Transport Virtualization (OTV) on the F3 Series modules. This feature is supported only in VDCs without OTV extended switched virtual interfaces (SVIs). VLAN translation and traffic depolarization are not supported.

OTV Traffic Depolarization

OTV traffic depolarization is enabled by default with Cisco NX-OS Release 6.2(6). Also, the OTV display shows the secondary addresses used by the overlay and adjacencies. You can disable route depolarization using the command-line interface (CLI). This is enabled for all modules except the F3 Series modules.

Interoperability Between Modules

With Cisco NX-OS Release 6.2(6), you can interoperate the F3 Series modules with other types of modules in the same VDC, in a lowest common denominator mode. The following combinations of modules are allowed:

- F3 Series modules with F2 Series and/or F2e Series modules
- F3 Series modules with M2 Series modules

In both cases, all modules perform full Layer 2 and Layer 3 forwarding based on their native capabilities (there is no proxy routing with F3 modules as on earlier M Series and F Series mixed VDCs). Therefore, only features that are common to all of the modules in the VDC are supported, and available hardware forwarding table sizes are generally equivalent to the smallest forwarding table sizes of any of the modules.

With Cisco NX-OS Release 6.2(6), you cannot interoperate the F3 Series plus the F2 and/or F2e Series plus the M2 Series modules in the same VDC.

FCoE Over Physical Port vPC

The Cisco NX-OS Release 6.2(6) supports Fibre Channel over Ethernet (FCoE) with the physical port virtual port channel (vPC) for the F2 Series and F2e Series modules. This feature is not supported on the F3 Series modules.

Physical Port vPC

The Cisco NX-OS Release 6.2(6) supports the physical port virtual port channel (vPC) for the F2 Series and F2e Series modules. This feature is not supported on the F3 Series modules.

Ingress NetFlow Sampling and DHCP Relay Together

With the Cisco NX-OS Release 6.2(6), you can configure ingress NetFlow sampling and DHCP relay on the same interface.

F3 Series and F2 Series Modules with Dynamic Fabric Automation

Both the F3 Series and the F2 Series modules can function with Dynamic Fabric Automation (DFA).

Automatic Bandwidth Adjustment for MPLS TE Tunnels

The automatic bandwidth adjustment for TE tunnels feature allows you to configure Multiprotocol Label Switching (MPLS) to automatically monitor and adjust the bandwidth allocation for TE tunnels based on their measured traffic load. The automatic bandwidth behavior changes the configured bandwidth in the running configuration. If automatic bandwidth is configured for a tunnel, TE automatically adjusts the tunnel's bandwidth. This feature is supported with Cisco NX-OS Release 6.2(6).

ACL Logged as a Permit or Deny Entry

The switch indicates if the logged ACL was a permit or deny entry with Cisco NX-OS Release 6.2(6).

Large SGT,DGT Pairs

The switch now supports downloading large SGT,DGT tables, with improved caching functionality.

ELAM Enhancement

With Cisco NX-OS Release 6.2(6), you do not have to specify the module to run Embedded Logic Analyzer Module (ELAM).

Cisco NX-OS Release 6.2(8) Software Features

Cisco NX-OS Release 6.2(8) includes the new features described in the following sections:

Display of I/O Rates

A new show command has been added to display only input/output rates for the interfaces:

show interface *ex/y* **counter brief load-interval** *load*

IPv4 Prefix over IPv6 in BGP

Beginning with Cisco NX-OS Release 6.2(8), BGP supports RFC 5549 which allows an IPv4 prefix to be carried over an IPv6 next hop. Because BGP is running on every hop and all routers are capable of forwarding IPv4 and IPv6 traffic, there is no need to support IPv6 tunnels between any routers. BGP installs IPv4 over an IPv6 route to the Unicast Route Information Base (URIB).

BGP Next Hop

By default, BGP puts itself as the next hop when announcing to an eBGP peer. When you enter the **set ip next-hop unchanged** command for an outbound route map that is configured for an eBGP peer, it propagates the received next hop to the eBGP peer.

BGP PIC Edge Active-Backup

Cisco NX-OS Release 6.2(8) introduces the **additional paths install backup** command which enables BGP to install the backup path to the routing table when you are using the BGP Prefix Independent Convergence (PIC) Edge feature.

BGP Prefix-Peer Wait Timer

Cisco NX-OS Release 6.2(8) introduces the **timers prefix-peer-wait** command that enables you to disable the peer prefix wait time so that there is no delay before BGP prefixes are inserted into the RIB.

Intelligent Traffic Director

Intelligent Traffic Director (ITD) is an intelligent, scalable clustering, and load-balancing engine that addresses the performance gap between a multiterabit switch and gigabit servers and appliances. The ITD architecture integrates Layer 2 and Layer 3 switching with Layer 4 to Layer 7 applications for scale and capacity expansion to serve high-bandwidth applications. ITD provides adaptive load balancing to distribute traffic to an application cluster. With this feature on the Cisco Nexus 7000 Series switch, you can deploy servers and appliances from any vendor without a network or topology upgrade.

LISP

Cisco NX-OS Release 6.2(8) introduces the LISP multi-hop mobility functionality.

LISP multi-hop mobility provides a mechanism to separate LISP host detection from Tunnel Router function, previously implemented in the same device: the Ingress Tunnel Router (ITR) and Egress Tunnel Router (ETR), also known as xTR. A LISP First-Hop Router (FHR) detects the presence of a dynamic host Endpoint Identifier (EID) and notifies the Site Gateway xTR, which registers the dynamic EID with a Map Server. The Site Gateway xTR performs Locator/ID Separation Protocol (LISP) encapsulation/decapsulation of the traffic from or to the dynamic EID to or from remote sites.

LISP supports redistributing host routes for servers discovered by LISP into Interior Gateway Protocol (IGP) via Open Shortest Path First (OSPF) protocol, Intermediate System-to-Intermediate System (IS-IS) protocol, Routing Information Protocol (RIP), Border Gateway Protocol (BGP), and Enhanced Interior Gateway Routing Protocol (EIGRP). LISP supports server detection based on receiving host routes updates on a Site Gateway xTR, using the same routing protocols listed above.

Beginning with Cisco NX-OS Release 6.2(8), the ITR supports load balancing based on map-cache weights for encapsulated packets, as specified in RFC6830. For example, if there are four locators in a map-cache entry, with the weights assigned as 30, 20, 20, and 10, the first locator destination gets 37.5 percent of the traffic, the second and third locators get 25 percent of the traffic, and the fourth locator gets 12.5 percent of the traffic.

OSPF Distribute List Enhancement

This enhancement allows you to filter next-hop paths for a given OSPF route from being programmed into the RIB using the **route-map** *map-tag* [**deny** | **permit**] command.

RISE Phase 2

The Remote Integrated Service Engine (RISE) makes a service appliance appear as a service module to the Cisco Nexus 7000 Series switches so that an appliance user can enjoy the same benefit of a service module’s simple configuration and operation. The RISE phase 2 feature set adds support for auto policy-based routing (APBR), including appliance high availability (HA) and vPC support.

RISE with NAM

This feature enables a direct connection between a NAM appliance and the Cisco Nexus 7000 Series switches, including the Cisco Nexus 7000 Series switches and Cisco Nexus 7700 Series switches, through the Remote Integrated Service Engine (RISE) mechanism.

RISE with FEX

Cisco NX-OS 6.2(8) has tested the Cisco Remote Integrated Services Engine (RISE) on the following Fabric Extenders, and the remaining FEX models are expected to work:

- Cisco Nexus 2248PQ-10GE Fabric Extender
- Cisco Nexus2248TP-1GE Fabric Extender

Python Enhancements

Beginning with Cisco NX-OS Release 6.2(8), you can change the Python socket operation to another VRF instance. It does not have to stay in the management VRF.

Additional enhancements to Python are as follows:

- You can use the import **hashlib** statement,
- You can import the logging module.
- Cisco now supports the Java Script Object Notation (JSON) module.

GOLD Corrective Action

You can configure the system to take disruptive action if the system detects a failure on specified health-monitoring online diagnostic tests.

OTV Tunnel Depolarization

Beginning with Cisco NX-OS Release 6.2(8), support is added for the OTV traffic depolarization feature on F3 Series modules. OTV traffic depolarization is enabled by default. The OTV display shows the secondary addresses used by the overlay and adjacencies. You can disable traffic depolarization using the **otv depolarization disable** command. See the *Cisco Nexus 7000 Series NX-OS OTV Configuration Guide* for more information.

Support for Increased Number of Policers

Beginning with Cisco NX-OS Release 6.2(8), you can have up to 4.096 class-maps per policy.

NAM Data Center Protocol Performance Improvements in FPGA 6.2

To take advantage of the Cisco Nexus 7000 Series NX-OS Release 6.2(8) performance improvements for Data Center protocols (VxLAN, FabricPath, OTV, LISP, Segment ID, VNTag, and FCoE), you can upgrade the FPGA image in your Cisco Nexus 7000 Series NAM-NX1. For details on how to upgrade your FPGA image, see the *FPGA/EPLD Upgrade Note for Cisco Prime NAM-NX1, 6.0*.

FabricPath Anycast

Prior to The Cisco NX-OS Release 6.2(8), The FabricPath Layer 2 IS-IS was advertising the anycast switch ID even with the overload bit set, which would incur longer convergence times for selected nodes. Beginning with Cisco NX-OS Release 6.2(8), the system does not advertise the configured anycast switch ID while the overload bit is set, which effectively improves the convergence times.

Cisco NX-OS Release 6.2(10) Software Features

Cisco NX-OS Release 6.2(10) includes the new features described in the following sections:

- LISP, page 50
- Enhancement for ACL TCAM Bank Mapping, page 50
- QoS Shared Buffering Queuing on F3 Series Modules, page 50
- GRE Tunnels on F3 Series Modules, page 50
- BPDU Guard for Pruned VLANs, page 50
- Enhancement to Online Diagnostics, page 50
- Control Plane Policing, page 50
- Native VLAN Tagging on Trunk Ports, page 51
- Remove FabricPath Headers from SPAN Packets, page 51
- SGT Support for F3 Series Modules, page 51

- Common Criteria, page 51
- Intelligent Traffic Director, page 47
- License Requirement MP-BGP RR for FabricPath (aka DFA), page 51

LISP

The ITR now supports dynamic Path MTU discovery for data packets received from local servers. It sets the Don’t Fragment (DF) bit value in LISP outer header and can adjust the MTU value used for encapsulation based on the particular Routing Locator (RLOC) destination capability.

Enhancement for ACL TCAM Bank Mapping

The **show hardware acces-list feature-combo** command is added to display the mix of classification features that customers plan to use and if they will work together. See the *Cisco Nexus 7000 Series NX-OS Security Command Reference* for more information.

QoS Shared Buffering Queuing on F3 Series Modules

You can split QoS buffers into dedicated and shared buffers. With only dedicated buffers based on the CoS value, one queue may have very high traffic even though memory associated with some of the other queues may be lying idle. The shared buffer pools address this problem.

The default is disabled for shared buffer queuing and it runs per port.

GRE Tunnels on F3 Series Modules

Support is added for GRE tunnels on F3 Series Modules.

BPDU Guard for Pruned VLANs

BPDUs are dropped if they are not in the allowed VLAN list and if BPDU Guard is enabled on the port. In Cisco NX-OS Release 6.2(10) and later releases, the port will be error disabled when an a BPDU is received on any VLAN and BPDU Guard is enabled on the port.

Enhancement to Online Diagnostics

A new nondisruptive health monitoring test per module, the Internal PortLoopback test, is added. This test verifies the integrity of packet paths from the supervisor to all front ports and runs while the ports are UP.

Control Plane Policing

You can classify and rate-limit IP unicast RPF failure packets by using the **match exception ip unicast rpf-failure** command. See the *Cisco Nexus 7000 Series NX-OS Security Command Reference* for more information.

Native VLAN Tagging on Trunk Ports

Beginning with Cisco NX-OS Release 6.2(10), you can specify whether control and data packets are tagged or untagged by using the **vlan dot1q tag native** command at the global configuration level, or the **switch port trunk native vlan tag** command at the port level. See the *Cisco Nexus 7000 Series NX-OS Interfaces Command Reference* for more information.

Remove FabricPath Headers from SPAN Packets

This release added new commands to remove the FabricPath header on traffic going to SPAN destinations.

SGT Support for F3 Series Modules

Support is added for security group tags (SGT) on F3 Series modules.

Common Criteria

Support is added for Common Criteria in this release.

Intelligent Traffic Director

Intelligent Traffic Director (ITD) has the following enhancements in Cisco NX-OS Release 6.2(10):

- Weighted load-balancing.
- Node-level standby.
- Layer 4 port load-balancing.
- Sandwich mode node-state synchronization across two VDCs on the same Cisco Nexus 7000 Series device.
- DNS probe.
- Start/stop/clear ITD statistics collection.
- The **show itd statistics** command output shows percentage of traffic.
- VRF support for the ITD service and probes.
- Scalability changes:
  - 32 ITD services per VDC.
  - 16 virtual IPs per ITD service.

License Requirement MP-BGP RR for FabricPath (aka DFA)

The MPLS license and MPLS feature set requirement for the Spine node with MP-BGP Route Reflector has been removed with Cisco NX-OS Release 6.2(10). You can configure **feature fabricpath-vpn** instead of feature-set MPLS and **feature mpls l3vpn** to enable the VPNv4/VPNv6 address families on the MP-BGP Route Reflector**.**

Cisco NX-OS Release 6.2(12) Software Features

Cisco NX-OS Release 6.2(12) includes the new features described in the following sections:

- Cisco Nexus 7706 FCoE Support, page 52
- ip dscp-lop Command, page 52
- VDC Command, page 52

Cisco Nexus 7706 FCoE Support

FCoE support has been added for the Cisco Nexus 7706.

ip dscp-lop Command

The following command has been introduced as a part of CSCup78075.

ip dcsp-lop

VDC Command

The following command has been added to prevent LASER from turning off when a link failure is encountered.

[**no**] **system default link-fail laser-on**

Once configured, the command will apply to every link failure that is detected in a VDC.

Note This feature is only supported on F3 line cards.

MIBs

This section contains the following topics:

- MIBs Added in Release 6.2(2), page 52
- MIB Updated in Release 6.2(6), page 53
- MIB Added in Release 6.2(8), page 53
- MIBs Added in Release 6.2(10), page 53

MIBs Added in Release 6.2(2)

Support for the following MIBs is added in Cisco NX-OS Release 6.2(2):

- CISCO-SWITCH-HARDWARE-CAPACITY-MIB
- CISCO-SWITCH-ENGINE-MIB
- CISCO-SWITCH-FABRIC-MIB
- CISCO-vPC-MIB
- CISCO-OTV-MIB

- CISCO-SWITCH-STATS-MIB
- CISCO-VDC-MIB
- CISCO-HARDWARE-IP-VERIFY-MIB
- CISCO-FABRICPATH-TOPOLOGY-MIB
- CISCO-VLAN-IFTABLE-RELATIONSHIP-MIB
- CISCO-L2L3-INTERFACE-CONFIG-MIB
- CISCO-BRIDGE-EXT-MIB
- CISCO-IF-EXTENSION-MIB (Enhancements)
- CISCO-IETF-VRRP-MIB
- IGMP MIB, IGMP-STD-MIB (RFC 2933)
- CISCO-IGMP-SNOOPING-MIB
- CISCO-MVPN-MIB
- Cisco Nexus Platform MIB
- LLDP MIB

MIB Updated in Release 6.2(6)

Support for the following MIB for the M2, F2, F2e, and F3 Series modules is added in Cisco NX-OS Release 6.2(6):

- CISCO-CLASS-BASED-QOS-MIB

MIB Added in Release 6.2(8)

Support for the following MIB for the F1, F2, F2e, F3, and M2 Series modules is added in Cisco NX-OS Release 6.2(8):

- CISCO-BGP-MIBv2 MIB

MIBs Added in Release 6.2(10)

Support for the following MIB is added in Cisco NX-OS Release 6.2(10):

- CISCO-SWITCH-QOS-MIB
- csfFabricCrcErrorNotif trap in CISCO-SWITCH-FABRIC-MIB

Licensing

Cisco NX-OS Release 6.2(2) includes the following changes to Cisco NX-OS software licenses:

- The MPLS feature license (N7K-MPLS1K9) includes support for VPLS and EoMPLS.

The following licenses are available for the Cisco Nexus 7718 chassis (N77-C7718) and Cisco Nexus 7710 chassis (N77-C7710):

- LAN_ENTERPRISE_SERVICES, N77-LAN1K9
- VDC_LICENSES, N77-VDC1K9,

- ENHANCED_LAYER2_PKG, N77-EL21K9
- STORAGE_ENT, N77-SAN1K9

For additional information, see the *Cisco NX-OS Licensing Guide*.

Limitations

This section describes the limitations in Cisco NX-OS Release 6.2 for the Cisco Nexus 7000 Series. It includes the following sections:

- Native VLAN Change Causes Link Flap, page 55
- MPLS over GRE, page 55
- MAC Address Is Not Learned Over Peer Link, page 55
- The no hardware ejector enable Command Is Not Recommended for Long-Term Use, page 55
- Saving VLAN Configuration Information, page 56
- Rebind Interfaces Command Is Not Automatically Executed When Replaying ASCII Configuration in Cisco NX-OS Release 6.2(x), page 56
- Fabric Module Removal on the Cisco Nexus 7700 Series, page 57
- Fabric Utilization on the Cisco Nexus 7700 Series, page 57
- MTU Changes Do Not Take Effect on FEX Queues, page 57
- Clearing FEX Queuing Statistics Is Not Supported, page 57
- Multicast Traffic Is Forwarded to FEX Ports, page 57
- F2 Connectivity Restrictions on Connecting Ports to a FEX, page 57
- ACL Capture on the Cisco Nexus 7000 Series Network Analysis Module, page 58
- Behavior of Control Plane Packets on an F2e Series Module, page 58
- Error Appears When Copying a File to the Running Configuration, page 58
- CISCO-TRUSTSEC-SXP-MIB Does Not Provide an Instance Number, page 58
- Reload with CTS Configuration on the Nondefault VDC Causes a syslog Message, page 58
- CTS Configuration Command Not Available, page 59
- ECMP Support in Hardware, page 59
- CLI Command for Breakout Capabilities, page 59
- WCCP Support in a Mixed Mode VDC, page 59
- DSCP Queuing with FEX and M1 Series Modules, page 59
- DHCP Snooping with vPC+ FEX, page 59
- Fabric Module Migration Errors, page 60
- Proxy Limitation for the N7K-F132XP-15 Module, page 60
- PONG in a vPC Environment, page 60
- PONG in a vPC Environment, page 60
- SVI Statistics on an F2 Series Module, page 60
- LISP Traffic, page 60

Native VLAN Change Causes Link Flap

Changing the native VLAN on an access port or trunk port will flap the interface. This behavior is expected.

MPLS over GRE

MPLS over GRE is not supported on F3 and M3 modules.

MAC Address Is Not Learned Over Peer Link

When port security is configured on an orphan port that belongs to a VPC VLAN, the MAC address does not sync up with the peer and this leads to flooding on the peer link.This behavior is expected. To overcome this issue do not use VPC VLAN port for port security configuration.

The no hardware ejector enable Command Is Not Recommended for Long-Term Use

The **no hardware ejector enable** command cannot be a configured command in both the startup configuration and the runtime configuration. This command is a debugging command and should not be configured for long-term use.

In Cisco NX-OS Release 6.2(6), if you have dual Sup2e supervisors and your configuration includes the **no hardware ejector enable** command, physically removing the active supervisor will cause the modules to reload.

To work around this limitation, do not physically remove an active supervisor. Instead, use the **system switchover** command to switch to the standby supervisor.

This applies only to the Cisco Nexus 7700 Series devices.

Saving VLAN Configuration Information

Because a VLAN configuration can be learned from the network while the VLAN Trunking Protocol (VTP) is in a server/client mode, the VLAN configuration is not stored in the running configuration. If you copy the running configuration to a file and apply this configuration at a later point, including after a switch reload, the VLANs will not be restored. However, the VLAN configuration will be erased if the switch is the only server in the VTP domain.

To work around this limitation, do one of the following:

- Configure one of the clients as the server.
- Complete these steps:
  - Copy the VTP data file to the bootflash: data file by entering the **copy** *vtp-datafile* **bootflash:***vtp-datafile* command.
  - Copy the ASCII configuration to the startup configuration by entering the **copy** *ascii-cfg-file* **startup-config** command.
  - Reload the switch with Cisco NX-OS Release 6.2(2) or a later release.

This limitation does not apply to a binary configuration, which is the recommended approach, but only to an ASCII configuration. In addition, this limitation applies to all Cisco NX-OS software releases for the Cisco Nexus 7000 Series.

Rebind Interfaces Command Is Not Automatically Executed When Replaying ASCII Configuration in Cisco NX-OS Release 6.2(x)

The **rebind interfaces** command introduced in Release 6.2(2), is needed to ensure the proper functionality of interfaces in certain circumstances. The command might be required when you change the module type of a VDC. However, because of the disruptive nature of the **rebind interfaces** command, for Releases 6.2(x) prior to Release 6.2(8), it is not automatically executed when replaying an ASCII configuration file. Beginning with Release 6.2(8), the **rebind interfaces** command is always automatically performed whenever necessary during the replay of an ASCII configuration file.

For those Releases 6.2(x) prior to Release 6.2(8), this limitation applies when only when all of the following conditions are met:

- The ASCII configuration file is replayed in the context of the default VDC or the admin VDC, and at least one VDC has an F2e Series or an F3 Series module listed as supported module types either before or after the replay.
- The **limit-resource module-type** commands listed in the ASCII configuration file requires that **rebind interfaces** command be performed.

To work around this limitation, take the following steps:

- Manually enter the **rebind interfaces** command wherever needed to the ASCII configuration file for replay.
- Enter the **rebind interfaces** command immediately after the you enter the l**imit-resource module-type** command.
- Ensure that the ASCII replay properly applies all interface configurations for all interfaces in the relevant VDCs.

Note If you boot up a switch without any startup configuration, this limitation might apply to an ASCII replay. The reason is that without a startup configuration, the default VDC might still have certain interfaces automatically allocated. Because of this possibility, follow the preceding approaches to work around the limitation.

Fabric Module Removal on the Cisco Nexus 7700 Series

When a fabric module is power cycled or removed momentarily during an online insertion and removal (OIR) from slot 5 or 6 on a Cisco Nexus 7700 Series switch, packet drops can occur.

Fabric Utilization on the Cisco Nexus 7700 Series

When traffic ingresses from a module on the Cisco Nexus 7700 Series switch at a rate much below the line rate, uniform fabric utilization does not occur across the fabric modules. This behavior is expected and reflects normal operation based on the fabric autospreading technology used in the Cisco Nexus 7700 Series switch.

MTU Changes Do Not Take Effect on FEX Queues

When you change the interface MTU on a fabric port, the configured MTU on the FEX ports are not configured to the same value. This issues occurs when the interface MTU changes on a fabric port.

The configured MTU for the FEX ports is controlled by the network QoS policy. To change the MTU that is configured on the FEX ports, modify the network QoS policy to also change when the fabric port MTU is changed.

Clearing FEX Queuing Statistics Is Not Supported

Cisco NX-OS Release 6.2(2) does not support clearing queuing statistics for FEX host interfaces.

Multicast Traffic Is Forwarded to FEX Ports

Multicast traffic that is sent to Optimized Multicast Flooding (OMF) Local Targeting Logic (LTL) is forwarded to FEX ports that are not part of the bridge domain (BD). This issue occurs when multicast traffic is sent to OMF LTL, which happens if an unknown unicast and flood occur when OMF is enabled.

FEX interfaces can support multicast routers, but OMF on those VLANs must be disabled. If there is a multicast MAC address mismatch on the VLAN, traffic will be flooded in the VLAN and will eventually reach the router behind the FEX port.

F2 Connectivity Restrictions on Connecting Ports to a FEX

If an ASCII configuration has incompatible ports, such as when the configuration is created with ports that are added to the FEX from different line cards or VDC type, the ports might be added without warnings.

When connecting F2 Series ports to the same FEX, make sure the VDC type is the same as in the source configuration that is being replayed.

ACL Capture on the Cisco Nexus 7000 Series Network Analysis Module

The Cisco Nexus 7000 Series Network Analysis Module (Cisco NAM-NX1) does not support ACL capture on a NAM interface in Cisco NX-OS Release 6.2(2).

Behavior of Control Plane Packets on an F2e Series Module

To support the coexistence of an F2e Series module with an M Series module in the same VDC, the F2e Series module operates in a proxy mode so that all Layer 3 traffic is sent to an M Series module in the same VDC. For F2e proxy mode, having routing adjacencies connected through F2e interfaces with an M1 Series module is not supported. However, routing adjacencies connected through F2e interfaces with an M2 Series module is supported.

Error Appears When Copying a File to the Running Configuration

Copying a file to the running configuration can trigger the following error:

```
"WARNING! there is unsaved configuration" message.
```

This issue can occur if the configuration contains SNMP related configurations to send traps or notifications, and if the file to be copied to the running configuration contains only EXEC **show** commands.

Enter **Yes** to the prompt “This command will reboot the system. (y/n)? [n] y.” There is no operational impact and no configuration loss when the switch reloads.

CISCO-TRUSTSEC-SXP-MIB Does Not Provide an Instance Number

The object ctsxSxpConnInstance does not provide the instance number of the CTS SXP connection. Currently this number is not maintained and therefore cannot be displayed.

Reload with CTS Configuration on the Nondefault VDC Causes a syslog Message

When Cisco Trusted Security (CTS) enforcement is enabled on VLANs and a VDC reload occurs, CTS tries twice to disable the enforcement on the VLANs. The second time, the following syslog message appears:

```
CTS-2-RBACL_ENFORCEMENT_FAILED:Failed to disable RBACL enf on vdc reload
```

This syslog message can be ignored for the VDC reload because the VLANs are deleted on reload and CTS also deletes the enforcement configurations for those VLANs.

CTS Configuration Command Not Available

The **no cts dev-id pswd dev-pswd** command is currently not supported in NX-OS software.

Once the **cts dev-id pass** command is configured, it can be replaced using the same command, but it cannot be deleted.

ECMP Support in Hardware

32-way ECMP is supported with F2 and F2e Series modules since Cisco NX-OS Release 6.2(2). It is supported for F3 Series modules since Release 6.2(6).

CLI Command for Breakout Capabilities

Cisco NX-OS Release 6.2(2) supports the **show interface breakout** command, and this command is visible on all modules even if the specific module does not support breakout capabilities. The following lists the software support requirements for the specific line cards that support breakout capabilities:

- N7K-M206FQ-23L module—Cisco NX-OS Release 6.2(2) is the minimum software requirement.
- N7K-F312FQ-25 module—Cisco NX-OS Release 6.2(6) is the minimum software requirement.
- N77-F324FQ-25 module—Cisco NX-OS Release 6.2(8) is the minimum software requirement.

Please refer to Table 5 for the list of supported optics.

WCCP Support in a Mixed Mode VDC

Web Cache Control Protocol (WCCP) redirect-in and redirect-out is fully supported in the Cisco NX-OS Release 6.2 in unmixed module VDCs. WCCP is also supported in mixed module VDC scenarios for most module combinations. For complete support details, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x* for complete information.

DSCP Queuing with FEX and M1 Series Modules

Differentiated services code point (DSCP) based queuing does not work for FEX uplinks to the 32-port 10-Gigabit Ethernet SFP+ I/O module (N7K-M132XP-12) or the 32-port 10-Gigabit Ethernet SFP+ I/O module XL (N7K-M132XP-12L). All FEX data traffic will be in the default queue.

This limitation applies only when a FEX is attached to ports on a N7K-M132XP-12 or N7K-M132XP-12L module. It does not affect COS based queuing.

DHCP Snooping with vPC+ FEX

DHCP snooping is not supported when the vPC+ FEX feature is enabled.

Fabric Module Migration Errors

When you remove a Fabric 1 module and replace it with a Fabric 2 module, errors might occur. On rare occasions, 1 to 10 packets can drop during the fabric module migration process.

To avoid this situation, enter the **out-of-service xbar** command before you remove each Fabric 1 module.

Once the Fabric 1 module is out of service, remove it and insert the Fabric 2 module.

Proxy Limitation for the N7K-F132XP-15 Module

When the 6-port 40-Gigabit Ethernet I/O module XL (M2 Series) (N7K-M206FQ-23L) acts as a proxy for more than 90 G traffic from the 32-port 10-Gigabit Ethernet I/O module XL (N7K-F132XP-15), packet drops can occur. You might experience this issue if ports are oversubscribed on the N7K-F132XP-15 F1 Series module.

PONG in a vPC Environment

There are two situations where PONG is not supported in a vPC environment:

- In a vPC environment, a PONG to an access switch or from an access switch might fail. To work around this issue, use the interface option while executing a PONG from an access switch to a vPC peer. The interface can be one that does not need to go over the peer link, such as an interface that is directly connected to the primary switch.
- When FabricPath is enabled and there are two parallel links on an F2 Series module, PONG might fail. To work around this issue, form a port channel with the two links as members.

SVI Statistics on an F2 Series Module

F2 Series I/O modules do not support per-VLAN statistics. Therefore, the **show interface** command will not display per-VLAN Rx/Tx counters or statistics for switch virtual interfaces (SVIs).

LISP Traffic

A Layer 3 link is required between aggregation switches when deploying LISP host mobility on redundant LISP Tunnel Routers (xTRs) that are part of a vPC. In rare (but possible) scenarios, failure to deploy this Layer 3 link might result in traffic being moved to the CPU and potentially dropped by the CoPP rate limiters.

Role-Based Access Control

- Beginning with Cisco NX-OS Release 5.2, you can configure role-based access control (RBAC) in the Cisco Nexus 7000 storage VDC using Cisco NX-OS CLI commands. You cannot configure RBAC in the Cisco Nexus 7000 storage VDC using Cisco Data Center Network Manager (DCNM). Note that RBAC in the storage VDC is RBAC for the Cisco Nexus 7000 Series switches, which is different from that for the Cisco MDS 9500 Series switches.

- RBAC CLI scripts used in Cisco MDS 9500 Series switches cannot be applied to the storage VDC configured for a Cisco Nexus 7000 Series switch.
- You cannot distribute the RBAC configuration between a Cisco MDS 9500 Series switch and the storage VDC configured for a Cisco Nexus 7000 Series switch. To prevent this distribution, make sure to assign RBAC in Cisco MDS and the Cisco Nexus 7000 storage VDC to different Cisco Fabric Services (CFS) regions.

Standby Supervisor Can Reset with Feature-Set Operation

The standby supervisor might reload when a feature-set operation (install, uninstall, enable, or disable) is performed if the HA state of the standby supervisor is not "HA standby" at the time of the feature-set operation. To prevent the reload, ensure that the state of the standby supervisor is "HA standby." To check the HA state for the specific VDC where the feature-set operation is performed, enter the **show system redundancy ha status** command on the active supervisor.

A reload of the standby supervisor has no operational impact because the active supervisor is not affected.

In addition, if you perform a feature-set operation while modules are in the process of coming up, then those modules are power cycled. Modules that are up and in the "OK" state are not power cycled when you perform a feature set operation.

Unfair Traffic Distribution for Flood Traffic

Uneven load balancing of flood traffic occurs when you have a seven-member port channel. This behavior is expected and it occurs on all M Series and F Series modules. In addition, M Series modules do not support Result Bundle Hash (RBH) distribution for multicast traffic.

BFD Not Supported on the MTI Interface

If bidirectional forwarding detection (BFD) on protocol independent multicast (PIM) is configured together with MPLS multicast VPN (MVPN), the following error might appear:

```
2012 Jan  3 15:16:35 dc3_sw2-dc3_sw2-2 %PIM-3-BFD_REMOVE_FAIL:  pim [22512]  Session
remove request for neighbor 11.0.3.1 on interface Ethernet2/17 failed (not enough memory)
```

This error is benign. To avoid the error, disable BFD on the multicast tunnel interface (MTI) interface.

DOM Support

Diagnostic Optical Monitoring (DOM) is supported on the F3 Series modules in Cisco Nexus Release 6.2(8).

N77-F348XP-23: 1 Gigabit Ethernet Support

On the Cisco Nexus 7700 48-port 1/10-Gigabit Ethernet SFP+ I/O module (N77-F348XP-23), 1 Gigabit Ethernet port speed is not supported on the F3 Series modules in Cisco Nexus Release 6.2(6). This is supported beginning in Cisco NX-OS Release 6.2(8).

Level 4 Protocol Entries on the M Series Modules

There is a limitation of using 7 entries for Level 4 protocols on the M Series modules.

DFA Conversational Learning Route Aging in Bidirectional Traffic Scenario

Conversational learning is used to scale more routes by deleting aged-out routes from hardware thereby providing space for new routes that are active in conversation.

In the case of bidirectional traffic from a source to a destination route, the route will not age out unless both of the traffic streams are stopped. In the case where only one stream in a bidirectional stream is stopped, the route never ages out because the ternary content addressable memory (TCAM) entries are used for both route lookup on the destination and the Reverse Path Forwarding (RPF) check on the source. If you delete the entry based on one direction, RPF functionality configured on the reverse flow does not work. This is the reason our current ASICs marks a TCAM hit if a TCAM lookup happens in both directions (when an IP address lookup is done for the destination address and for the source address).

Caveats

This section includes the following topics:

- Open Caveats—Cisco NX-OS Release 6.2, page 62
- Resolved Caveats—Cisco NX-OS Release 6.2(12), page 114
- Resolved Caveats—Cisco NX-OS Release 6.2(10), page 142
- Resolved Caveats—Cisco NX-OS Release 6.2(8b), page 210
- Resolved Caveats—Cisco NX-OS Release 6.2(8a), page 211
- Resolved Caveats—Cisco NX-OS Release 6.2(8), page 212
- Resolved Caveats—Cisco NX-OS Release 6.2(6b), page 227
- Resolved Caveats—Cisco NX-OS Release 6.2(6a), page 228
- Resolved Caveats—Cisco NX-OS Release 6.2(6), page 229
- Resolved Caveats—Cisco NX-OS Release 6.2(2a), page 248
- Resolved Caveats—Cisco NX-OS Release 6.2(2), page 251

Note Release note information is sometimes updated after the product Release Notes document is published. Use the Cisco Bug Toolkit to see the most up-to-date release note information for any caveat listed in this document.

Open Caveats—Cisco NX-OS Release 6.2

- CSCuq39448

  Symptom: In A Nexus 5000, 6000, or 7000 switch may reload unexpectedly due to an EIGRP process crash. Reset reason will look similar to the following:

```
show system reset-reason`
----- reset reason for Supervisor-module 1 (from Supervisor in slot 1) ---
1) At 306358 usecs after Wed Sep  3 04:32:21 2014
 Reason: Reset triggered due to HA policy of Reset
 Service: __inst_001__eigrp hap reset
 Version: 7.0(2)N1(1)
```

Conditions: In Cisco NX-OS Release 6.2(12), this crash occurs when a **distribute-list** is configured on a physical interface or SVI. It will occur when a new neighborship is formed on that interface.

Workaround: None.

This issue will be rectified in Cisco NX-OS Release 6.2(14).

- CSCuq12660

  Symptom: If active Supervisor is running an NPE image, issuing the **reload module** *standby-Sup* **force-dnld** command will sync up a non NPE image to the standby Supervisor.

  The two Supervisors will form a full HA-standby, however, one is running a NPE image while the other is running a non NPE image.

  Conditions: This issue might be seen when you perform a Supervisor netboot or when the new Supervisor isn't preloaded with NPE image and it netboots a non-NPE image from active Supervisor.

  Workaround: perform the following tasks:

  1. Manually sync up the NPE image to the standby Sup.
  2. Set up the boot var to point to the NPE image and save the configuration.
  3. Issue the **reload module** *standby-Sup* command.

- CSCuq18447

  Symptom: One or more of the following syslog messages is displayed:

```
%RADIUS-3-RADIUS_ERROR_MESSAGE: Dropping response (packet ID 143) from server x.x.x.x
%RADIUS-3-RADIUS_ERROR_MESSAGE: RADIUS server x.x.x.x failed to respond
%RADIUS-3-RADIUS_ERROR_MESSAGE: RADIUS server x.x.x.x failed to respond evenafter all
retries
```

  Conditions: This issue might be seen if AAA authentication and accounting are configured for Radius and your Cisco Nexus 7000 Series switch is running Release 6.0(2).

  Workaround: Removing the **aaa accounting default group** command from the configuration to stop the traffic from being sent to the radius server.

  Also, adding the svc-isan user to the ACS server might resolve the issue.

- CSCui20722

  Symptom: The Cisco Nexus 7000 Series Network Analysis Module (Cisco NAM-NX1) and supervisor inband packet counters increment at a higher rate than expected.

  Conditions: This symptom might be seen only if all of the following conditions are met.

  - The monitor session source is supervisor inband.
  - The monitor session destination is Cisco NAM-NX1.

  - GOLD is enabled within the chassis.

  If the monitor session source includes supervisor inband RX, it is possible for internal diagnostic packets to continuously loop from the supervisor to the Cisco NAM-NX1 to the supervisor.

  Workaround: Configure a VLAN filter on the monitor session so that only the desired VLAN or VLANs are spanned. If you wish to monitor all traffic, use 1 to 3967,4048 to 4093 as the filter.

- CSCue53247

  Symptom: There are invalid entries for the number of routers and the number of cache engines in the output of the **show ip wccp** command.

  Conditions: This issue might be seen in scale configurations, but it is not consistent. After a reload or loss of service, the issue might occur.

  Workaround: Disable the feature and enable it again.

- CSCuf95718

  Symptom: High CPU utilization can occur on the active supervisor.

  Conditions: This symptom might be seen when Virtual Private LAN Service (VPLS) or EFP Ethernet virtual circuits are configured. The pktmgr process on the active supervisor registers high CPU utilization because it is busy dropping STP BPDUs that are entering the virtual circuit from the customer edge-side device.

  Workaround: Prevent STP BPDUs that originate from the customer-edge device from entering the provider-edge device where Ethernet virtual circuits/VPLS are configured.

- CSCuh11224

  Symptom: An egress RACL is configured on a switch virtual interface (SVI) and pushed to all line cards, even though there are no members associated with the VLAN (both the default and nondefault VLAN).

  Conditions: This symptom might be seen when SVI egress policies are pushed to all line cards even though the member count is zero. There is no functional impact.

  Workaround: None.

- CSCuh23543

  Symptom: MAC addresses are not synchronized between vPC peers after a switchover in both vPC peers.

  Conditions: This symptom might be seen when a switchover (or a Layer 2 FM process restart) occurs on one peer, and a switchover (or a Layer 2 FM process restart) occurs immediately within 5 to 6 seconds on the other peer. With a sizable MAC address table scale (greater than approximately10,000 addresses), the symptom occurs.

  Workaround: Enter the **clear mac address-table dynamic** command. If the switchover occurs on one peer and the local MAC database recovery is complete, the other peer does not have the issue. Enter the **show mac address-table count** command to confirm that the local MAC database recovery is complete.

- CSCuh53048

Symptom: When VPLS encapsulates L2PT encapsulated BPDUs, the EXP field in the MPLS header might be set to 0.

Conditions: This symptom might be seen if the BPDU that is L2PT encapsulated does not have a 802.1Q header.

Workaround: None.

- CSCuh94778

  Symptom: In a scale setup with a high number of logical ports, many processes compete for the CPU immediately after a system switchover. As a result, the Spanning Tree Protocol (STP) has less CPU to be able to rebuild its own database and to send its time-critical BPDUs every 2 seconds. This situation causes a CBL port state change and traffic drop.

  Conditions: This symptom might be seen when a large number of logical ports are in either Rapid Spanning or Multiple Spanning tree configurations.

  Workaround: Make the scale lower on a single switch.

- CSCui07745

  Symptom: The following error message is displayed after an ISSU:

  ```
  %NETSTACK-3-IPV6_MTS_DROP:  netstack [9694]  Error returned from mts_drop(), errno:
  Invalid argument
  ```

  Conditions: This symptom might be seen on a Cisco Nexus 7700 Series switch.

  Workaround: None. There is no functional impact.

- CSCui08461

  Symptom: Ports can remain in a suspended state. The output of the **show cdp neighbors** command for these ports appear as neighbors to themselves.

  Conditions: In Cisco NX-OS Release 6.1(3), an issue exists where a port can go into a suspended state when the online diagnostics feature fails to reset the port programming meant for the port loopback tests. This situation can occur in a rare scenario when online diagnostics is running port loopback tests and the same ports are being allocated to a different VDC.

  When the device in this state is upgraded to a newer release using an ISSU, the same issue also occurs in the newer release.

  Workaround: Manually run the GOLD tests again on the particular line card or port to reset the state of the port and bring it back to a functional state.

- CSCui08526

  Symptom: The traffic rate on the NAM module data port 1 momentarily drops to almost zero.

  Conditions: This symptom might be seen immediately after an ISSU or SSO completes.

  Workaround: None.

- CSCui20722

  Symptom: The Cisco Nexus 7000 Series Network Analysis Module (Cisco NAM-NX1) and supervisor inband packet counters increment at a higher rate than expected.

Conditions: This symptom might be seen only if all of the following conditions are met.

- The monitor session source is supervisor inband.
- The monitor session destination is Cisco NAM-NX1.
- GOLD is enabled within the chassis.

If the monitor session source includes supervisor inband RX, it is possible for internal diagnostic packets to continuously loop from the supervisor to the Cisco NAM-NX1 to the supervisor.

Workaround: Configure a VLAN filter on the monitor session so that only the desired VLAN or VLANs are spanned. If you wish to monitor all traffic, use 1 to 3967,4048 to 4093 as the filter.

- CSCui21769

  Symptom: After a peer link is shut in the secondary peer in a vPC+, some MAC addresses in a few VLANs can point to a vPC path that is operationally down. As a result, traffic might be silently dropped.

  In the output of the **show mac address-table** command, the interface column is empty and the output of the **show hardware mac address-table** command points the MAC address to the LID that is down.

  Conditions: This symptom might be seen when a vPC+ peer link is shut and there is an active traffic flow. Due to a race condition that can occur when the vPC paths in the secondary peer are brought down, the issue might be seen.

  Workaround: Enter the **clear mac address-table dynamic** command for the MAC address and VLAN where the issue occurs.

- CSCui22991

  Symptom: The queuing configuration becomes incorrect when a policy is removed.

  Conditions: The symptom might be seen only under the following conditions:

  - Enable the DSCP from the CLI.
  - Configure a user-defined DSCP to the ingress queue on an M2 Series module.
  - On an M2 Series module interface, apply a queuing policy that is consistent with the dscp2q mapping.
  - Change the dscp2q mapping to the default.
  - Remove the queuing policy from the M2 Series module interface.
  - ACLQOS and the hardware are incorrect.

  Workaround: Apply and remove any queuing policy on the affected interfaces.

- CSCui25889

  Symptom: When the **peer-gateway exclude** *vlans* command is used to add or remove VLANs, the MAC address table G bit does not get updated.

  Conditions: This symptom might be seen if the number of configured SVIs is greater than 3660.

  Workaround: To work around this issue, do one of the following:

  - Flap the peer link after the configuration changes.

  - Before entering the command, shut down some SVIs to make the total number of SVIs less than 3660.

- CSCui27283

  Symptom: The BRIDGE-MIB des not support Layer 2 VPN interfaces.

  Conditions: Pseudowire, EFP, and VFI membership are not supported in the BRIDGE-MIB.

  Workaround: None.

- CSCui27887

  Symptom: GISCM logs multiple SC_ONLINE messages like the following to the console for the same module after a switchover.

  ```
  2013 Jul 24 09:51:11 SITE1-Agg-6 %$ VDC-1 %$ %GISCM-2-AGNIBOOTSTRAP: MOD 4: SC ONLINE
  2013 Jul 24 09:51:11 SITE1-Agg-6 %$ VDC-1 %$ %GISCM-2-AGNIBOOTSTRAP: MOD 8: SC ONLINE
  2013 Jul 24 09:51:11 SITE1-Agg-6 %$ VDC-1 %$ %GISCM-2-AGNIBOOTSTRAP: MOD 4: SC ONLINE
  2013 Jul 24 09:51:11 SITE1-Agg-6 %$ VDC-1 %$ %GISCM-2-AGNIBOOTSTRAP: MOD 8: SC ONLINE
  2013 Jul 24 09:51:11 SITE1-Agg-6 %$ VDC-1 %$ %GISCM-2-AGNIBOOTSTRAP: MOD 4: SC ONLINE
  ```

  For modules 4 and 8, SC_ONLINE appears more than once.

  Conditions: This symptom might be seen only when a switchover occurs and the new active supervisor is coming up. It does not occur during normal operation.

  Workaround: None. There is no functional impact of these redundant logs.

- CSCui28437

  Symptom: In a scale setup when debug logs are enabled, a few flush requests occur on particular port channels.

  Conditions: This symptom might be seen in a scale setup where the access layer contains a set of Catalyst 4000 switches.

  Workaround: Avoid any kind of debug logging during a stateful switchover.

- CSCui30896

  Symptom: In a scale setup with a high number of logical ports and with multiple VDCs, the number of processes competing for the CPU immediately after a switchover increases with a factor of the number of VDCs. This situation does not allow the Spanning Tree Protocol (STP) enough of the CPU to send out its time-critical and sensitive bridge protocol data units (BPDUs).

  Conditions: This symptom might be seen when there are a large number of logical ports and multiple VDCs in either Rapid Spanning or Multiple Spanning Tree configurations.

  Workaround: Use multiple physical switches instead of VDCs for a large scale setup.

- CSCui33310

  Symptom: In a vPC setup, a peer-link flap will bring down the secondary vPCs. They come back up when the peer link is up. During the time the secondary vPCs are going down, if a switchover occurs and then the peer link comes up (in this corner case scenario), the secondary vPCs come up without active VLANs on them. A few vPCs might come up with active VLANs on them, or many or all vPCs might come up with no active VLANs on them.

Conditions: This symptom might be seen on a Cisco Nexus 7000 Series device with a vPC setup and 4000 VLANs in use in the vPCs. The peer link goes down and before all the secondary vPCs go down, a switchover occurs.

Workaround: When the secondary vPCs are up with no active VLANs on them, enter the **shut** command followed by the **no shut** command on those vPCs to recover all the active VLANs.

- CSCui41285

  Symptom: An error should display when the **show spanning-tree bridge-domain 100** command is entered because 100 is a VLAN.

  Conditions: This symptom might be seen when the **show spanning-tree bridge-domain 100** command is entered on a Cisco Nexus 7000 Series device.

  Workaround: Enter this command for only those values that apply. If 100 is a bridge domain, enter the **show spanning-tree bridge-domain 100** command. If 100 is VLAN, enter the **show spanning-tree vlan 100** command.

- CSCui45691

  Symptom: In a scale setup with debug logs enabled, a few flush requests occur on particular port channels.

  Conditions: This issue is seen in a scale setup where the access layer contains a set of Catalyst 4000 switches.

  Workaround: Avoid any kind of debug logging during a stateful switchover.

- CSCui49735

  Symptom: In a multi-VDC setup with a large number of VLANs, immediately after a stateful switchover there are many processes competing for the CPU. As a result, Spanning Tree Protocol (STP) has very little of the CPU to be able to rebuild its own database and send its time-critical BPDUs every 2 seconds and process the incoming BPDUs. This issue causes BPDU to time out and causes STP disputes for a short time.

  Conditions: This symptom might be seen when there are a large number of logical ports, particularly in Rapid Spanning Tree configurations.

  Workaround: Try to use Multiple Spanning Tree and if that does not work, lower the scale on a single switch.

- CSCui53933

  Symptom: In a large scale vPC setup when a peer link port channel (with many member links) is shut on one peer or one of the peers is reloaded, the other peer experiences an Spanning Tree Protocol (STP) bridge protocol data unit (BPDU) timeout.

  Conditions: This issue might be seen when the **shut** command is entered on the peer link (or it is brought down due to a peer reload) in a large scale vPC setup. In this situation, there were 4000 VLANs in the setup.

  Workaround: None.

- CSCui56786

Symptom: Virtual Fabric Interfaces (VFIs) do not transition from standby to active.

Conditions: This symptom might be seen if the vPC peer link is shut when the VFIs are configured (for a longer period than the pseudowire active wait timer which has a default value of 300 seconds), and the **no shut** command is entered on the peer link.

Workaround: Enter the **clear l2vpn service vfi all** command to correct the issue.

- CSCuj17443

  Symptom: The **inherit** command is not working with TACACS authorization enabled.

  ```
  switch(config)# interface port-channel1006
  switch(config-if)# inherit port-profile Blade-Servers
  ERROR: Failed to write VSH commands
  ```

  Conditions: This symptom might be seen with the Cisco Nexus7000 C7009 (9 Slot) Chassis ("Supervisor Module-1X") with Cisco NX-OS Releases 6.0(4) to 6.2(1) and Cisco ACS 5.3 patch 6

  Workaround: Remove the TACACS authorization commands.

- CSCtz15300

  Symptom: On a scaled setup with 20,000 multicast route entries, tracebacks and malloc failed messages are seen after repeatedly entering the **clear ip mroute *** command.

  ```
  2012 Sep 10 16:06:21 l3sys2-agg1-A3 %PIM-3-SLAB_LIB_SLAB_ERR: Slab error [double free
  attempted] in pim_routetype
  2012 Sep 10 16:06:21 l3sys2-agg1-A3 %PIM-3-SLAB_ERR: -Traceback: librsw.so+0x9ee8f
  librsw.so+0x9f173 0x81822ac 0x81c0d5d 0x812f533 librsw.so+0xb76ee librsw.so+0x9bfef
  librsw.so+0xa6bdf libpthread.so.0+0x6140 libc.so.6+0xca8ce
  2012 Sep 10 16:06:21 l3sys2-agg1-A3 %PIM-2-SLAB_LIB_SLAB_ELEM_ERR: Slab element Alloc
  PC: 0x819cbef, Element index: 1097
  ```

  Conditions: Scaled set-up, messages are displayed when clear is applied in quick successions.

  Workaround: Clear it once and wait for the convergence before clearing again.

- CSCui52204

  Symptom: The port numbers displayed from the **show hardware internal statistics pktflow** command output for the XBAR ASIC (SAC) are not correct because of different port mappings across the two SAC ASICs.

  Conditions: This symptom might be seen when you are working with an F3 Series module.

  Workaround: None.

- CSCul56620

  Symptom: CDP does not work on the Layer 3 side when an F2 Series or M1 Series module is configured in the Layer 3 >> Layer 2 fashion.

  Conditions: This symptom might be seen when an F2 Series or M1 Series module is configured in the Layer 3 >> Layer 2 fashion.

  Workaround: Configure the Layer 2 side as access ports.

- CSCum06140

Symptom: A VDC in which the NAM resides has EIGRP neighbors configured on the interface VLAN. When you are reloading the NAM module, those SVIs will see the EIGRP neighbor go up and down when the NAM comes online.

Conditions: This symptom might be seen on SVI interfaces with EIGRP neighbors configured. Those SVIs with EIGRP configured but no neighbors are not affected. The EIGRP neighbors on the Layer 3 physical interface are not affected.

Workaround: Configure the EIGRP neighbor on a Layer 3 physical interface instead of an SVI.

- CSCuj72242

  Symptom: Unicast flood MAC addresses are missing from unicast-only overlay VLANs on all the remote edge devices.

  Conditions: This symptom might be seen when you are clearing the OSPF neighbors multiple times.

  Workaround: Bring the OTV overlay down and up.

- CSCul77115

  Symptom: Copy run start fails with the following error message

  ```
  SYSMGR-3-CFGWRITE_SRVFAILED  Service "xmlma" failed to store its configuration
  (error-id 0x41470001).
  SYSMGR-2-CFGWRITE_ABORTED  Configuration copy aborted.
  SYSMGR-3-CFGWRITE_FAILED  Configuration copy failed (error-id 0x401E0000).
  ```

  Conditions: This symptom might be seen when you are running Cisco NX-OS.

  Workaround: Reload the device.

- CSCum03543

  Symptom: Hosts using GLBP might lose connectivity after the forward timeout expires.

  Conditions: This symptom might be seen when you are running GLBP.

  Workaround: None.

- CSCum04595

  Symptom: In a vPC setup, if an NLB packet comes in through the Cisco Nexus 700 Series switch 1 and the Cisco Nexus 7000 Series switch 2 on the peer-link, the packet is flooded on Cisco Nexus 7000 Series switch 2 (and vice versa).

  Conditions: This symptom might be seen if you are working with the Cisco NX-OS Release 6.1(4).

  Workaround: None.

- CSCum17554

  Symptom: The switch reboots with the reset-reason "vshd hap reset" when using regex in an EEM applet.

  Conditions: This symptom might be seen when you are using any illegal characters at the end of the string.

Workaround: Avoid using illegal characters on the end of the command. (With EEM in the Cisco NX-OS software, all keywords must be expanded and only the * symbol can be used for argument replacement.)

- CSCum80422

  Symptom: You might see a kernel-related syslog message while you are reloading the switch or performing ISSU to Cisco NX-OS Release 6.2(6a).

  Conditions: You might see this symptom after you perform a reload of the ISSU

  Workaround: None.

- CSCum77431

  Symptom: When you apply an IGMP report policy on a switched VLAN interface (SVI), the policy does not block any joins that were established before the application of the policy.

  Conditions: You might see this symptom when you have an IGMP report policy configured on an SVI.

  Workaround: Bring down the SVI and bring it up again.

- CSCun10505

  Symptom: WCCP redirection might fail after you perform an ISSU from Cisco NX-OS Release 6.1(1) to Release 6.1(4) to Release 6.2(6as7).

  Conditions: You might see this symptom when you are working with bank mapping and preform an ISSU from Cisco NX-OS Release 6.1(1) to Release 6.1(4) to Release 6.2(6as7).

  Workaround: Remove the WCCP policy and reapply it.

- CSCun74253

  Symptom: RISE DIRECT or vPC mode service might not come up.

  Conditions: You might see this when you are working with VLAN dot1q TAG NATIVE enabled on the switch.

  Workaround: Disable VLAN dot1q TAG NATIVE. Ensure that your switching is not disrupted by disabling this feature.

- CSCun60330

  Symptom: The Netstack core might be observed if a remote race condition is hit.

  Conditions: In rare instances, you might see the Netstack core if the following conditions occur when you have an interface in a VRF with an IP command do a breakout:

  - As part of the breakout, no version of IP commands are received by Netstack (and this issue happens before the system moves the interface on which it is configured to the correct VRF).
  - The system moves the interface to the correct VRF.
  - The status of the interface in question changes as part of breakout.

  Workaround: Netstack process comes back up.

- CSCuj20185

  Symptom: When you issue the **clear ip ospf** *id* **neighbor** command, you may see the ospf neighbor flap.

  Conditions: In rare instances, you might see this when you have BFD enabled with OSPF.

  Workaround: None; but the system recovers.

- CSCum61174

  Symptom: When a sub-mode which is in use is deleted, it might lead to a VSH process crash.

  Conditions: This issue might be seen in the very rare occasion when one user configures a given sub mode and second user deletes that submode at the same time. This behavor is expected.

  Workaround: None.

- CSCun93531

  Symptom: The **show policy-map interface brief** command does not show the inherited policies from port-channel logical interface to the physical ports.

  Conditions: You might see this symptom when custom queuing policies are applied on port channels.

  Workaround: Use the **show policy-map interface** *x/y* command instead.

- CSCun06941

  Symptom: VTP and CDP packets might not pass through Layer 2 VPNs.

  Conditions: You might see this symptom when you are working with Layer 2 VPNs.

  Workaround: None.

- CSCuo10896

  Symptom: A secure port might fail to learn the MAC address as a secure MAC address.

  Conditions: You might see this symptom when a MAC address moves from an unsecure to a secure port.

  Workaround: Clear the dynamic MAC address for the address in question.

- CSCun98035

  Symptom: The 1-Gigabit link might not come up on the N77-F348XP-23 module.

  Conditions: You might see this symptom when you have upgraded from Cisco Release NX-OS 6.2(6) or Release 6.2(6a) or the module is not reloaded with Release 6.2(8) or later images.

  Workaround: Reload the module with the Cisco NX-OS Release 6.2(8) module.

- CSCun06100

  Symptom: When you enter the **no shutdown command**, active RISE service might down or RISE service not come up.

Conditions: You might see this symptom when you create a new service with a conflicting RISE IP addresses.

Workaround: Change the IP address for the new RISE service.

- CSCun00920

  Symptom: You might see NX-OS log messages of the following type:

  ```
  ISCM-4-APBR_WARNING: RISE APBR: slot id: 397, reason: % PBR is already active on the
  interface with route-map _rise-system-rmap-Vlan1214
  ```

  Conditions: You might see this symptom when multiple PBRs are pushed by NetScaler for the same switched virtual interface (SVI).

  Workaround: None. This has no functional impact.

- CSCuo05808

  Symptom: You might see the RSVP core after an SSO test following reloading a module with TE/RSVP running.

  Conditions: You might see this symptom during an RSVP SSO test after you have reloaded a module running TE/RSVP.

  Workaround: None.

- CSCun99366

  Symptom: You might see the ingress dropped bytes counter value at zero or less than dropped packets.

  Conditions: This symptom might be seen when you are working in any configuration.

  Workaround: None.

- CSCuo30517

  Symptom: When you perform an ISSD from Cisco NX-OS Release 6.2(8) to Release 6.2(x). the switch reloads.

  Conditions: This symptom might be seen when you have LISP configured on the switch before the downgrade.

  Workaround: Remove LISP configuration prior to ISSD and reconfigure it after downgrade completes.

- CSCuo20562

  Symptom: You might see the sysmgr process go down after you perform an ISSU from Cisco NX-OS Release 6.2(6) to Release 6.2(8) and you apply a a 'copp profile dense' configuration.

  Conditions: This symptom might be seen after you perform an ISSU from Cisco NX-OS Release 6.2(6) to Release 6.2(8) with BFD, followed by a change in the copp profile configuration to dense mode. The problem is seen with SUP1 supervisor module and M Series modules.

  Workaround: Reload the M Series modules after you perform the ISSU with BFD enabled in the NX-OS Release 6.2(6) from which you are upgrading and before the copp profile dense configuration is applied.

- CSCun06187

  Symptom: When you are working with vPC configurations, you might see indirect RISE services with different name appear as active.

  Conditions: This symptom might be seen when you are working in vPC environment, and the primary and secondary switches have RISE service configured with different names.

  Workaround: Ensure that you have identical RISE service configurations on both vPC peers. NetScaler rejects different names for the service.

- CSCum10680

  Symptom: When you issue the **show rise** command, the display might show some APBR entries as having the next hop as NSIP instead of SNIP.

  Conditions: This symptom might be seen when the default route is defined as NS IP sub net. When a specific route to realer server subnet is removed, NetScaler uses the default route to reach realer server and install APBR entry with NS IP address as next hop. This is an issue when NetScaler is configured in HA pair, since NS IP is unique to each NetScale.

  Workaround: Do not configure a default gateway address in NS IP subnet.

- CSCuo19731

  Symptom: The RISE-NAM service shutdown might cause the data port configuration to be lost, and entering the **no shutdown** command does not bring back the configuration. Shutdown works as removing the service.

  Conditions: This symptom might be seen under the following conditions:

```
RISE-n77# sh rise
Name            Slot Vdc Rise-Ip         State        Interface
                  Id  Id
--------------- ---- --- --------------- ------------ ----------------
RISE-NAM-77     300  1   172.23.228.125  active       Eth1/39
RISE-n77# sh run int e1/41
!
interface Ethernet1/41  <===== Data port is configured
  switchport
  switchport mode trunk
  switchport monitor
  no shutdown

RISE-n77# conf t
RISE-n77(config)# service type rise name RISE-NAM-77 mode direct
RISE-n77(config-rise)# shut
RISE-n77(config-rise)#
RISE-n77(config-rise)# 2014 Apr  8 22:31:13 RISE-n77 %$ VDC-1 %$
%ISCM-2-RISE_SERVICE_INACTIVE: service 'RISE-NAM-77' (slot id 300) became inactive.

RISE-n77(config-rise)#
RISE-n77(config-rise)# sh run int e1/41
!
interface Ethernet1/41  <======= Data port config is gone after service is shut
  no shutdown
RISE-n77(config-rise)# no shut
RISE-n77(config-rise)# service type rise name RISE-NAM-77 mode direct2014 Apr  8
22:39:09 RISE-n77 %$ VDC-1 %$ %ISCM-2-RISE_SERVICE_ACTIVE: service 'RISE-NAM-77' (slot
id 300) became active.
```

```
RISE-n77(config-rise)#
RISE-n77(config-rise)#
RISE-n77(config-rise)#
RISE-n77(config-rise)# sh rise
Name            Slot Vdc Rise-Ip         State        Interface
                  Id  Id
--------------- ---- --- --------------- ------------ ----------------
RISE-NAM-77     300  1   172.23.228.125  active       Eth1/39
RISE-n77(config-rise)# sh running-config int e1/41
!
interface Ethernet1/41 <==== not restored after unshut
  no shutdown
RISE-n77(config-rise)#
```

Workaround: Enter the **data** *interface* command again in the RISE configuration submode.

- CSCun06330

  Symptom: When you are working with F2 and M1 Series modules in the same mixed-mode VDC with an interface configured under RISE and you add a SPAN-related configuration, you might not see the SPAN configuration removed when you remove the RISE configuration.

  Conditions: This symptom might be seen when you are working with F2 and M1 Series modules in the same mixed-mode VDC with an interface configured under RISE and you add a SPAN-related configuration.

  Workaround: Save the original configuration and reapply after you shut down the RISE service.

- CSCup70873

  Symptom: The DSCP2Q mapping is broken after you upgrade by reload.

  Conditions: This issue might be seen if you upgrade from Release 6.2(6) to Release 6.2(6a) by reload. This symptom persists even if you then upgrade from 6.2(6a) to Release 6.2(8).

  Workaround: First configure the **hardware qos dscp-to-queue ingress** command and then configure the **copy running-config startup-config** command.

- CSCuo90474

  Symptom: ISSU fails when a supervisor switch is upgraded and the F2 line card is not. leaving the port-client core files.

  Conditions: This symptom was seen on Cisco 7000 Series switches running Cisco NX-OS Release 6.1(3) to NX-OS Release 6.2(8).

  Workaround: Reload the configuration.

- CSCuo98502

  Symptom: A port-channel interface might not set the MTU correctly even though the **show run interface** command output shows the MTU configured under the port-channel interface. For example, jumbo frames may be hardware rate-limited even though the MTU is configured to allow jumbo frames.

  Conditions: This issue may be seen if the MTU is configured on the port-channel member interface before creating the port-channel interface. The MTU may not be set correctly on the port-channel interface itself.

Workaround: Configure the MTU on the actual port-channel interface instead of the member interface.

- CSCuo90941

  Symptom: The following error is returned when executing the **switchport trunk allowed vlan add** command:

  ```
  N7KA(config-port-prof)# switchport trunk allowed vlan add 6
  ERROR: This form of the command has no effect on the current system
  N7KA(config-port-prof)#
  ```

  Conditions: This issue was seen under the following conditions:

  - Port-profile being modified must have been created prior to the ISSU.
  - Cisco Nexus Series device must have been upgraded via ISSU from Cisco NX-OS 6.0(1) to NX-OS Release 6.2(2a), This condition is not present in switches upgraded via disruptive upgrade.

  Workaround: To avoid this symptom, delete and re-create the port profile from scratch after the ISSU. Rebooting the device is not an effective workaround as the condition remains after reloading the switch.

- CSCue77464

  Symptom: The MAC address is present in the hardware but is missing in the software database.

  Conditions: This issue has been seen on M1 Series modules.

  Workaround: Clear the MAC address-table.

- CSCup53102

  Symptom: Netflow doesn't account for 1:100 multiplier in Sampler Options.

  Conditions: This issue might be seen for Netflow in your Nexus 7000 switch with an F2, F2e, or F3-only VDC. This issue is not seen in M series modules where there is not additional/system default sampler. This issue is not seen in M series modules or an F3 module with a mixed VDC of F2/F2e modules.

  Workaround: None.

- CSCup59971

  Symptom: End hosts that are sending packets through the Cisco ASR 9000 Series routers and the Cisco Nexus 7000 Series switches experience packet loss in the range of 6-8 seconds because the switchover time from Secondary to Primary is a slow.

  Conditions: The MC-LAG is configured on the router-to-switch setup with a VPC.

  Workaround: Remove VPC on the switch and remove the **mlacp switchover maximize links** configuration on router.

- CSCup66561

  Symptom: NTP daemon starts crashing and the following messages are displayed:

```
2014 Jun 12 19:18:42 switch %$ VDC-1 %$ %SYSMGR-2-LAST_CORE_BASIC_TRACE: : PID 686
with message non-sysmgr(non-sysmgr) crashed, core will be saved .
```

```
2014 Jun 12 19:18:42 switch %$ VDC-1 %$ %SYSMGR-2-LAST_CORE_BASIC_TRACE: : PID 688
with message non-sysmgr(non-sysmgr) crashed, core will be saved .
```

Conditions: This issue might b seen after you configure one or more NTPv6 server.

Workaround: Remove the NTPv6 configuration and roll back to NTPv4.

- CSCup81717

  Symptom: The secondary VLAN shut does not take effect on the private-VLAN host and promiscuous ports. On a PVLAN host port and promiscuous port, traffic goes through to the egress ports

  Conditions: This issue might be seen when a secondary VLAN is shut but traffic is seen egressing out from the host or promiscuous ports. The trunk promiscuous and trunk secondary ports are fine.

  Workaround: None

- CSCup85222

  Symptom: A host may not be able to resolve ARP for its HSRP GW if it is single-homed to a VPC secondary device with a dual-active exclude interface-VLAN configuration is on the SVI that hosts the HSRP VIP.

  Conditions: This issue might be seen if the following conditions apply:

  - The Cisco Nexus 7000 Series switches are in a vPC.
  - The peer-keepalive is up and the peer-link is down.
  - The **dual-active exclude interface-vlan** command is configured for the SVI.
  - There are no other Layer 2 route between the VPC peers for the Vlan.
  - The VPC secondary is HSRP active.

  Workaround: None.

- CSCup97564

  Symptom: There is a difference in time between parent switch and FEX switch. After you clear the counter on FEX host interface port, the show-interface counter does not reset to 0 sec. Instead the counter resets to a value that is equal to the time difference between the parent switch and the FEX.

  Conditions: This issue might be seen if your FEX is a 2248PQ.

  Workaround: None.

- CSCuq05308

  Symptom: After you reload your switch or after you reload your nondefault VDCs, the breakout configuration and breakout ports for the nondefault VDC are out of sync.

  Conditions: This issue might be seen during any of the following scenarios:

  1. An interface breakout is done in the nondefault VDC and one of the following occurs before the switch is reloaded:

  – Either the **copy running-config startup-config** command is not issued in the nondefault VDC or the **write erase** command is not issued in the nondefault VDC

- The **copy running-config startup-config** command is issued in the default VDC.

2. An interface breakout is done in the nondefault VDC and the switch is running Release 6.2(8a) or an earlier release and one of the following occurs before the switch is reloaded:

- The **copy running-config startup-config** command is not issued in the nondefault VDC.
- The **copy running-config startup-config** command is issued in the default VDC.

Note In Release 6.2.10, you must issue the **copy running-config startup-config** command in the default VDC before your can issue the command in the nondefault VDC.

3. An interface breakout is done in the nondefault VDC and one of the following occurs before the switch is reloaded: Either the **copy running-config startup-config** command is not issued in the nondefault VDC or the command is issued in the default VDC and the **no interface breakout** command is issued on the interface that is broken.

Workaround: Perform the one of the following workarounds:

- To avoid the issue any time that a breakout/no breakout is done in a nondefault VDC: Either issue the c**opy running-config startup-config** command in both the default and nondefault VDCs or issue the **copy running-config startup-config vdc-all** command in the default VDC
- For scenario 2 (An interface breakout is done in the nondefault VDC and the switch is running Release 6.2(8a) or an earlier release): Issue the **no interface breakout** command on the port in the nondefault VDC for which breakout configuration and breakout ports are out of sync and issue the **interface breakout** command on the port in the default VDC for which breakout configuration and breakout ports are out of sync.
- Perform the following steps:

1. Power off the module on which the breakout configuration and breakout ports are out of synch.
2. Purge the running configuration on the module by using the **purge module** *module #* **running-config** command.
3. Power on the module and allocate the interfaces.
4. Issue either the **copy running-config startup-config** command in both the default and nondefault VDCs or issue the **copy running-config startup-config vdc-all** in the default VDC.

- CSCup43137

  Symptom: Timeout occurs, when VLAN mode changes for all VLANs from Fabricpath to non-Fabricpath.

  Conditions: Occurs in a configuration with Fabricpath scale and VLANs configured in all ports with FEX.

  Workaround: Flap all FEX ports that failed.

- CSCup70887

  Symptom: Flowcontrol statistics incrementing when TxPPP/RxPPP priority-flow-control signals being sent/received.

  Conditions: Occurs when per-priority pause is configured.

  Workaround: Use the 'show interface flowcontrol' command to verify TxPPP/RxPPP signals.

- CSCuq12104

  Symptom: MAC in VLAN 254 moves from 103.11.4514 to 103.11.65535.

  (4514 is the vpc LID and 65535 (0xffff) is the flood LID when the packet is snooped by the CPU)

  Conditions: Occurs for F1 modules when the source MAC addresses are incoming on VPC+.

  Workaround: Disable the snooping for the feature.

- CSCuq20660

  Symptom: 'show interface' command displays incorrect MTU on subinterface

  Conditions: Occurs when the network QoS based MTU is configured to changed the port MTU. (This is only a display issue. The hardware is programmed correctly.)

  Workaround: Apply the port MTU on the port-channel parent and then on the subinterfaces.

- CSCuq50149

  Symptom: No latency buffer to accommodate "no drop" behavior from FCoE with default FCoE qos policy.

  Conditions: Occurs when the F2/F2e linecards support the long-haul (>=2km) FCoE multihop.

  Workaround: Modify the default service policy support lossless FCoE over long distance.

- CSCuq50479

  Symptom: ORIB generates core file when killed.

  Also seen while simultaneously restarting the otv-isis process and polling the otv mibs.

  Conditions: Occurs intermittently.

  Workaround: N/A

- CSCuq69858

  Symptom: Layer 2 instabilities occur in 2 VDCs connected by VPC. Some BPDUs are lost.

  Conditions: Occurs in a large scale STP setup (800 VLANs with rapid-PVST).

  Workaround: Use MST instead of rapid-PVST (reduces number of BPDUs).

- CSCuq96822

  Symptom: Heartbeat looping and crashes due to timeout.

  Conditions: Occurs when 4000 VLANs are assigned to an interface.

  Workaround: Assign smaller number of VLANs.

- CSCur00775

  Symptom: For an SNMP clear, some SNMP related counters are not cleared.

  Conditions: Occurs for an M1 linecard.

  Workaround: N/A

- CSCur03111

  Symptom: Without MACSEC traffic passes fine and the customer can push roughly 8.6 gig of traffic on the link. However, When MACSEC is configured initially the connection is established as long as there is roughly 5 gigs of traffic or less the link stays established.

  At about 6.5 gigs of traffic, CRC errors may start being seen on the peer port. After that SAP rekey may fail and the MACSEC link go down. Link will stay down in SAP AUTHEN INCOMPLETE state.

  Conditions:

  1. MACSEC is enabled.

  2. Egress traffic rate exceeds the effective line rate for encrypted traffic.

  Workaround: Limit the traffic egressing out of MACSEC ports to less than effective 10 G linerate (including MACSEC header overhead). This may be roughly 6.5G for 64B packets.

- CSCun87659

  Symptom: Long time for VDC creation.

  Conditions: Occurs when the set up has 400+ VLANs and 500 layer 2 ports or port-channels.

  Workaround: N/A

- CSCun60330

  Symptom: Netstack core observed if a remote race condition occurs.

  Conditions: When an interface ethernet x/y belongs to a VRF X, there is an IP command present on the interface and a breakout is performed, a Netstack core occurs when the following events happen:

  - As a part of the breakout, no version of the IP commands are received by Netstack. This happens before the system moves the interface on which it is configured to the correct VRF.
  - The system moves the interface to correct the VRF of which it is a member.
  - The status of the interface in question changes as part of the breakout.

  Workaround: Netstack process comes back up.

- CSCuq92989

  Symptom: Input discards are observed on interfaces carrying FCoE traffic.

  Conditions: Occurs while performing ISSU (Cisco Nexus release 6.2(2) to Cisco Nexus release 6.2(8a)) with F2E linecards in an FCoE multihop environment with long-distance (2 km or greater) links.

  Workaround: N/A

- CSCur03681

  Symptom: Enabling hardware resource pooling results in error %ACLMGR-3-ACLMGR_PPF_ERROR: PPF error: DDB Error: 0x41170040 (ddb_srv_rtrmtln_remove/460)

Conditions: Occurs when hardware resource pooling is enabled after changing ACL currently applied to a WCCP policy.

Workaround: Do change ACLs applied to WCCP policies and do not enable resource pooling.

- CSCur11844

  Symptom: Links flap repeatedly until they reach link flap error-disabled status.

  Conditions: Occurs on Cisco Nexus release 6.2.8(a) when bringing up a 10G DWDM link between Cisco Nexus 7000 devices connected over a Siemens WAN.

  Workaround: Increase the link debounce timer starting at 500 ms and increasing it in 500 ms increments until the link stabilizes. Flap the link (on both ends) with the shut/no shut commands after each increase. (On the DWDM/UVN/WAN circuit make sure to disable the automatic link suspension (ALS).)

- CSCur16606

  Symptom: SVI traffic rate resets to 0.

  Conditions: Occurs when issuing the “sh inter vlan 200” command.

  Workaround: N/A

- CSCuj04442

  Symptom: Switchport trunk allowed VLAN configuration is not propagated to inherited interfaces.

  Conditions: Problem happens in large scale configuration (80+ interfaces in profile) when there are multiple add statements and the system is also processing interface up/down or peer link up/down.

  Workaround: Add VLANs on port manually.

- CSCuq01463

  Symptom: After flapping core link which is a port channel interface, EIGRP is not registering with BFD but remaining protocols such as PIM, OSPF, and ISIS are fine.

  Conditions: After flapping core link.

  Workaround: Restart EIGRP Neighborships (clear ip eigrp adj).

- CSCur17682

  Symptom: When Qos network template, Cos2q,DSCP2q mapping is changed with heavy congested traffic that can cause heartbeat failure on stats client and stat's client crashes. Also the configuration change might not be applied on all ports because traffic needs to be stopped to apply the configuration and due to congestion traffic was not stopped.

  Conditions: Heavy congestion on multiple ports and very corner case.

  Effected Image : 6.2.10

  Product : Nexus 7k and Nexus 7700

  module : applicable for all modules

  Workaround: Stop the traffic/prevent congestion when changing network qos Template,cos2q,dscp2q.

- CSCur18408

  Symptom: WCCP egress policy programming happened only on first instance and not programmed on the remaining instances.

  Conditions: Resource pooling and bank mapping CLI are toggled multiple times. WCCP policy is applied/removed/modified several times.

  Workaround: Remove the egress RACL, Egress WCCP policy. Apply it back.

- CSCur18519

  Symptom: On FEX setups or setups with ports in port fast mode, quick range VLAN delete/adds can cause missing GWMACs in Linecard.

  Conditions: The Linecard could be under stress due to previous command like VDC reload, PL flap. Under these conditions, issue is more likely.

  Workaround: When a range of VLANs are deleted and added, user must have a reasonable delay of a few minutes between the delete and add. Test command < test l2fm dump smac > can be used in case if error happens. Please contact l2fm before you run the test command.

- CSCur22182

  Symptom: snmpwalk aborts when performing a snmpwalk on CISCO-BGP4-MIB, when any bgp peer is configured with both VPNv4 Unicast AF(afi=1, safi=128), IPv6 Unicast AF (afi=2, safi=1) and/or IPv6 Multicast AF(afi=2, safi=2). Snmpwalk on the following tables are affected and will not be queried completely and aborts when querying VPNv4 unicast AF of any peer with above mentioned configuration:

  - cbgpPeerAddrFamilyTable

  - cbgpPeerAddrFamilyPrefixTable

  - cbgpPeer2AddrFamilyTable

  - cbgpPeer2AddrFamilyPrefixTable

  This issue will abort snmpwalk on cbgpPeer (1.3.6.1.4.1.9.9.187.1.2) or any oid tree that includes these tables.

  Example:

  .1.3.6.1.4.1.9.9.187.1.2.7.1.3.1.4.19.0.101.6.2.1 = IPv6 Unicast

  .1.3.6.1.4.1.9.9.187.1.2.7.1.3.1.4.19.0.101.6.1.128 = VPNv4 Unicast

  <snmpwalk aborts here>

  Conditions: Occurs when any bgp peer is configured with both VPNv4 Unicast AF(afi=1, safi=128), IPv6 Unicast AF (afi=2, safi=1) and/or IPv6 Multicast AF(afi=2, safi=2)

  Workaround: Since this issue breaks the snmpwalk to cbgpPeer and the other oid trees which queries the tables in CISCO-BGP4-MIB, snmpwalks can work only with oid trees which do not include CISCO-BGP4-MIB, under this configuration/scenario without aborting. The

CISCO-BGP4-MIB tables can be queried separately to get the information from the unaffected tables. Further, snmpget is not affected by this issue and will provide the correct outputs but can only be used to query specific oids.

- CSCur22627

  Symptom: EIGRP neighborship flaps after SSO due to authentication failure on the peer.

  Conditions: EIGRP authentication is configured on the interface.

  Workaround: Disable authentication or increase hello/hold timers on all the peering devices.

- CSCur22754

  Symptom: VLAN manager times out on PVLAN during deletion of 2000 primary VLANs.

  Conditions: To re-create, configure

  ```
  RJ-N7K-1-new-vdc(config)# vlan 1000-3000
  RJ-N7K-1-new-vdc(config-vlan)# mode fabricpath
  RJ-N7K-1-new-vdc(config-vlan)# private-vlan primary

  RJ-N7K-1-new-vdc(config)# no vlan 1000-3000
  ```

  Workaround: When there are 2000 primary VLANs, and if all of them need to be deleted, the workaround is to delete 1000 VLANs at once, instead of all the 2000 VLANs at once. For example in the above sequence, the following configuration will go through fine:

  ```
  RJ-N7K-1-new-vdc(config)# vlan 1000-3000
  RJ-N7K-1-new-vdc(config-vlan)# mode fabricpath
  RJ-N7K-1-new-vdc(config-vlan)# private-vlan primary

  RJ-N7K-1-new-vdc(config)# no vlan 1000-2000
  RJ-N7K-1-new-vdc(config)# no vlan 2001-3000
  ```

- CSCur23572

  Symptom: ETPM sequence timeout and interface err-disabled.

  Conditions: This happens during stress test after hundreds of interface changes, VLAN add removal.

  Workaround: None

- CSCur23775

  Symptom: ACL QoS process might crash on LC during PBR/QOS changes.

  Conditions: During high stress - ACL/QOS/PBR policies changes ACLQOS in LC might crash

  the steps required to hit this issue.

  Remove or add each of the below:

- VLAN

- QoS

- ACL from PBR

- delete add VLAN

- delete/add SVI

- delete/add inherit with QoS configuration

Workaround: None

- CSCur24427

  Symptom: VLAN MGR sequence timed out syslog seen on doing a delete for a range of PVLANS. This is seen only when the number of PVLANS is around 40 in the system.

  Conditions: This happens when there are 40 or more PVLANS configured and a range delete operation on VLANS is attempted.

  Workaround: Even though the event seq for PVLAN times out, the delete is successful. The VLANs are deleted from the system.

- CSCur25911

  Symptom: Stale FTAG entries in PIXM leading to wrong hardware programming with triggers.

  Due to stale entry in the PC, FTAG 0 entry will be cleaned up from the FTAG-ERBID table.

  This affects the traffic driving FTAG 0. Which results in packet drop on VPC leg.

  Conditions: VPC+ to VPC conversion in F2 based setup.

  Workaround: VDC reload or deconfigure and configure back vPCs.

- CSCur28138

  Symptom: ipfib process might core after ISSU from 6.2(8) to 6.2(10).

  Conditions: M1-XL or M2 series line cards when they have more than 200,000 prefixes in them.

  Workaround: Reload the module to recover from this error.

- CSCur28449

  Symptom: Core generated by vPC component

  Conditions: This core can be generated under scenarios when the vPC ports are coming up and failure happens in other parts of the system such as a linecard reload.

  It is triggered by a failure to respond in time by the sdb component.

  Workaround: No functional impact. The vPC process restarts and continues the processing it was doing.

- CSCur28450

  Symptom: Rollback fails the verification phase, saying "flowcontrol send on" is present in the running-config on the port-channel.

```
switch# rollback running-config checkpoint checkpoint_name
<Fails with following message>
Verification patch contains the following commands:
---------------------------------------------------
!!
interface port-channel###
  no flowcontrol send on
  exit
```

Conditions: When trying to rollback to a checkpoint where a current HifPC (a port-channel with FEX host interfaces as its members) becomes a simple port-channel (no FEX host interfaces as its members), rollback will fail the verification phase.

Workaround: Rollback running checkpoint checkpoint_name best-effort.

So that it will not do verification and will not revert back to the original running configuration.

And then do "no flowcontrol send on" on the affected interfaces

- CSCur29530

  Symptom: Static route is floated again after deleting static bfd configuration even we already did shutdown interface.

  Conditions: Device Info:

  SUP2

  F248XP

  Version: 6.2(6b)

  To remove static bfd configuration after shut down interface.

  Workaround: To remove "static bfd" first before shutdown interface, the problem symptom doesn't happen.

- CSCup45045

  Symptom: ITD configuration might cause a crash with Nexus switches.

  Conditions: This issue has been seen with several triggers such as enabling the peer, vdc switchovers, scaled testing, and shut/no shut of the service.

  Workaround: None

- CSCuq91958

  Symptom: On F2/F3 cards, at the tunnel termination point, the frames are replicated further towards the trunk interfaces. As the bug describes, the default CoS value will not be retained (what was set in tunnel head end) and instead would be observed to be 0.

  Conditions: The frames as observed on a downstream trunk port at the point of tunnel termination.

  Workaround: None

- CSCur30373

Symptom: On F3 module only:

We are seeing that SGT with only lower 8 bits of the configured SGT is derived from the IPFIB.

Example of SGT not being reflected.

```
n77k1-dc2-core1(config)# cts role-based sgt-map 10.21.10.56 21456
n77k1-dc2-core1(config)# sh logging ip access-list cache detail
SGT    Src IP       Dst IP     S-Port   D-Port   Src Intf        Protocol         Hits
-------------------------------------------------------------------------------------------
208    10.21.10.56     192.168.2.220     0      0       Ethernet1/35 (1)ICMP             4
```

Conditions: On F3 module only:

The issue occurs if using SGT value Source tag > 255

Workaround: None

- CSCur25124

  Symptom: Route updates from a Layer 2 VPN are not processed and there might be traffic loss. No further updates will be processed.

  Conditions: This problem might occur on an ISSU to a 6.2.10 version.

  Workaround: A SUP switchover will correct the problem.

- CSCuq91381

  Symptom: The programming in hardware is incorrect if IPv4 and IPv6 statements are used in a single route-map attached at pbr.

  Conditions: When a route-map attached in pbr contains both IPv4 and IPv6 statements.

  Workaround: Use separate route-maps for IPv4 and IPv6.

- CSCur32662

  Symptom: During an ISSU from 6.2.8a to 6.2.10, if there are ACLs applied on port channel members, the aclmgr will keep crashing on the supervisor card.

  Conditions: This issue happened during ISSU from 6.2.8a to 6.2.10, if there are ACLs applied on port channel members, the aclmgr will keep crashing. This can happen when we applied ACLs on some interfaces and later on we change these interfaces to port channel members. In this situation, due to an existing issue in 6.2.8a (which has been fixed in 6.2.10), the ACLs on the member interface will not be removed, which will cause the aclmgr crash during ISSU.

  Workaround: In old release (6.2.8a or earlier), if there are ACLs applied on port channel members, we should remove these ACLs before attempting the upgrade.

- CSCur32745

  Symptom: L3 route does not forwarded on Port-channel interfaces.

  Conditions: With two PC member from different Linecards

Having another L3 port in the same ASIC.

Removing one of the member from the PC causes LIF not published to SDB.

Route added later to the L3 Port-channel will fail.

Workaround: Flap the L3 port-channel after PC member removal.

- CSCur09585

  Symptom: There is false High voltage alarms on DWDM Sumitomo SFP 10G

  Conditions: Normal condition

  Workaround: None

- CSCur32239

  Symptom:

  Conditions: PVLAN configuration. VLAN range delete with range more than 2000 VLANs.

  Workaround: None

- CSCur34707

  Symptom: The Ingress Buffer (IB) block of a SoC on a Nexus 7000 F2 module might become stuck and will not pass packets towards the fabric.

  Conditions: This has been seen with (but is not limited to) the following software and hardware:

  - NX-OS version 6.2(8a) for the Nexus 7000

  - Nexus 7000 F2 Series Linecards

  - Nexus 7000 Supervisor 2

  - Nexus 7000 Fabric 2

  Workaround: Reload of the module has addressed this so far.

- CSCur44048

  Symptom: Multidestination traffic egressing a vPC+ port-channel in I state is dropped by the Nexus 7000

  Conditions: vPC+

  No lacp suspend-individual enabled on the port-channel

  Reload Nexus 7000

  vPC Port-channel members come up in 'I' state following reload

  Workaround: Toggle the vPC members between mode on and LACP enabled. Ensure that once configured for mode on, that both peers report the members in the 'P' state. At that point enable LACP on the members again. The interfaces will return to the 'I' state, However, now the tag_to_erbdg_or_dvif table will now be programmed properly.

- CSCur57084

Symptom: Nexus 2000 may fail to copy the core file to the Nexus 7000 during a crash but continues to try over and over:

N7k-2 SYSMGR-FEX101-3-CORE_OP_FAILED Core operation failed: send_msg_to_ccdmon: Could not send to CORE_DMON return -1 errno 32

N7k-2 SYSMGR-FEX101-5-SUBPROC_TERMINATED "System Manager (core-client)" (PID 1903) has finished with error code SYSMGR_EXITCODE_CORE_CLIENT_ERR (11).

Conditions: When the Nexus 2000 connected to a non-default VDC crashes.

Workaround: Contact Cisco TAC.

- CSCur61854

  Symptom: SNMP configuration loss after multiple snmpd crashes leading to hap-reset and supervisor failover.

  Conditions:

  Workaround: None

- CSCur63689

  Symptom: OSPF process may crash on N7K with a core file

```
 %SYSMGR-2-SERVICE_CRASHED: Service ?__inst_001__ospf " (PID 6381) hasn't caught signal 11 (core will be saved).
```

```
 show cores
 VDC  Module  Instance  Process-name     PID       Date(Year-Month-Day Time)
 ---  ------  --------  ---------------  --------  -------------------------
 1    9       1         ospf-1999        6347      2014-10-15 10:48:13
```

  Conditions: OSPF might crash after a routing update. For example, ospf cost change.

  This issue has been seen in 6.1(4a).

  Workaround: None

- CSCur66762

  Symptom: URIB might use (backup) routes with BGP PIC EDGE.

  Conditions: When BGP PIC EDGE is enabled and backup route has a better AD or metric.

  Workaround: None

- CSCur66914

  Symptom: The aclqos process may crash on a linecard.

  Conditions: This has been seen when changing the port speed.

  Workaround: None

- CSCur67753

  Symptom: Linecard reload during the switchover.

  Conditions: Not known.

  Workaround: None

- CSCur77001

  Symptom: STP BPDUs are not being sent shortly from a VDC holding STP Root Switch role when entering the **copy run start vdc-all** command in the admin VDC.

  Conditions: The problem could be seen when "copy run start vdc-all" is executed in the admin VDC.

  Workaround: Enter the **copy run start** command on each VDC separately.

- CSCur87739

  Symptom: FabricPath overload bit gets set to 1 with some consequences on the convergence.

  Conditions: Remove the active supervisor from the chassis.

  Workaround: To avoid this issue, it is best to make sure to remove the standby supervisor, and enter the **system switchover** command to make the active supervisor standby and then remove it from the chassis.

- CSCur93820

  Symptom: F2 module might reset due to an aclqos crash after IP ARP inspection is removed on the Cisco Nexus 7000.

  Conditions: The following command was entered: **no ip arp inspection vlan**

  Workaround: None

- CSCuq99456

  Symptom: At distributed port-channel (all member interfaces are located at M1 linecard), one of the linecard's mac address entry may become inconsistency unexpectedly and it will lead to packet drop.

  Conditions: N/A

  Workaround: When the problem happened, clearing MAC entry manually will recover.

- CSCur22130

  Symptom: Interface discards erroneously increment for some interfaces when polling via SNMP but those same discard counters don't increment for the CLI

  Conditions: There is no known trigger for this problem at this time. "Input Discards" or "Output Discard" will increment for SNMP but not when you do a 'show interface eX/Y'

  Workaround: config

  no snmp-server counter cache enable

  snmp-server counter cache enable

- CSCur22840

  Symptom: When macsec is configured, we see dropped packets on some interfaces with Bad CRC signature. This is seen after switchover and when linerate exceeds stipulated limit for macsec.

  Conditions: 1. MACSEC is enabled.

  2. egress traffic rate exceeds the effective line rate for encrypted traffic.

  3. Switchover is performed

  Workaround: Limit the traffic egressing out of Macsec ports to less than effective 10G linerate (including Macsec header overhead). This may be roughly 6.5G for 64B packets.

  Enable global pause on the ports in questions.

- CSCur60498

  Symptom: traffic rate of SPAN dest port is double when using "both" option

  Conditions: The traffic rate of SPAN destination port is twice as much as that of SPAN source port when using "both" option at monitor session.

  - when use "both" option.

  - input rate of SPAN source port become twice.

  - output rate of SPAN source port do NOT become twice. (even if use "both" option)

  Workaround: use "rx" option when monitor input traffic.

- CSCur97147

  Symptom: Fabricpath adjacency not coming up on core port

  Conditions: Following initial setup on 6.1(5a)

  Workaround: shut/no shut the impacted interface

- CSCur99327

  Symptom: Rollback to a previously created checkpoint might fail at "no license grace-period" command.

  Conditions: This only happened when performing rollback after "write erase" and on scale setup.

  This issue is not always reproducible.

  Workaround: Issue another rollback with "best-effort" option first. And do "no license grace-period" manually if necessary.

- CSCus04213

  Symptom: ISSU Failure

  Conditions: ISSU

  Workaround: None

- CSCus05123

  Symptom: Access-list applied on line vty stops working after a supervisor switch over.

  Conditions: This issue happens only on SUP2 & not on Sup 1 supervisor switchover.

  Workaround: Remove & reconfigure the ACL on line vty.

- CSCus06624

  Symptom: Monitor service crash on nexus switch upon show run command with monitor session in configuration

  ```
   Reason: Reset triggered due to HA policy of Reset
    System version: 6.1(2)
    Service: monitor hap reset
  ------
  VDC  Module  Instance  Process-name     PID       Date(Year-Month-Day Time)
  ---  ------  --------  ---------------  --------  -------------------------
  1    2       1         monitor          5254      2014-12-05 14:01:54
  1    2       1         monitor          28739     2014-12-05 14:01:59
  1    1       1         monitor          26093     2014-12-05 14:19:53
  1    1       1         monitor          5191      2014-12-05 14:19:56
  ```

  Conditions: SPAN configured on the switch

  Workaround: Avoid the show run command on the switch with monitor session configuration

- CSCus08870

  Symptom: After an ISSU to 6.2(10), traffic to some destination IP addresses which are passing through the F2 modules might be disrupted.

  Conditions: F2 module. ISSU from 6.2(8) or 6.2(8a) to 6.2(10)

  Workaround: Reload the module.

- CSCus00306

  Symptom: VLANs are not enabled on peer-link after ISSU 6.1(2)->6.2(10).

  Conditions: Direct ISSU 6.1(2) to 6.2(10).

  Workaround: Remove and re-add VLANs to configuration.

- CSCus13750

  Symptom: Traffic is received on WAAS / Proxy even when WCCP is off from Cisco Nexus 7000 switch.

  Conditions: WCCP flaps. In some cases BFD Move and LC installed.

  Workaround: Possible 2 workarounds.

A: "no feature wccp" this would work, if SPM is not crashed yet.

B: Reload the complete Switch, when SPM is crashed.

- CSCus18323

  Symptom: The traffic destined for CPU never reaches the CPU but dropped in the switch itself.ARP table goes incomplete after disruptive upgrade 6.2(8) to 6.2(10).

  ```
  F340.11.24-C7700-1-sdi-1# sh ip arp vrf vpc-keepalive

  Flags: * - Adjacencies learnt on non-active FHRP router
         + - Adjacencies synced via CFSoE
         # - Adjacencies Throttled for Glean
         D - Static Adjacencies attached to down interface

  IP ARP Table for context vpc-keepalive
  Total number of entries: 1
  Address         Age       MAC Address     Interface
  10.55.128.86    00:00:14  INCOMPLETE      port-channel4
  ```

  Conditions: Disruptive upgrade from 6.2(8) to 6.2(10)

  Workaround: Perform ISSU upgrade instead of disruptive upgrade.

- CSCuo31207

  Symptom: In an F1 only VDC in a Cisco Nexus 7000, the **default interface ethernet** *x/y* command does not remove all configuration from interfaces which are port-channel members.

  Conditions: This issue only occurs in F1 only VDCs configured with "limit-resource module-type f1".

  Workaround: Entering the **default interface** command removes any of the remaining configuration.

- CSCur24174

  Symptom: Allocating interfaces in a VDC fails.

  Conditions: MPLS-TE is enabled in the VDC.

  Workaround:

  1 - Ensure MPLS-TE is disabled in the VDC.

  2 - Allocate interfaces in VDC.

  3 - Enable MPLS TE in the VDC.

- CSCur97641

Symptom: Sometimes when egress policy-map is applied to a VLAN, the "packets" count will not increment in "show policy map vlan"

```
N7K-1# sh policy-map vlan 100
***SNIP***
  Class-map (qos):   test (match-any)


   Aggregate forwarded :
    0 packets  82687055970 bytes   <----bytes increment but packets don't
```

Conditions: Egress policy-map applied on vlan

Workaround: None

- CSCus09312

  Symptom: Port-channels which have

  1) PVLAN trunk secondary config

  and

  2)LACP or other control protocols running,

  could flap continuously, due to BPDU's not flowing. They don't flow because the native vlan is in CBL disabled state, instead of being in CBL Blocking state.

  Conditions: The issue is specific to M1 module since the programming model is different on F2/F3 LC's.

  There is no issue on F2 and F3 modules.

  Even if the customer uses M1 module there is NO issue, if customer is allowing native VLAN on VPC Leg.

  Below are the 3 conditions that need to be satisfied to hit this bug:

  1) PVLAN port mode should be TRUNK Secondary

  2) Native VLAN is NOT allowed on VPC Leg

  3) LC Module should be M1 module

  Workaround: Workaround is to have customer have the native vlan in allowed list for the port, by configuration.

  For a private-vlan port, the command to add trunk allowed vlan 1 would be:

  ```
   switchport private-vlan trunk allowed vlan 1
  ```

- CSCus14797

  Symptom: A Nexus 7000 running 6.2(8a) has faced an unexpected reset of one of its VDCs. After the reset, that VDC was not booting up properly and its interfaces were listed in state 'unknown'. A few minutes later, an aclmgr core file was written:

```
`show cores`
VDC  Module  Instance  Process-name     PID       Date(Year-Month-Day Time)
---  ------  --------  --------------  --------  -------------------------
3    5       1         aclmgr          22993     2014-12-08 15:27:29
```

Conditions: ACL changes were made prior to the issue.

Workaround: None

- CSCus21952

  Symptom: Layer 3 traffic does not egress out of trunk secondary port/PC/VPC.

  Conditions: Routed traffic gets dropped and does not egress out of PVLAN trunk secondary port.

  Workaround: None

- CSCus22805

  Symptom:

  1. DEVICE_TEST-2-COMPACT_FLASH_FAIL: Module 5 has failed test CompactFlash 20 times on device Compact Flash due to error the compact flash power test failed.

  2. Unable to save configuration.

  Conditions: Affects sup2/sup2e.

  Workaround: Strongly advocate opening a TAC case to help analyze the logs and recommend appropriate workaround.

- CSCus24678

  Symptom: For a specific port-channel number, an interface can error-disable when trying to add a normal (non-fex) interface to a normal (non-fex) port-channel for reason type "error". The fail will happen when the Cisco Nexus 7000 thinks that a "regular port is

  being added to a fex port-channel" as seen from the logs in 'show system internal spm event-history errors" even though the port-channel is configured as a normal one.

  Conditions: -This issue has been seen on Cisco Nexus 7000 SUP2 on 6.1(3) and on 6.1(4a), but exact conditions are unknown.

  Workaround: Two workarounds:

  1. Add a FEX HIF interface to the port-channel. Then remove the FEX HIF interface and remove the port-channel. Reconfigured the port-channel from a normal (non-FEX) interface.

  Steps:

1. Create affected port-channel number
2. Add FEX HIF interface to the port-channel (ethx/x/x)
3. Remove the FEX HIF interface
4. Remove port-channel (no interface port-channel [x])
5. Under non-FEX interface (Ethx/x) create the port-channel (channel-group mode x [active])

Should work at this point.

- CSCus25882

  Symptom: No traffic rate on interface.

  As soon as interface comes up, log occurs.

  Conditions: N/A

  Workaround: Hard reset of module would resolve the issue for few days but issue re-occurs.

- CSCus26929

  Symptom: SAP negotiation is failing, responder does not send the response for message #0 - even if validation of that packet is successful.

  Port is pending with authorization:

  ```
  cap-r10-n7k-2(config-if)# show int e10/18
  Ethernet10/18 is down (Authorization pending)
  ```

  (the same status on the other side). Cts status shows:

  ```
  cap-r10-n7k-2(config-if)# show cts interface e10/18
  CTS Information for Interface Ethernet10/18:
    CTS is enabled, mode:   CTS_MODE_MANUAL
    IFC state:              CTS_IFC_ST_SAP_NEGOTIATING_STATE
    Authentication Status:  CTS_AUTHC_SKIPPED_CONFIG
      Peer Identity:
      Peer is:              Unknown in manual mode
      802.1X role:          CTS_ROLE_UNKNOWN
      Last Re-Authentication:
    Authorization Status:   CTS_AUTHZ_SKIPPED_CONFIG
      PEER SGT:             0
      Peer SGT assignment:  Not Trusted
    SAP Status:             CTS_SAP_INCOMPLETE
      Configured pairwise ciphers:
      Replay protection:
      Replay protection mode:
      Selected cipher:
      Current receive SPI:
      Current transmit SPI:
    Propagate SGT: Disabled
  ```

Shut/no shut does not help to resolve the issue.

Conditions:

Workaround: None

- CSCus28111

  Symptom: 1. Port gets error disabled

  SYS-N7K2-AGG-LEAF# show interface ethernet 2/39

  Ethernet2/39 is down (Internal-Fail errDisable, libeventseq: sequence timeout)

  2. VLAN type change fails

  SYS-N7K2-AGG-LEAF(config-vlan)# vlan 1520

  SYS-N7K2-AGG-LEAF(config-vlan)# private-vlan isolated

  ERROR: VLAN 1520 : type change failed

  Conditions: When Private-vlan VPCs are already up and a "no vlan [vlan-range]" and then vlans are created again and changed to pvlan type

  Workaround: 1. Do a shut no-shut on the error disabled port to bring it up

  SYS-N7K2-AGG-LEAF(config)# interface ethernet 2/39

  SYS-N7K2-AGG-LEAF(config-if)#shut

  SYS-N7K2-AGG-LEAF(config-if)#no shut

  2. If VLAN type change to PVLAN type did not succeed for a particular vlan or range of vlans, configure the type change through for those vlans yet again:

  SYS-N7K2-AGG-LEAF(config-vlan)# vlan 1520

  SYS-N7K2-AGG-LEAF(config-vlan)# private-vlan isolated

  SYS-N7K2-AGG-LEAF(config-vlan)# exit

- CSCus29110

  Symptom: When we are polling for the oid ipNetToPhysicalPhysAddress for certain IPv6 neighbors, the MAC addresses are not returned. This issue is seen only for neighbors that are more than five weeks old.

  Conditions: Some neighbors (age around five weeks) are missing in the SNMP data base and the same is present in icmpv6 and ND databases.

  Workaround: Clear the IPv6 neighbor.

- CSCus31599

  Symptom: On a M108 module, ports 1 and 2 were reporting as "rate-mode dedicated" while ports 3-8 reported as "rate-mode shared".

Conditions: All conditions unknown at this time.

When the issue was hit in the customer environment, they had upgraded from M132 module to M108. It is assumed that some of the ports on the M108 module came up with "rate-mode shared" post hardware upgrade.

Workaround: Create dummy VDC.

Allocate interfaces to new/dummy VDC. (interfaces report correct rate-mode at this time).

Allocate interfaces back to originating VDC. (interfaces should still report correct rate-mode now).

- CSCus32483

  Symptom: The **switchport mode fex-fabric** command is not available under interface portchannel.

  ```
  N7k(config)# int po104
  N7k(config-if)# switchport
  N7k(config-if)# switchport mode ?
    access        Port mode access
    dot1q-tunnel  Port mode dot1q tunnel  <-- 'switchport mode fex-fabric' option missing
    trunk         Port mode trunk
  ```

  This is when creating a new FEX.

  Conditions: Dual N7K-SUP2 running 6.2.8b. Configuring new fex on N7K-M224XP-23L.

  Workaround: None

- CSCus32949

  Symptom: During boot after NX-OS downgrade 6.2(10) -> 6.2(2a), "flowcontrol receive on" and "flowcontrol send on" is not set on physical interface.

  These configurations on port-channel interface is set without problem.

  Therefore physical interface cannot join port-channel for configuration differential.

  ```
  ===================
  Wed Dec 24 22:27:26 2014:type=update:id=vsh.7083:user=admin:cmd=configure terminal ; interface Ethernet3/7 ; flowcontrol receive on (FAILURE)
  Wed Dec 24 22:27:26 2014:type=update:id=vsh.7083:user=admin:cmd=configure terminal ; interface Ethernet3/7 ; flowcontrol send on (FAILURE)
  Wed Dec 24 22:27:26 2014:type=update:id=vsh.7083:user=admin:cmd=configure terminal ; interface Ethernet3/7 ; channel-group 54 mode active (FAILURE)
  ===================
  ```

  Conditions: When downgraded.

  Workaround: Once set "priority-flow-control mode off", then can configure flowcontrol configuration on physical interface.

- CSCus33041

  Symptom: 1. HQ01OTV and HQ02OTV are otv edge devices within one side.

2. IDC01OTV is the remote otv edge.

3. All above 3 OTV edge devices are running with n7000-s2-dk9.6.2.8a, working in otv unicast mode, with otv stp-synchronization enabled.

4. HQ01OTV and HQ02OTV are vpc peers.

5. we find out that all vlans are active on both HQ01OTV and HQ02OTV,

```
HQ01OTV#  show otv vlan | ex inac
OTV Extended VLANs and Edge Device State Information (* - AED)
Legend:
(NA) - Non AED, (VD) - Vlan Disabled, (OD) - Overlay Down
(DH) - Delete Holddown, (HW) - HW: State Down
 (NFC) - Not Forward Capable

VLAN   Auth. Edge Device                   Vlan State              Overlay
----   ----------------------------------   ----------------------      ------
-
   2*  HQ01OTV                          active              Overlay17
   4*  HQ01OTV                          active              Overlay17
   5*  HQ01OTV                          active              Overlay17
   6*  HQ01OTV                          active              Overlay17

HQ02OTV# show otv vlan | ex inac
OTV Extended VLANs and Edge Device State Information (* - AED)

Legend:
(NA) - Non AED, (VD) - Vlan Disabled, (OD) - Overlay Down
(DH) - Delete Holddown, (HW) - HW: State Down
 (NFC) - Not Forward Capable

VLAN   Auth. Edge Device                   Vlan State              Overlay
----   ----------------------------------   ----------------------      ------
-
   2*  HQ02OTV                          active              Overlay17
   4*  HQ02OTV                         active              Overlay17
   5*  HQ02OTV                          active              Overlay17
   6*  HQ02OTV                          active              Overlay17
```

Conditions: otv stp-synchronization is enabled on otv edge devices which is working as vpc peers.

Workaround: 1.shut overlay ports

2. remove otv stp-synchronization

3. no shut overlay ports.

- CSCus34965

  Symptom: Memory leak on STP causing the STP process to crash.

  Conditions: N/A

  Workaround: N/A

- CSCus35514

  Symptom: Telnet hostname not working. Other protocols are working (ping/ssh).

  Conditions: N/A

  Workaround: Use host IP address instead of host name.

- CSCus36072

  Symptom: Multicast traffic is not received by the host connected on the FEX HIF.

  Conditions: Inter-VLAN multicast traffic would be affected. Traffic within same VLAN would work.

  VPC

  Receiver should be connected on FEX.

  Workaround:

- CSCus36521

  Symptom: SIS Neighbor with network type p2p not coming up between NXOS and IOS

  Conditions: With network type p2p the adjacency is not coming up.

  It shows INIT State on IOS Device and does not list the adjacency on NXOS end.

  Workaround: Remove p2p network type command from the interfaces on both ends.

- CSCus38183

  Symptom: When using the ethanalyzer to capture traffic toward supervisor with following options, the file rotation stops at around 1135th iteration and no more packets are captured.

  [command]

  ```
   # ethanalyzer local interface inband limit-captured-frames 0 capture-ring-buffer files 30 write bootflash:///aaa
  ```

  Conditions: # dir

  ```
          0    Dec 08 09:34:33 2014  ethanalyzer-aaa
     ... snip ...
     10742694    Dec 09 04:31:06 2014  ethanalyzer-aaa_01131_20141209043006
     10740950    Dec 09 04:32:06 2014  ethanalyzer-aaa_01132_20141209043106
  ```

```
10742767    Dec 09 04:33:06 2014  ethanalyzer-aaa_01133_20141209043206
10736733    Dec 09 04:34:07 2014  ethanalyzer-aaa_01134_20141209043306
      24    Dec 09 09:29:45 2014  ethanalyzer-aaa_01135_20141209043407  <== capture stops here
```

Workaround: None

- CSCus39311

  Symptom: The SNMP agent in the Nexus 7000 series always returns a value of 1 for the object dot3adAggAggregateOrIndividual, regardless of how the device is actually configured.

  Conditions: None

  Workaround: None

- CSCus40378

  Symptom: Broadcast traffic is not flooded out the last 24 FEX ports of N2K-C2248PQ-10GE device attached to N7K-F248XP-25;

  At the same time unicast traffic is sent correctly out of these interfaces;

  No issue is seen with ingress traffic on these interfaces, only egress broadcast is affected;

  The first 24 ports (1-24) do not have such issue;

  Conditions: The problem could be seen on N2K-C2248PQ-10GE device attached to N7K-F248XP-25 line card;

  Workaround: None

- CSCus41209

  Symptom: Bios verification fails for "show system verify bios flash/protection" on N7K-SUP1 running 6.2(2) :

  ```
  DDC1-USON-DSW1# sho system verify bios flash 0
  Return code : 0xffffffff

   BIOS Verification Failed
  DDC1-USON-DSW1# sho system verify bios flash 1
  Return code : 0xffffffff

   BIOS Verification Failed
  DDC1-USON-DSW1# sho system verify bios pro 0
  Return code : 0xffffffff

   BIOS Verification Failed
  DDC1-USON-DSW1# sho system verify bios pro 1
  Return code : 0xffffffff
  ```

```
 BIOS Verification Failed
!
DDC1-MDFE1-DSW1# sho system verify bios flash 0
Return code : 0xffffffff

BIOS Verification Failed
DDC1-MDFE1-DSW1# sho system verify bios flash 1
Return code : 0xffffffff

 BIOS Verification Failed
DDC1-MDFE1-DSW1# sho system verify bios protection 0
Return code : 0xffffffff

 BIOS Verification Failed
DDC1-MDFE1-DSW1# sho system verify bios protection 1
Return code : 0xffffffff
```

Conditions: NX-OS running :6.2(2) on N7K-SUP1

Workaround: None

- CSCus42535

  Symptoms: After pulling out the active supervisor from the Nexus 7000, some of the Anycast-HSRP bundles are in Active/Active state.

  This is causing the Anycast-HSRP state to be in Active/Active and their standby state to be unknown on both sides.

  Conditions: Physically pull out the active supervisor from the Nexus 7000.

  Workarounds: Reinsert the pulled supervisor.

- CSCUS44449

  Symptom: "mode fabricpath" not seen under running configuration of certain VLANs though the "show vlan" lists these VLANs as fabricpath vlans.

  No impact observed .

  Conditions: VLANs configured as FP VLANs but "mode fabricpath" missing under running config of these VLANs.

  Workaround: None

- CSCus45372

  Symptom: You will see a build up of DCOS_SSHD/VSH process pairs when using DCNM. TheSSHD process is not terminating due to VSH even with the exec-timeout configured (with and without).

Conditions: DCNM is used and that does SSH connections to the Nexus 7000.

Workaround: The **no feature sshd** command has to be entered and SSH reconfigured. This clears for a period of time.

- CSCus45517

  Symptom: BGP MED not used in path selection for paths with LOCAL AS Neighbors.

  Conditions: Device running NX-OS.

  Workaround: configure :

  ```
  bestpath always-compare-med
  ```

- CSCus47263

  Symptom: A Nexus 7000 pair in vPC will suspend all vPCs on the secondary when the operational primary is reloaded.

  Conditions: When the vPC peer-link is configured on an F3 series line card and the peer-keepalive is configured on an M series line card.

  Workaround: Configure the peer-keepalive on the F3 card or management interface. The management interface will only prevent this if running 6.2(10) otherwise use the F3 card.

- CSCus49797

  Symptom: 6PE bgp peering establishes and just after that juniper sends notification message saying that Cisco Nexus sends non-negotiated IPv6 unicast update packets , tears down the BGP session.

  Conditions: 6PE BGP session between Cisco Nexus 7000 and Juniper device.

  Workaround: None

- CSCuo90184

  Symptom: ARP packets will not processed and all ARP packets will be dropped due to block ACL due to the following ARP access-list,

  ```
  N7k-TEST(config)# arp access-list OTV-BLOCK-HSRP-ARP
  N7k-TEST(config-arp-acl)#   10 deny ip any mac 0000.0c07.ac00 ffff.ffff.ff00
  N7k-TEST(config-arp-acl)#   20 deny ip any mac 0000.0c9f.f000 ffff.ffff.f000
  N7k-TEST(config-arp-acl)#   30 permit ip any mac any
  ```

  Without calling the arp inspection filter(ip arp inspection filter OTV-BLOCK-HSRP-ARP vlan), the ARP access-list will be applied to all vlans and block ARP.

  Conditions: N/A

  Workaround: None

- CSCus45511

  Symptom: Add Debug Messages in MSDP API

This bug is raised to add debug messages in the MSDP Application on NX-OS which could help to isolate the issue further whether the MSDP not coming up or sending FIN Packets is actually caused by MSDP Application or Netstack

Conditions: Running MSDP on NX-OS

Workaround: This issue is resolved.

- CSCus51508

  Symptom: With 2 Cisco Nexus 7000s using VSAN 15 and 18, with IVR for VSAN 18, when a VF port (connected to a storage array) comes back up, no RSCN is sent to the ports in VSAN 15 on the remote switch. On the local switch, all ports are receiving a proper RSCN and on the remote switch, there is a RSCN sent for VSAN 18.

  Conditions: This problem is present in 6.2(8) and 6.2(8b), but not in 6.2(10).

  Workaround: Upgrade to 6.2(10).

- CSCus54213

  Symptom: Service not responding followed by crash.

  Conditions: applying L4 Protocol ACL to interface with stale entry.

  Workaround: None

- CSCus54797

  Symptom: Backup adjacency server flap links with all neighbors.

  Conditions: Primary adjacency server is at site B and backup adjacency server is at site A.

  If primary adjacency server goes down and backup adjacency server takes over, it will flap links with all neighbors after 10 minutes.

  Workaround: None

- CSCus57051

  Symptom: Failure recovery action::

  "Standby will be rebooted to force netboot and image download".

  Install has failed. Return code 0x4093005E (system switchover failed).

  Please identify the cause of the failure, and try 'install all' again.

  Conditions: ISSU from 6.2(8a) to 6.2.10.

  Workaround: Upgrade without ISSU.

- CSCus57079

  Symptom: Link up/down events might show up on N7K-F248XP-25E ports with no cable/fibre, with/without dust plug.

  Conditions: N/A

Workaround: Keep interface admin shut.

- CSCus57608

  Symptom: When route leaking between VRFs, the destination VRF where the route is being imported fails to remove the old/stale route.

  Conditions: The symptom is observed when route leaking between VRFs on the Cisco Nexus 7000. The route is removed from the parent VRF, but not removed from the child VRF.

  Workaround: Clear IP route in the destination VRF.

- CSCus61120

  Symptom: When upgrading NX-OS policy-map on GRE interfaces are no longer enforced.

  Conditions: The symptom is observed when upgrading from previous NX-OS releases to 6.2(8a), 6.2(8b), and 6.2(10).

  Workaround: Remove the affected tunnel interface and reapply along with the QoS configuration (service-policy) after the upgrade.

- CSCus61813

  Symptom: Loop in the network after upgrade 6.1(4)->6.2(8a).

  Conditions: Upgrade 6.1(4)->6.2(8a).

  Workaround: Break physical loop, as spanning-tree cannot do that.

- CSCus61885

  Symptom: (configuration)

  ```
  ipv6 route 2001:2d8:d1:201::/64 Ethernet3/13 2001:2d8:d8:1::2 name IPv6_UNION
  ipv6 route 2001:2d8:d1:201::/64 2001:2d8:d8:3::e name IPv6_UNION(B) 100


  (RIB)
  KRSS02-PUBS01-7009L3(config)# sh ipv6 route 2001:2d8:d1:201::
  IPv6 Routing Table for VRF "default"
  '*' denotes best ucast next-hop
  '**' denotes best mcast next-hop
  '[x/y]' denotes [preference/metric]


  2001:2d8:d1:201::/64, ubest/mbest: 1/0
      *via 2001:2d8:d8:1::2, [1/0], 04:40:07, static    >>>>> eth3/13 is not
  installled.
       via 2001:2d8:d8:3::e, Po4, [100/0], 05:21:33, static
  ```

  Conditions: IPv6 static route with interface and IPv6 BFD

  Workaround: 1. Remove static and add the static route again,

2. Do not use IPv6 BFD.

- CSCus61975

  Symptom: SUP2/2E will down BGP session with "holdtimer expired error" by configure as path list.

  Conditions: - Switch is learning large BGP entry.

  - Timer is very aggressive.

  - This issue is occurred when configure as path list as a wild card (".*")

  Workaround: Use more defensive BGP timer.

- CSCus64084

  Symptom: DFA-Spine1# show int trunk

```
--------------------------------------------------------------------------------
Port          Vlans Forwarding on FabricPath
--------------------------------------------------------------------------------
Eth3/1         10-20
Eth3/2         none          ====> Port 3/2 show missing FP Vlan
Po100          10-20
!
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
100   Po100(SU)   Eth      LACP      Eth3/1(P)    Eth3/2(P)
!
interface port-channel100
  switchport
  switchport mode fabricpath
```

  Conditions: - FP core port in a port-channel with multiple interfaces.

  - PO interface is in switchport mode fabricpath

  Workaround: This is a cosmetic issue.

  Shut/no shut on physical port might show VLAN forwarding but shut/no shut PO might show missing VLAN again.

- CSCur31394

  Symptom: aclmgr crash is seen when applying/removing large ACL configuration to SVI, for example:

```
interface vlan 1-800
```

no ip access-list X

will see something like:

SYSMGR-2-SERVICE_CRASHED: Service "aclmgr" (PID 6058) hasn't caught signal 6 (core will be saved)

Conditions: This is not traffic impacting.

Workaround: Apply ACL configuration in smaller chunks. For example:

```
interface vlan 1-100
no ip access-list X
```

- CSCus71454

  Symptom: In a private VLAN VPC setup in private-vlan host mode, when peer link flaps, the vPC leg in private-vlan host mode also flaps and comes back up in some time. There is traffic loss from the vPC leg until the leg bringup happens again.

  Conditions: The vPC legs have to be private-vlan host mode as follows:

  ```
  switchport mode private-vlan host
  ```

  Example configuration:

  interface port-channel10

  switchport

  switchport mode private-vlan host

  switchport private-vlan host-association 2 3

  vpc 1

  Workaround: None

- CSCus69869

  Symptom: Traffic black-hole for receivers in secondary VLAN on non-forwarder peer when source is in core.

  Conditions: On forwarder peer vPC leg is down and there are no other ports on secondary VLAN on this box.

  Workaround: Bring the vPC leg up on the forwarding peer or make sure source has better metric to non-forwarding peer.

- CSCus52559

  Symptom: Whenever multicast traffic is flooded with MD flood LTL as the DI and if the chassis has any M modules that are powered up, the MD flood received in the M module is flooded back to the fabric.

  This packet on reaching the F2e/F3 cards triggers egress MAC learn which sometimes could overwrite an existing MAC entry with a wrong destination port.

  Conditions: Conditions in which this issue could happen are:

1. Any M-series module should be up in the chassis

AND

2. MD flood LTL will be used always if Private VLAN SVI part of OIF list for the mutlicast group

OR

3. Until an optimised multicast MD LTL is obtained for a mutlicast group, MD flood LTL will be used for a very brief period.

Workaround: None

- CSCus76724

  Symptom: On M1XL line cards, when some VLAN configuration causes a private-VLAN association to be non-operational , private-VLAN trunk secondary port sees traffic loss. Similarly, when the trunk association is unconfigured and re-configured on private-VLAN trunk-secondary port, the issue might be observed.

  Conditions: This issue is seen on M1XL linecards. Will not be seen with M1 and F-series line cards

  Example config and trigger:

  Config:

  ```
  switch(config-if)# show running-config interface e3/3

  !Command: show running-config interface Ethernet3/3
  !Time: Wed Feb  4 00:38:51 2015

  version 6.2(12)

  interface Ethernet3/3
    switchport
    switchport mode private-vlan trunk secondary
    switchport private-vlan association trunk 2 3
    no shutdown
  ```

  The issue will be seen after any of the following triggers

  1. Delete and recreate of primary vlan

  ```
  switch(config-if)# no vlan 2
  switch(config)# vlan 2
  ```

```
switch(config-vlan)# private-vlan primary
switch(config-vlan)# private-vlan association 3
switch(config-vlan)# ex
```

2. Delete and recreate secondary vlan

```
switch(config-if)# no vlan 3
switch(config)# vlan 3
switch(config-vlan)# private-vlan isolated
switch(config-vlan)# ex
```

3. Delete and re-add trunk association on the port

```
switch(config)# int e3/3
switch(config-if)# no switchport private-vlan association trunk 2 3
switch(config-if)#  switchport private-vlan association trunk 2 3
```

Workaround: Workaround is to do a shut no-shut on the port or PC or VPC leg where the issue is observed

```
switch(config)# int e3/3
switch(config-if)# shutdown
switch(config-if)#  no shutdown
```

- CSCus56036

  Symptom: BGP session would continuously flap with FD read error syslogs. fd Read error is the signature of this defect with varying error no :-

  Switch %BGP-5-ADJCHANGE: bgp-65111 [5639] (default) neighbor 10.x.x.x Down - fd read error

  or

  bgp 65111 [5639]: [5717]: (default) EVT: Read error from peer 10.x.x.x: No buffer space available

  BGP tracebacks can be seen as well but not necessary

  Conditions: Reload of chassis

  or

  SSO

  or

  Module Reload

  or

  Interface flap

  Workaround: Restarting BGP process should bring back BGP sessions.

  Restart BGP < AS >

This is a corner case timing issue and is seen rarely. However, there is no way to avoid this issue to be seen.

- CSCug81339

  Symptom: Security Daemon crashes with heartbeat failure with the following output in logs

  %SYSMGR-3-HEARTBEAT_FAILURE: Service "Security Daemon" sent SIGABRT for not setting heartbeat for last 6 periods. Last heartbeat 3

  88.97 secs ago.

  %SYSMGR-2-SERVICE_CRASHED: Service "Security Daemon" (PID 6144) hasn't caught signal 6 (core will be saved).?

Note This bug applies to Nexus 1000V switch also for 5.2(1)SM1, and 4.2(1)SV2.

  Conditions: No specific conditions, can occur randomly.

  Workaround: None

- CSCur99790

  Symptom: NX-OS "show run vrf DE" will show the running configuration for vrf default.

  Conditions: vrf name is DE.

  Workaround: complete show run.

- CSCus62502

  Symptom: If OTV Tunnel Depolarization is implemented, traffic is dropped when several OTV tunnels are down.

  Conditions: None

  Workaround: None

- CSCus64703

  Symptom: N7018

  6.1.4aE1

  --When trying to change ospf costs on few port channels from 33568 to 800. Ospf cost did not take affect on PO 47

  --On Po 47, we see that running configuration has right cost, but ospf does not have it (has old cost)

  --Ospf process crashed when the cost change was happening

  --config change with ospf/bgp/lacp.. etc, issue is seen.

  Conditions: N/A

Workaround: Reconfigure the ospf costs and it will work.

- CSCus64899

  Symptom: BFD stops working.

  Conditions: ++ when BFD echo is enabled

  ++ one side is f2

  ++ other side is F3

  Workaround: disable BFD echo.

- CSCus65380

  Symptom: Customer is using n7000-s1-dk9.6.2.8a.bin, once traceroute IP with source-ip specified, then will see alarms,

  ```
  n7k-2# traceroute 1.1.1.1 source 123.1.1.123
  traceroute to 1.1.1.1 (1.1.1.1) from 123.1.1.123 (123.1.1.123), 30 hops max, 40 byte packets
   1  123.1.1.123 (123.1.1.123)  0.41 ms !N  0.176 ms !N  0.144 ms !N
  2015 Jan 27 05:23:05 n7k-2 dcos-traceroute[4674]: IP-3-IP_FAILURE: Failed to Failed call urib_api_init()
  ```

  Conditions: Traceroute and specify the source IP address.

  Workaround: None

- CSCus66234

  Symptom: BGP process might crash following constant IGP neighbor adjacency flaps.

  Conditions: Unstable IGP/LDP session.

  Workaround: None

- CSCus66235

  Symptom: Match statements in a single sequence of route-map are supposed to behave like AND. It is not happening when 'match ip next-hop' or 'match interface' is present in the sequence. The behaviour is fine with any other combination of match statements.

  Conditions: In a route-map, when 'match ip next-hop' or 'match interface' statement is used in combination with other match statements. This is not behaving as AND. The final result is TRUE if 'match ip next-hop' or 'match interface' turns out to be true and even if the other match statements are evaluated to FALSE.

  ```
  route-map OSPF-FILTER1 permit 10
    match interface Ethernet5/39.321
    match metric 20    <-- final result is TRUE even if this match statement is FALSE.
  ```

  Workaround: Add another sequence in a 'deny' statements for the other match statements.

```
route-map OSPF-FILTER1 permit 10
  match interface Ethernet5/39.321
  match metric 20
route-map OSPF-FILTER1 deny 20
  match metric 10   < we can add unwanted attribute values for deny here.
```

- CSCus68473

  Symptom: Cisco Nexus 7000 running 6.2(8E10) crash on the urib process after clearing all routes in the VRF rib.

  Conditions: Clearing all routes in the RIB when there is high route count.

  Workaround: Clear individual routes, instead of the entire table using *.

- CSCus71342

  Symptom: MAC address not learned in VLAN with SVI from ARP bcast only in F3 cards. Impact of issue is unexpected flooding or longer time to learn mac as it will need bidirectional traffic.

  Conditions: Happen with conversational learning enabled FP core port received ARP bcast in vlan with SVI

  Workaround: None

- CSCus72364

  Symptom: Cisco Nexus 7000 BFD brings down additional BFD peers.

  After manually shutting down of some other BFD link, other sessions also get flapped with below message in syslog.

  ```
  2014 Oct  7 14:33:12 BGL.G.03-N7K-1-DIST %BFD-5-SESSION_STATE_DOWN: BFD session 1140850705 to neighbor 10.233.255.102 on interface Eth3/12.2 has gone down. Reason: Neighbor Signaled Session Down.
  ```

  Conditions: BFD optimize subinterface.

  Code- 6.2(8a) / also seen in 6.2.10

  Mod 9 - N7K-M224XP-23L

  Mod 3 - N7K-F248XP-25E

  Workaround: None

- CSCus73066

  Symptom: An M2 line card might be reset by the supervisor due to the EOBC heartbeat being missed by the line card.

  Conditions: Unknown.

  Workaround: None

- CSCus77610

Symptom: Link might go to errdisable state with "UDLD empty echo" very rarely when line card reload.

Conditions: On 10G board, configure:

1. UDLD protocol enabled

2. Option "system default link-fail laser-on" enabled

3. interface debounce time is set to 0

then reload the line card.

Workaround: 1. Enter shut/no shut the port that in "errdisable" state, or

2. Configure the link debounce time to 10ms or larger, or

3. Disable the UDLD protocol, or

4. Configure "no system default link-down laser-on" option

- CSCus78697

  Symptom: Logging source-interface seems to be non-working with v6 syslog server on Cisco Nexus 7000 after device reload even the loggingsource-interface pointing to the loopback0 interface.

  Conditions: After device reload.

  Workaround: Reapply logging source-interface loopback0.

- CSCus81533

  Symptom: LACP PDUs not received by the CPU. As a results, port-channel members are suspended.

  ```
  Nexus# show system internal pktmgr internal event-history lcache-err
          No Memory available for pcm entry ifindex <hexa ID>
  ```

  Conditions: This issue is seen in Nexus7000 with 6.0 releases.

  No other known conditions.

  Workaround: None

- CSC us82364

  Symptom: The **show ip bgp regexp** command could cause the issue that bgp neighbor goes down

  ```
  Switch# sh clock ; sh ip bgp regexp "_12345_" ; sh clock
  17:19:51.055 JST Fri Feb 06 2015

  BGP routing table information for VRF default, address family IPv4 Unicast
  BGP table version is 539791, local router ID is 202.213.196.212
  Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
  Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
  ```

```
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

   Network            Next Hop            Metric     LocPrf     Weight Path
17:20:01.362 JST Fri Feb 06 2015
Switch# 2015 Feb  6 17:20:01 Switch %BGP-5-ADJCHANGE:  bgp-2527 [7785] (default) neighbor
202.213.197.101 Down - holdtimer expired error
2015 Feb  6 17:20:14 Switch %BGP-5-ADJCHANGE:  bgp-2527 [7785] (default) neighbor
202.213.197.101 Up
```

Conditions: Cisco Nexus 7000 with SUP2/2E learns ~500K routes and attributes from BGP.

Aggressive BGP keep-alive timer is set, such as hello 3sec and hold time 10sec.

Workaround: Default keep-alive timer, hello 60sec and hold 180sec, does not cause the issue.

- CSCus82982

  Symptom: When using the **mpls ldp autoconfig** command in ISIS, modifying the 'is-type' can cause an interface to de-register as an MPLS interface even with ISIS enabled on that interface and even though the 'is-type' is permitted.

  Conditions: Modifying the 'is-type' in the ISIS configuration causes an interface to no longer show as a MPLS-enabled interface.

  BEFORE:

  ```
  Nexus# show mpls interfaces detail
  Interface Ethernet8/25:
      ldp enabled
      MPLS operational
      Label space id 0x10000001
      MPLS sub-layer Ethernet8/25-mpls layer(0x26000001)
  ```

  AFTER:

  ```
  Nexus# show mpls interfaces detail
  Nexus#
  ```

  Workaround: Occasionally, modifying the 'mpls ldp autoconfig' configuration or 'is-type' again will cause it to re-register. In other instances, 'mpls ip' must be explicitly configured on the interface.

- CSCus83776

  Symptom: The ITD can't advertise the route of VIP.

  Conditions: ' advertise enable' does not work.

  Workaround: Do not enable ' advertise enable' and configure static route.

Resolved Caveats—Cisco NX-OS Release 6.2(12)

- CSCuo82450

  Symptom: An egress FCoE interface log output discards during congestion even though pause frames are sent upstream on the ingress interface. Pause frames received on the egress interface do not prevent the output discards.

  Affected ingress interfaces can be identified when the 'ENABLED' field is 1 in the output of the following module-level command:

  ```
  <b>show hardware internal qengine inst <i>inst-num</i> table vq_voq_td</b>
  ```

  where <i>inst-num</i> = quotient of (the port number - 1) / 4. For example, to verify Ethernet1/1 is affected using 'slot 1' and 'inst 0' as arguments to the above command:

  ```
  <pre>
   switch# slot 1 show hardware internal qengine inst 0 table vq_voq_td | include
  "port|ENABLE|^0"
   | Inst 0; port(s) 1-4
   INDEX  QUEUE  PKT TYPE  VL  THRESHOLD  ENABLE
   0      8q     CRD DE    0   14400      1
  </pre>
  ```

  Conditions: This issue only applies to interfaces with a 'no drop' CoS, that is, FCoE interfaces. An interface will be affected by this issue only after a supervisor switchover (this includes ISSU/ISSD switchovers) and then the interface flaps for any reason (this includes moving the interface into a port channel).

  For Nexus 7000/7700 switches, the first affected release is NX-OS release 6.2(2).

  For MDS 9500/9700 switches, the first affected release is NX-OS release 6.2(7).

  Workaround: To nondisruptively restore the 'no drop' functionality, set the priority flow control to 'on' and back to 'auto' for each affected ingress interface. If the interface is a member of a port channel then the change should be done at the port channel interface level. For example:

  ```
   switch(config-if)# <b>interface port-channel 1</b>
   switch(config-if)# <b>priority-flow-control mode on</b>
   switch(config-if)# <b>priority-flow-control mode auto</b>
  ```

  The above workaround can only be applied to interfaces which are up. This will restore the potency of pause frames on the Ethernet interfaces. However, further port flaps will cause the issue to recur on the interface.

- CSCus68892

  Symptom: On January 27, 2015, a buffer overflow vulnerability in the GNU C library (glibc) was publicly announced. This vulnerability is related to the various gethostbyname functions included in glibc and affect applications that call these functions. This vulnerability may allow an attacker to obtain sensitive information from an exploited system or, in some instances, perform remote code execution with the privileges of the application being exploited. This vulnerability is documented in CVE-2015-0235.

A Cisco Security Advisory has been published to document this vulnerability at:

http://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20150128-ghost

This bug has been opened to address the potential impact on this product.

Conditions: Exposure is not configuration dependent.

Workaround: This issue is resolved.

- CSCus55175

  Symptom: The VDC bind for the breakout ports is not failed upon execution of the breakout CLI.

  Conditions: The ACLQoS returning failure for the breakout TLV verify message

  The ACLQoS not returning response within the timeout period.

  Workaround: This issue is resolved.

- CSCur26436

  Symptom: Nexus 7000 and MDS 9000 switches include a version of SSL that is affected by the vulnerability identified by the Common Vulnerability and Exposures (CVE) IDs:

  CVE-2014-3566

  Conditions: A POODLE exploit requires a man in the middle attack between the switch (the LDAP client utilizing the SSL client) and the LDAP server. Nexus 7000 and MDS 9000 both contain an SSL client with SSLv3 support. The client supports fall back to SSLv3 if negotiation with TLS 1.0 fails.

  The LDAP feature may be configured to utilize this client. This feature is disabled by default. Hence, this vulnerability only exists if the LDAP feature is enabled.

  Workaround: This issue is resolved.

- CSCur30049

  Symptom: BGP sets next-hop-self for self-originated routes. not RFC4271 compliant, we need RPM to add bgp-redist-unchanged for setting ip next-hop.

  Conditions: While redistributing routes, the next-hop is being changed. This option is needed to keep the next-hop unchanged.

  Workaround: This issue is resolved.

- CSCur64055

  Symptom: When lookup mode is changed from IP to MAC.

  Conditions: ISIS crashes and restarts.

  Workaround: This issue is resolved.

- CSCus26870

  Symptom: The following Cisco products:

NEXUS 7000

NEXUS 6000

NEXUS 5000

MDS

include a version of NTPd that is affected by the vulnerabilities identified by the Common Vulnerability and Exposures (CVE) IDs:

CVE-2014-9293

CVE-2014-9294

CVE-2014-9295

CVE-2014-9296

This bug has been opened to address the potential impact on this product.

Conditions: This issue is configuration dependent and applies only when the following command is configured:

feature ntp

It affects all versions of NX-OS that support NTP prior to the first fixed release of code for each product.

Workaround: This issue is resolved.

- CSCus54220

  Symptom: This was found when verifying CSCur31394. Service not responding followed by router crash occurs when applying ACLs to many SVIs at the same time. (1000 SVIs and 140 port-channels).

  Conditions: This happened when ACL is applied to large number of SVIs.

  Workaround: This issue is resolved.

- CSCus56042

  Symptom: The Cisco Nexus 7000 unexpectedly wrote out a TACACS core file when doing TACACS accounting.

  Conditions: TACACS accounting is configured.

  Workaround: This issue is resolved.

- CSCus51304

  Symptom: With an underlying MPLS (IGP over MPLS) , EBGP routes are marked inconsistent without any route inconsistent by the consistency checker.

If MPLS is removed, there is no inconsistency.

FIB seems to be consistent with RIB.

Traffic forwarding seems to have no impact.

Conditions: Underlying MPLS

Workaround: This issue is resolved.

- CSCus48046

  Symptom: OID is not increasing and the getnext returns the same 1st entry of the SFP indices.

```
2    0.006410    10.210.126.13     10.199.15.9 SNMP  115   get-next-request
1.3.6.1.4.1.9.9.91.1.1.1.1.3.300000002 1.3.6.1.4.1.9.9.91.1.1.1.1.4.300000002
3    0.009069    10.199.15.910.210.126.13     SNMP  117   get-response
1.3.6.1.4.1.9.9.91.1.1.1.1.3.300000002 1.3.6.1.4.1.9.9.91.1.1.1.1.4.300000002
4    0.014416    10.210.126.13     10.199.15.9 SNMP  115   get-next-request
1.3.6.1.4.1.9.9.91.1.1.1.1.3.300000002 1.3.6.1.4.1.9.9.91.1.1.1.1.4.300000002
5    0.017114    10.199.15.910.210.126.13     SNMP  117   get-response
1.3.6.1.4.1.9.9.91.1.1.1.1.3.300000002 1.3.6.1.4.1.9.9.91.1.1.1.1.4.300000002
```

  Conditions: The problem occurs when the following conditions are met:

  1. Inserted a transceiver to e1/1

  2. MIB walk CISCO-ENTITY-SENSOR-MIB against the device

  3. Remove the transceiver from e1/1

  Workaround: This issue is resolved.

- CSCus50392

  Symptom: Multicast Packets may be dropped when traversing an F3 module.

  This will be an "all or none" symptom for the suspect flows vs an intermittent packet drop scenario.

  Conditions: All of the following conditions are required to see this issue:

  a) Having enabled "feature otv" at any point, even if "no feature otv" has been applied after, or even if no other OTV config exists

  b) Suspect multicast traffic is traversing F3 module

  c) Running versions 6.2(10) or earlier, or having ISSU'd from 6.2(10) or earlier to a later version w/o reloading the linecard (or chassis if preferred)

  Workaround: This issue is resolved.

- CSCup90186

  Symptom: When queuing policy on port-channel and port moving to it are same, policy on physical port resets to default.

Conditions: Only when queuing policy on the port-channel and port moving to it are same.

Workaround: This issue is resolved.

- CSCup43628

  Symptom: ISSU from 6.2(8) to a newer version fails with a timeout in LISP.

  Conditions: LISP feature must be enabled.

  Workaround: This issue is resolved.

- CSCui36126

  Symptom: Same subnet primary & secondary on interface.

  Advertising the same in EIGRP. We see the connected route in EIGRP topology. However, if I remove the secondary IP address from the interface, it breaks EIGRP for primary IP too.

  EIGRP topology will not show the connected prefix even though it is configured on that interface.

  Conditions: Same subnet primary & secondary on interface.

  Workaround: Shut/no shut on the interface fixes the issue.

- CSCui51673

  Symptom: After removing EIGRP from an interface the network is no longer injected into the EIGRP topology table.

  Conditions: EIGRP must have been configured on the interface.

  "redistribute direct" must be configured under the same EIGRP process

  the route-map attached to "redistribute direct" must permit the interface network

  Workaround: This issue is resolved.

- CSCuj29140

  Symptom: While applying acl with L4 protocol filtering, aclmgr crashes frequently causing swover and vdc reload

  Conditions: applying L4 Protocol ACL to interface

  Workaround: This issue is resolved.

- CSCuj70143

  Symptom: The following error can be seen when applying ACL changes:

  "Failed to complete Verification: Link exists" or

  "Database Internal error"

  Conditions: Applying a large amount of config change.

  Workaround: This issue is resolved.

- CSCuj96361

  Symptom: BGP/TCP dependant applications may NOT work as expected after Netstack crash/core is seen on NX7k running 6.2.8a and earlier releases.

  Conditions: Issue is hit, when the netstack service crashes with the log.

  %SYSMGR-2-SERVICE_CRASHED: Service "netstack"

  Workaround: This issue is resolved.

- CSCul38702

  Symptom: Configuring <b>max-metric on-startup</b> under OSPFis applied to the config, but does not cause OSPF to restore normal metrics.

  Conditions: OSPF has unexpectedly restarted since the last reboot.

  OSPF is currently configured for max-metric.

  The above conditions are not required to happen at the same time.

  Workaround: This issue is resolved.

- CSCun21805

  Symptom: The following messages are seen:

  2014 Jul 3 16:09:30 N7K %IGMP-3-PIM_PIM_LIB_API_INIT: Couldn't initialize PIM-SHM-LOCK API

  2014 Jul 3 16:09:40 N7K %IGMP-4-SYSLOG_SL_MSG_WARNING: PIM-3-PIM_LIB_API_INIT: message repeated 1 times in last 13 sec

  Conditions: - PIM feature not enabled

  - IGMP configured on an interface

  Workaround: This issue is resolved.

- CSCun86910

  Symptom: SNMP walk on N7k returns incomplete information after a change in EIGRP neighborship.

  Conditions: The CLI however shows the complete list of EIGRP neighbors.

  Workaround: This issue is resolved.

- CSCuo24886

  Symptom: Pings to Loopback learnt over MPLS is not reachable

  Conditions: a) Route is reachable via MPLS

  b) While upgrading by ISSU, the MPLS enabled interface is in down state

  c) upgrade from 6.2(2) to 6.2(2a) Via ISSU

Issue was seen on both SUP1 & SUP2

Workaround: This issue is resolved.

- CSCuo44890

  Symptom: OTV routes are not exchanged between a Nexus 7000 switch running 6.2(x) image and an ASR1K device. The overlay adjacency comes up fine, however no OTV routes are exchanged.

  Conditions: Nexus 7000 switch running 6.2(x) image

  Workaround: The fix has been incorporated in 6.2(12) by using a specific command. Here is the procedure to make the config change:

  ```
  N77-1(config)# inter overlay 1
  N77-1(config-if-overlay)# shut
  N77-1(config-if-overlay)# otv-isis default
  N77-1(config-router)# interop-enable
  N77-1(config-router)# end
  N77-1(config)# inter overlay 1
  N77-1(config-if-overlay)# no shut
  N77-1(config-if-overlay)# end
  ```

- CSCup39948

  Symptom: After reload with maximum routes X warning-only devices stops installing routes into routing table saying it has reached the limit configured.

  Even after the command stays the same in running config. after reload device ignores warning-only and set the number as actual limit for the routing table.

  And you will see following msg says configured limit reached even though the running config will say warning only.

  ```
  URIB-4-ROUTELIMIT_EXCEEDED  urib [5997] (test-base) Number of routes (4) reached or exceeds configured limit (4); dropped (1)
  ```

  ```
  vrf context test
    address-family ipv4 unicast
      maximum routes 4 warning-only >>>>>>>>>>> warning only
  ```

  after reload

  ```
  show routing vrf test internal
  ```

```
SUP state: ACTIVE
Clean Restart: YES
Ksink: PSS URI: volatile:/dev/shm/urib_runtime_pss
Servers   ; CLIS: Up  ; L3VM: Up  ; RPM: Up  ; LISP: Down; UFDM: Up  ;
Registered; CLIS: Yes ; L3VM: Yes ; RPM: Yes ; LISP: No  ; UFDM: Yes ;
SHM Min: 96; Size: 96
MPLS VPN MIB Trap Threshold: disabled by default
debug-filter vrf: default (0x1)
          client: all (index 0x0) (mask 0x1)
          prefix-list: none

VRF "test" (0x00000003) State: UP (Up) UFDM Enabled: Yes
   TIB converged: Yes
   Route limit 4, warning limit 100% (4)>>>>>>>>>>>>>> set the number as an actual limit
   Reinstall at 0% (4), pending: Yes
   Last MIB trap 00:05:49 ago; Max reset: No; Warn reset: Yes
  Current count 4, added 4, deleted 0, dropped 5
```

Conditions: When maximum routes X command configured for a routing table with warning-only option.

Workaround: This issue is resolved.

- CSCup44514

  Symptom: The MACSEC port of the N7K-M148GS-11 module with 4500 port shows port flapping or packet loss (for instance EIGRP flapping). It happens intermittently (once per day, few times per day, once in few weeks) and the port comes up immediately afterwards.

  Conditions: This issue might be seen if you have 4500 1G interop with an N7K-M148GS-11 module used with MACSEC configuration

  Workaround: This issue is resolved.

- CSCup64117

  Symptom: When a new fan or power supply unit (PSU) is inserted into the Cisco Nexus 7000 chassis, the SystemMgmtBus diagnostic test is not restarted properly.

  Conditions: When power supply or fan are hot swapped in the Cisco Nexus 7000 chassis.

  Workaround: This issue is resolved.

- CSCup65293

  Symptom: AR-CAR1-VDC1# sh ipv6 prefix-list v6-att-cbb-in

  ipv6 prefix-list SB_l3fiber_v6: 2 entries

```
  seq 5 permit cafe:1890:e000:4510::/64
  seq 10 permit cafe:1890:e000:4512::/64
ipv6 prefix-list v6-att-cbb-in: 10 entries
  seq 10 permit 0::/0
  seq 20 permit ::1/128
  seq 30 permit 0::/128
  seq 40 permit 0::/96
  seq 60 permit 0::/8 le 64
  seq 70 permit fe80::/10 le 64
  seq 80 permit fec0::/10 le 64
  seq 90 permit fc00::/7 le 64
  seq 100 permit ff00::/8 le 64
  seq 110 permit 2001:db8::/32 le 64
ipv6 prefix-list vlan_70_v6: 1 entries
  seq 5 permit cafe:1890:e000:4200::/64
ipv6 prefix-list vlan_71_v6: 1 entries
  seq 5 permit cafe:1890:e000:4201::/64
ipv6 prefix-list vlan_72_v6: 1 entries
  seq 5 permit cafe:1890:e000:4202::/64
ipv6 prefix-list vlan_73_v6: 1 entries
  seq 5 permit cafe:1890:e000:4203::/64
ipv6 prefix-list vlan_80_v6: 1 entries
  seq 5 permit cafe:1890:e000:4500::/64
ipv6 prefix-list vlan_81_v6: 1 entries
  seq 5 permit cafe:1890:e000:4501::/64
ipv6 prefix-list vlan_82_v6: 1 entries
  seq 5 permit cafe:1890:e000:4502::/64
ipv6 prefix-list vlan_83_v6: 1 entries
  seq 5 permit cafe:1890:e000:4503::/64
ipv6 prefix-list vlan_84_v6: 1 entries
  seq 5 permit cafe:1890:e000:4504::/64
ipv6 prefix-list vlan_85_v6: 1 entries
  seq 5 permit cafe:1890:e000:4505::/64
```

Conditions: every time

Workaround: This issue is resolved.

- CSCup67884

  Symptom: Service "aclmgr" crashed

Conditions: Seen during the upgrade from 4.2.8 to 5.2.9.

Customer got three of these instances

a) After failover to sup in slot 6

B) Reload of standby sup in slot 6

c) Finally during the EPLD upgrade, again when slot 6 tried to come back online & it failed to come fully on-line and was stuck in power-on state until a manual reload was done.

Workaround: This issue is resolved.

- CSCuq31499

  Symptom: FEX attaching to Cisco Nexus 7000 might crash:

  1) At 659678 usecs after Thu Jul 31 09:03:30 2014

  Reset Reason: Reset triggered due to HA policy of Reset (16)

  Service (Additional Info): satctrl hap reset

  Image Version: 6.2(8)

  Conditions: there is no trigger can be found for now.

  Workaround: This issue is resolved.

- CSCuq37760

  Symptom: On VPC+ peer F2E/M2 proxy routing, if HSRP active switch is reloaded, the HSRP standby will remove the remote bit from the HSRP VMAC, correct VMAC will only be installed when it will become HSRP active.

  Conditions: On VPC+ peer F2E/M2E proxy routing, HSRP active switch is reloaded

  Workaround: This issue is resolved.

- CSCuq40179

  Symptom: Cisco Nexus 7000 on 6.2.8a.E code might experience a slow memory leak. The problem was due to NX-OS incorrectly writing to a debug file without a size limitation.

  Conditions: N/A

  Workaround: This issue is resolved.

- CSCuq42845

  Symptom: ERROR: Below error messages seen on applying queuing policy.

  switch-outoforder(config-if)# service-policy type queuing input FAB_all-4q-7e-in --> when applying the policy directly

  ERROR: Unable to perform the action due to incompatibility: Module 7 returned status "Error reading from internal database"

```
switch-outoforder(config-if)#   service-policy type queuing output FAB_all-4q-7e-out
```

ERROR: Unable to perform the action due to incompatibility: Module 7 returned status "Error reading from internal database"

Error message is seen when a heirarchical queuing policy (in which the order of the inner service policy and queue-limit is changed) is applied. For eg. the error is seen with the below policy.

```
 policy-map type queuing nonwork-4q-7e-in
   class type queuing c-4q-7e-drop-in
     queue-limit percent 70
     service-policy type queuing dup4q-7e-drop-in
   class type queuing c-4q-7e-ndrop-in
     queue-limit percent 30
     service-policy type queuing dup4q-7e-ndrop-in
```

In the above policy, the order of the queue-limit and the service policy is changed.

Conditions: Custom Policy applied in incorrect order.

Workaround: This issue is resolved.

- CSCuq60138

  Symptom: FEX HIF with "medium p2p" configuration cannot be cleared and CLI cannot enter config-int mode for HIF .

  Conditions: Occurs for a FEX HIF with "medium p2p".

  Workaround: This issue is resolved.

- CSCuq96869

  Symptom: Some routes may either be missing or mis-programmed.

  Conditions: At VDC reload some early new route updates might be processed and dropped before processing the vdc-create. this causes ipfib to ignore the messages since they will be treated as invalid. Issue is seen during line card reload and during issu as well.This is a corner case timing issue and very rare

  Workaround: This issue is resolved.

- CSCur03040

  Symptom: Nexus 7000 may experience a process crash with OTV configuration.

```
Reason: Reset triggered due to HA policy of Reset
System version: 6.2(8a)
Service: Service "__inst_001__isis_otv" in vdc 2 hap reset
```

Conditions: OTV must be configured. This issue was first seen on 6.2(8).

Workaround: This issue is resolved.

- CSCur07245

  Symptom: When configuring IPv6 NTP on NXOS 6.2.8a using a 7010 chassis with a Sup-1 module, a core file is being generated.

  Conditions: When configuring IPv6 NTP on NXOS 6.2.8a using a 7010 chassis with a Sup-1 module, a core file is being generated.

  Workaround: This issue is resolved.

- CSCur07262

  Symptom: Port-channels using LACP fast-rate will not pass traffic.

  Conditions: Occurs when completing a supervisor switchover, the “show lacp interface ethernet x/x” command displays that the port is invalid for LACP and that the channel-group is invalid with a value of 0. The port counters for LACP PDUs do not move in the TX or RX direction, hence the device connected to the switch shows its member ports in a suspend state.

  Workaround: This issue is resolved.

- CSCur07911

  Symptom: OTV fails to send MAC updates causing connectivity issues.

  Conditions: Occurs when OTV fails to send MAC updates to peers in unicast mode from local site to remote.

  Workaround: This issue is resolved.

- CSCur12074

  Symptom: Current CLI session hangs, to get back CLI, either need to power cycle or need to reload the router by telnet to the mgmt IP.

  Conditions: This happens only with "system admin-vdc migrate xxx" when user has banner motd configurations with metacharacters as delimiter. Note not all metacharacters cause this issue. The following characters are known to cause this issue:

  [$ ^ . | ? \ >

  The following characters works fine

  { }]*+<

  Workaround: This issue is resolved.

- CSCur23435

  Symptom: VLAN configuration could get locked and not able to process any VLAN configuration changes anymore.

  Conditions: While configuring private VLAN associations from one terminal, if the user simultaneously deletes primary VLANs from another terminal PVLAN (Private-Vlan) get locked and cause lock of complete VLAN database.

Issue can happen only when configuration are done over multiple terminal session where one session configure private VLAN association and second deleting private VLANs before first session exit VLAN configuration.

Workaround: This issue is resolved.

- CSCur23769

  Symptom: After device reload some VDCs might failed to allocate interface

  Conditions: Issue might happen after device reload with high scale of VLAN/SVI/PBR on multiple LCs.

  Workaround: This issue is resolved.

- CSCur24077

  Symptom: Learned MAC addresses might still flood.

  Conditions: Ingress LC missing HW entry for forwarding engine.

  Workaround: This issue is resolved.

- CSCur24099

  Symptom: This is a feature enhancement to improve the handling of ejector open/close events. Open events of both top and bottom ejectors (On Fabric or Interface modules) will trigger the card to power down in preparation to be removed from the system.

  The following error message will be seen in the log:

  %PLATFORM-3-EJECTOR_STAT_CHANGED: Ejectors' status in slot <slot #> has changed, Top Ejector is CLOSE, Bottom Ejector is OPEN

  %PLATFORM-3-EJECTOR_STAT_CHANGED: Ejectors' status in slot <slot #> has changed, Top Ejector is OPEN, Bottom Ejector is OPEN

  Conditions: It is possible under certain conditions (excessive noise and/or interference) that ejector signals can be falsely interpreted. This can lead to ejectors being declared as open when in fact they are closed. It's also possible to see ejector open/close messages for slots without cards.

  If all Fabric modules see such a problem at the same time and power down, this will in turn trigger all Interface Module to also power down.

  The system will recover by itself and once the ejectors close again the module(s) will be powered back on.

  Workaround: This issue is resolved.

- CSCur24212

  Symptom: dhcp_snoop process crashing due to a memory leak.

Conditions: If there are drops in relay packet due to interface error then memory leak happens for the process.

Interface error counter can be seen in show ip dhcp relay statistics(drop reasons).

Workaround: This issue is resolved.

- CSCur26119

  Symptom: EIGRP interface/adjacency flap may cause some EIGRP prefixes not to be sent.

  Conditions: This has been seen with when using route tagging with distribute list, however this is not an absolute cause and effect.

  Occurrence of the issue is unpredictable given it also requires rare internal EIGRP condition when forming route-update packets.

  Workaround: This issue is resolved.

- CSCur29527

  Symptom: Cisco Nexus 7000 forwards IPv6 packets with link-local as source IPv6 address.

  Conditions: Cisco Nexus 7000 learns the default route from its BGP peer over MPLS link. The packet with link-local as source-address is seen to be forwarding.

  Workaround: This issue is resolved.

- CSCur33701

  Symptom: On a Cisco Nexus 7000 with LISP configured, after a line card reload, OIR, or insertion, not all of the LISP /32 routes may be downloaded to the card once it is powered up and is brought online. This results in traffic being forwarding to the supervisor and being limited by the COPP policy on the switch. Traffic flows may be forwarded at a reduced rate or not at all. This can be verified with a "show forwarding route <prefix>" command, which will show the /32 route installed on other line cards with a valid next-hop address, but on the problem linecard, it may show a /25 route with "receive" as it's next-hop.

  Conditions: This can occur with a Cisco Nexus 7000 running LISP Discovery, where /32 hosts are learned on the Cisco Nexus 7000 via LISP. The problem can appear when a line card is reloaded, OIRd, or a new card is installed, and the card is powered up and brought online. This affects traffic coming inbound to the ports off the affected card.

  Workaround: This issue is resolved.

- CSCur35982

  Symptom: LISP xTR doesn't register eid prefix, even though more specific prefixes exist in the routing table. We should be registering the prefixes. which have more specific prefixes reachable as well.

  Conditions: See example:

  R2# show ip lisp database

  LISP ETR IP Mapping Database for VRF "default" (iid 0), global LSBs:

  0x00000001

```
EID-prefix: 10.2.0.0/16, instance-id: 0, LSBs: 0x00000001, no-route
  Locator: 20.0.2.2, priority: 10, weight: 100
           Uptime: 00:04:32, state: up, local
R2# show ip route 10.2.0.0
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]

10.2.0.0/24, ubest/mbest: 2/0
    *via 20.1.2.1, Eth2/1, [110/20], 00:10:01, ospf-1, type-2
    *via 20.1.2.3, Eth2/1, [110/20], 00:10:01, ospf-1, type-2
R2#
```

In this case 10.2.0.0/16 should still be registered to the mapping-system, even though no 10.2.0.0/16 route exists in the rib.

Workaround: This issue is resolved.

- CSCur38749

  Symptom: A vulnerability in Cisco Locator/ID Separation Protocol (LISP) feature of Cisco NX-OS could allow an authenticated, remote attacker to use SSH connection to a non-existent host covered by lisp mobile prefix

  from outside of the DC and get presented with the xTR login prompt.

  An attacker still needs to have login credentials for NX-OS device in order to be able to log in.

  Conditions: LISP mobility and lisp enabled interface need to be configured.

  Workaround: This issue is resolved.

- CSCur40130

  Symptom: A user enters the command "show ip bgp regex". After 15 seconds of the command being entered, the exec command prompt appears, and there is no show output displayed.

  Conditions: This occurs to software release of 6.2.10 with over 400k entries in the BGP routing table.

  Workaround: This issue is resolved.

- CSCur50305

  Symptom: Vtp dropped into transparent mode due to some unexpected error condition during VLAN creation/modification,

  but the reserved VLAN 1002-1005 were not removed.

This is preventing vtp mode changes till the reserved VLANs are removed. In vtp transparent mode is possible to remove these VLANs but a second symptom is observed. Vlan manager is logging that vtp is client/server mode so the changes can not be done.

As a summary there are two symptoms due to failing to create/update VLAN :

1) Reserved VLAN 1002-1005 are not removed. Show VLAN will still displaying them.

2) When vtp drops into transparent mode, VLAN mgr is not able to update VLAN.

Ex:

swith(config)# no vlan 20-41

ERROR: VTP VLAN configuration not allowed when device is in CLIENT mode

Conditions: Feature vtp has to be enabled and an error condition has to happen during VLAN creation/modification.

Workaround: This issue is resolved.

- CSCur52940

  Symptom: Continuous snmpd process crashes leading to Supervisor HAP reset.

  Conditions: - poll CiscoEntitySensorMIB.

  - have 40G/100G ports and transceivers.

  - reload switch or reload 40G/100G module or OIR of 40G/100G transceiver

  Workaround: This issue is resolved.

- CSCur53280

  Symptom: On a Cisco Nexus device, ipqosmgr might crash during bootup. This specific bug would not cause a system running fine to initially reboot. This is only seen while booting.

  Conditions: The Cisco Nexus device is running QoS.

  Workaround: This issue is resolved.

- CSCur54182

  Symptom: Cisco Nexus 7000 configured for TACACS may face crash due to &quot;Tacacs

  Daemon hap reset&quot;

  Reason: Reset triggered due to HA policy of Reset

  Service: Tacacs Daemon hap reset

  Conditions: On a switch running NX-OS 6.2(10), if a very long command is

  given with remote authorization using TACACS enabled, a crash is seen in TACACS. Because

  TACACS expects the strings to be of size 255, it is unable to handle strings greater than 255.

  Workaround: This issue is resolved.

- CSCur54213

  Symptom: There is an inconsistency on behavior for vPC consistency-check on "Vlan xlt mapping".

+ Case 1 :

```
n7kb(config-if)# sh vpc consistency-parameters vpc 30 | i xlt
Vlan xlt mapping          1     Enabled  10->30        -
n7kb(config-if)#
n7kb(config-if)# sh vpc | i role|Po30
vPC role                        : secondary, operational primary
30   Po30      up     success     success                30
```

+ Case 2 :

```
n7kb(config-if)# sh vpc consistency-parameters vpc 30 | i xlt
Vlan xlt mapping          1     Enabled  10->30        Enabled
n7kb(config-if)#
n7kb(config-if)# sh vpc | i role|Po30
vPC role                        : secondary, operational primary
30   Po30      up     failed     Vlan mgr VPC compat check  30
```

Above two cases should behave in a same manner.

Both "-" and "Enabled" means VLAN translation is enabled but no parameters are configured.

Conditions: When using VLAN translation on a vPC.

Workaround: This issue is resolved.

- CSCur57243

  Symptom: OTV traffic gets black-holed.

  Conditions: Issue is seen when the OTV VDC is an F3 VDC and "no otv suppress-arp-nd" is configured and can be triggered after a switch reload, vdc reload, system switchover, adding an extended VLAN, and "clear ip route" in the otv vdc.

  Workaround: This issue is resolved.

- CSCur57579

  Symptom: 2014 Oct 25 03:14:49.506320 %LACP-5-LACP_SUSPEND_INDIVIDUAL: LACP port Ethernet7/18(0x1a311000) of port-channel Ethernet7/18(0x1a311000) not receiving any LACP BPDUs suspending (individual) port

  ```
  show inter eth152/1/10
  Ethernet152/1/10 is down (Internal-Fail errDisable, port-channel: required service is not responding)
  ```

  ```
  sh system internal mts buffers summary
  ```

```
node    sapno   recv_q  pers_q  npers_q log_q
sup     9589    0       0       1       0
sup     5445    14      0       0       0
sup     347     0       128281  0       0
sup     284     0       5       0       0
```

SAP 347 is LACP.

Conditions: Currently the known condition is when you enter the **show lacp interface** command of a FEX interface that is using LACP. LACP is able to operate, but with limited capability. Hence, LACP normal rate works fine, but fast rate members might flap or go suspended.

Workaround: This issue is resolved.

- CSCur61265

  Symptom: If you have a FEX uplink port that is 1/1 or 1/1/1 (using breakout) as part of the FPC then an ISSU to 6.2(10) will cause the ISSU to terminate.

  Conditions: This only happens is 1/1 or 1/1/1 are using as part of the FEX port-channel for the uplinks.

  Workaround: This issue is resolved.

- CSCur68350

  Symptom: A MAC address is present in hardware but not present in software. Since the MAC is not in software other processes and features are not notified of new MAC address. This can result in black-holed traffic.

  Conditions: This issue is triggered via ISSU

  This issue has been seen on F3 modules

  Workaround: This issue is resolved.

- CSCur69114

  Symptom: HSRP filtered packets are flooded to layer2 domain.

  This can result in HSRP flaps between DCI sites.

  Conditions: HSRP filtering applied on DCI interfaces.

  Workaround: This issue is resolved.

- CSCur70861

  Symptom: If F3 module experiences repeated single bit ECC errors it will error-disable the associated ports with that forwarding instance.

```
exception information --- exception instance 1 ----
Module Slot Number: 5
Device Id          : 197
Device Name       : Flanker L3 Driver
Device Errorcode  : 0xcc503600
Device ID         : 197 (0xc5)
Device Instance   : 03 (0x03)
Dev Type (HW/SW)  : 06 (0x06)
ErrNum (devInfo)  : 00 (0x00)
System Errorcode  : 0x429b0026 failure recovery threshold
Error Type        : Minor error
PhyPortLayer      : Ethernet
Port(s) Affected  : Ethernet5/25-32
Error Description : FLN_FW_INT_STATUS_TCAM_MATCH0_ECC_0
DSAP              : 0 (0x0)
UUID              : 0 (0x0)
Time              : Tue Nov 11 16:37:55 2014
                    (Ticks: 546281B3 jiffies)
```

Conditions: When repeated single-bit ECC errors are detected.

Workaround: This issue is resolved.

- CSCur71657

  Symptom: On a Cisco Nexus 7000/7700 series switch, multicast received on GRE tunnel is not sent to receivers

  Conditions: The multicast traffic is received on F3 series modules.

  Workaround: This issue is resolved.

- CSCur75014

  Symptom: When shutting down ports with promiscuous access/trunk configuration using range command, "sequence timeout" is seen for some ports.

  ```
  7004-4(config)# int e4/1-32
  7004-4(config-if-range)# shut
  7004-4(config-if-range)# show log log | eg seq
  %ETHPORT-5-IF_SEQ_ERROR: Error ("sequence timeout") communicating with MTS_SAP_PVLAN for opcode MTS_OPC_ETHPM_PORT_LOGICAL_CLEANUP (RID_PORT: Ethernet4/27)
  %ETHPORT-5-IF_SEQ_ERROR: Error ("sequence timeout") communicating with MTS_SAP_PVLAN for opcode MTS_OPC_ETHPM_PORT_LOGICAL_CLEANUP (RID_PORT: Ethernet4/26)
  ```

  Bringing up those ports in this state will result in misprogramming in CBL for those ports as well as egress translation in eltm will be lost for those parts.

```
vlan 93
  name 10.243.32.0/20-NFS-Primary
  private-vlan primary
  private-vlan association 94
vlan 94
  name 10.243.32.0/20-NFS-Isolated
  private-vlan isolated

interface Ethernet4/1
  switchport
  switchport mode private-vlan trunk promiscuous
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  flowcontrol receive on
  flowcontrol send on
  switchport private-vlan trunk native vlan 998
  switchport private-vlan trunk allowed vlan 94,99
  switchport private-vlan mapping trunk 93 94
  no shutdown

interface Ethernet4/6
  switchport
  switchport mode private-vlan trunk promiscuous
  spanning-tree port type edge trunk
  spanning-tree bpduguard enable
  flowcontrol receive on
  flowcontrol send on
  switchport private-vlan trunk native vlan 998
  switchport private-vlan trunk allowed vlan 94,99
  switchport private-vlan mapping trunk 93 94
  no shutdown
```

Working Port

```
show system internal eltm info interface e4/1

ELTM Detailed info for Interface Ethernet4/1
```

```
 cr_flags = INTF LIF , LIF = 16420 (0x4024), LTL = 56 (0x38) (S 0x0, P 0x0)
 IF_INDEX = 437780480 (0x1a180000), SHARED 0x0
 SEC_LTL info:
       sec_ltl = 0 (0x0), A/D progress = FALSE
       sec_ltl_cnt = 0, trig_ifidx = NULL
 State = UP FRR_NTFN_SENT = FALSE

Vlan Egress Translation: 0 1-------------------Egress Translation Present
       94       93
 Layer = L2, Mode = TRUNK
 Operational VLAN's (3):
  93,99,998
 Configured VLAN's (0):

 Line Card VLANs (2):
  93,99
 egress_vsl_drop = 0
 Interface Features:
  ipv4_en = 0, ipv4_mcast_en = 0,df_mask = 0, mpls_en 0
  v4_table_id = 0 (0x0), ipsg_en = 0
  per_pkt_ls_en = 0, icmp_redirect = 1, v4_same_if_check = 0
  ipv6_redirect = 1 ipv6_en = 0, ipv6_mcast_en = 0
  v6_table_id = 0 (0x0), v6_same_if_check = 0
  mtu = 1500 (0x5dc), f_index = 0 (0x0)
  bd = 2 (0x2)
  sub_type = 0x0, vpws_en = 0 fp_core_learn = 0x0
  mgmt_svi = 0x0,  admin_port_mode = 0x0,    bd_uuid[0] = 0 (0x0)
       num group ids = 0
       num ext group ids = 0
       num ext group ids = 0
    bd_uuid[1] = 0 (0x0)
       num group ids = 0
       num ext group ids = 0
       num ext group ids = 0
    bd_uuid[2] = 0 (0x0)
       num group ids = 0
       num ext group ids = 0
       num ext group ids = 0
  bd_uuid[3] = 0 (0x0)
     num group ids = 0
     num ext group ids = 0
     num ext group ids = 0
  bd_uuid[4] = 0 (0x0)
     num group ids = 0

module-4# show hardware internal ashburton port 1 state
Port:  1, Int-State:  Enabled, XAUI-Status: 0xfc,        Rate-mode: OSM
CFG0: Dest_flood: Disabled,    Dest_index: 0x0000
CFG1: Vlan:  998,              Virtual_chassis_id: 0
CFG2: Source_flood: Disabled,  Source_index: 0x0038
CFG3: Trusted_port: OFF,       L3_rewrite: OFF,          Notify_new_learn:  ON
    Disable_new_learn: OFF,  Disable_index_learn: OFF, Index_direct: OFF
    Conditional_learn: OFF,  Bundle_bypass: OFF,       Qos_override: OFF
    CSH-Type: 0x0
CFG4: Notify_index_learn: OFF, Bundle_id: 0, Rx_span: 0, Rx_capture: 0
    Control2_byte: 0x00
CFG5: Dont_forward_override_value: 0,        Dont_forward_override: OFF
    Dont_learn_override_value: 0,          Dont_learn_override: OFF
    Drop_untagged: OFF,                  Drop_tagged: OFF
    Control3_byte: 0x00,                 Decap_mode: AUTO
    Bundle_select: 0xff,                 Encap_mode: NATIVE
CFG6: SGT: 0x0000
CFG7: SPI_override: OFF, SPI_LSB1: 0x00, SPI_LSB2: 0
    Dest_index2: 0x0, Source_index2: 0x0
VSL: Ingress VSL: <OFF, 0>, Egress VSL drop: OFF
CBL:
CBL state: RX CBL disabled, TX CBL disabled

    VLANs on CBL state ASHBURTON_CBL_PORT_DISABLED(0):
    0, 2 - 92, 95 - 98, 100 - 997, 999 - 4095,
    4091 VLANs in ASHBURTON_CBL_PORT_DISABLED(0) state.

    VLANs on CBL state ASHBURTON_CBL_PORT_BLOCKED_LISTENING(1):
    1, 998,
    2 VLANs in ASHBURTON_CBL_PORT_BLOCKED_LISTENING(1) state.
    0 VLANs in ASHBURTON_CBL_PORT_LEARNING(2) state.
    VLANs on CBL state ASHBURTON_CBL_PORT_FORWARDING(3):
    93 - 94, 99,
    3 VLANs in ASHBURTON_CBL_PORT_FORWARDING(3) state.
CTS: inst 7        port_sel 3
    PortMode:      OSM
    CtsMode:       OFF
    PeerTrusted:   NO

=========================================================================
==
Broken port for which sequence timeout is seen

show system internal eltm info interface e4/6

ELTM Detailed info for Interface Ethernet4/6
 cr_flags = INTF LIF , LIF = 16425 (0x4029), LTL = 61 (0x3d) (S 0x0, P 0x0)
 IF_INDEX = 437800960 (0x1a185000), SHARED 0x0
 SEC_LTL info:
     sec_ltl = 0 (0x0), A/D progress = FALSE
     sec_ltl_cnt = 0, trig_ifidx = NULL
 State = UP FRR_NTFN_SENT = FALSE ----------------------Egress Translation Not seen
 Layer = L2, Mode = TRUNK
 Operational VLAN's (3):
  93,99,998
 Configured VLAN's (0):

 Line Card VLANs (2):
  93,99
 egress_vsl_drop = 0
 Interface Features:
  ipv4_en = 0, ipv4_mcast_en = 0,df_mask = 0, mpls_en 0
  v4_table_id = 0 (0x0), ipsg_en = 0
  per_pkt_ls_en = 0, icmp_redirect = 1, v4_same_if_check = 0
  ipv6_redirect = 1 ipv6_en = 0, ipv6_mcast_en = 0
  v6_table_id = 0 (0x0), v6_same_if_check = 0
  mtu = 1500 (0x5dc), f_index = 0 (0x0)
  bd = 2 (0x2)
 sub_type = 0x0, vpws_en = 0 fp_core_learn = 0x0
 mgmt_svi = 0x0,  admin_port_mode = 0x0,    bd_uuid[0] = 0 (0x0)
       num group ids = 0
       num ext group ids = 0
       num ext group ids = 0
   bd_uuid[1] = 0 (0x0)
       num group ids = 0
       num ext group ids = 0
       num ext group ids = 0
   bd_uuid[2] = 0 (0x0)
       num group ids = 0
       num ext group ids = 0
       num ext group ids = 0
   bd_uuid[3] = 0 (0x0)
       num group ids = 0
       num ext group ids = 0
       num ext group ids = 0
   bd_uuid[4] = 0 (0x0)
       num group ids = 0
       num ext group ids = 0
       num ext group ids = 0


module-4# show hardware internal ashburton port 6 state
Port:  6, Int-State:  Enabled, XAUI-Status: 0xfc,        Rate-mode: OSM
CFG0: Dest_flood: Disabled,    Dest_index: 0x0000
CFG1: Vlan:  998,              Virtual_chassis_id: 0
CFG2: Source_flood: Disabled,  Source_index: 0x003d
CFG3: Trusted_port: OFF,       L3_rewrite: OFF,          Notify_new_learn:  ON
      Disable_new_learn: OFF,  Disable_index_learn: OFF, Index_direct: OFF
      Conditional_learn: OFF,  Bundle_bypass: OFF,       Qos_override: OFF
      CSH-Type: 0x0
CFG4: Notify_index_learn: OFF, Bundle_id: 0, Rx_span: 0, Rx_capture: 0
      Control2_byte: 0x00
CFG5: Dont_forward_override_value: 0,        Dont_forward_override: OFF
      Dont_learn_override_value: 0,          Dont_learn_override: OFF
      Drop_untagged: OFF,                    Drop_tagged: OFF
      Control3_byte: 0x00,                   Decap_mode: AUTO
```

```
     Bundle_select: 0xff,                  Encap_mode: NATIVE
CFG6: SGT: 0x0000
CFG7: SPI_override: OFF, SPI_LSB1: 0x00, SPI_LSB2: 0
     Dest_index2: 0x0, Source_index2: 0x0
VSL: Ingress VSL: <OFF, 0>, Egress VSL drop: OFF
CBL:
CBL state: RX CBL disabled, TX CBL disabled

     VLANs on CBL state ASHBURTON_CBL_PORT_DISABLED(0):
     0, 2 - 92, 94 - 98, 100 - 997, 999 - 4095,
     4092 VLANs in ASHBURTON_CBL_PORT_DISABLED(0) state.

     VLANs on CBL state ASHBURTON_CBL_PORT_BLOCKED_LISTENING(1):
     1, 998,
     2 VLANs in ASHBURTON_CBL_PORT_BLOCKED_LISTENING(1) state.
     0 VLANs in ASHBURTON_CBL_PORT_LEARNING(2) state.

     VLANs on CBL state ASHBURTON_CBL_PORT_FORWARDING(3):
     93, 99, ==========================================>Vlan 94 is missing
     2 VLANs in ASHBURTON_CBL_PORT_FORWARDING(3) state.
CTS: inst 0        port_sel 2
     PortMode:       OSM
     CtsMode:        OFF
     PeerTrusted:    NO
```

This will result in packets not being forwarded out of this port in isolated VLAN.

Conditions: Pre-condition to hit issue:- Having five or more ports (L2 trunk or PC) in a

system with private VLAN with any other PVLAN mode than "host" can be affected with this bug upon seeing following]triggers:-

```
BOTTOM-7010(config-if)# switchport mode private-vlan ?
  host          Port mode pvlan host
  promiscuous  Port mode pvlan promiscuous
  trunk         Private-vlan trunk
```

Known Triggers :-

if the module is poweroff and then power-on using cli interface range command to shutdown five or more than five ports

interface range command to change port mode (mode change can be related or unrelated to private VLAN) on five or more that five ports

reload module x

Physically reseating of the line card.

Workaround: This issue is resolved.

- CSCur78552

  Symptom: No LACP suspend-individual is added on all ports.

  Conditions: ISSU upgrade from 6.1.x train to 6.2.10.

  Workaround: This issue is resolved.

- CSCur79328

  Symptom: Cisco Nexis 7000 time display is always 35 seconds ahead than the GM.

  Conditions: The time display issue is always observed prior to this fix.

  Workaround: This issue is resolved.

- CSCur83912

  Symptom: After switch reload PVLAN was crashing. This crash will only be seen in gdb images, and will not affect final images. This happens when PVLAN is already in an errored state and no PVLAN configuration is going through due PVLAN being in locked state.

  Conditions: The setup was already in an errored state where some of the port-channels in PVLAN db were in an errored state, whereas the member ports of the same port-channels were in a good state with the PVLAN mapping saved in PSS. When box reload was done, the PC and member port configuration did not match, hence when HW programming request for PC was sent after reload, it did not find any PVLAN mappings on PC from PSS and asserted.

  Workaround: This issue is resolved.

- CSCur89700

  Symptom: Debug with filter applied does not print all output it should.

  Conditions: Running debug in 6.2(10)

  Workaround: This issue is resolved.

- CSCur91392

  Symptom: ipfib cored when checking for routes with mpls adjacency.

  Conditions: Issue is seen only when entering the **show forwarding ipv4 route** [**next-hop** *next-hop* | **interface** *interface*] command for a route having at least one recursive next-hop and if one of the connected next-hops for this route is an MPLS adjacency.

The **interface** or **next-hop** keyword *must* be used in the **show forwarding ipv4 route** command to encounter this issue.

Workaround: This issue is resolved.

- CSCur96497

  Symptom: A newly configured port-channel for VLAN translation does not seem to forward traffic.

  Conditions: Order of configuration where the port-channel member has the VLAN translation commands defined before the port is added to a channel-group

  Workaround: This issue is resolved.

- CSCus02026

  Symptom: PIM may crash with high number of sources per group in VPC setup

  Conditions: High scale of sources per group

  Workaround: This issue is resolved.

- CSCur35650

  Symptom: Shortly after installing an N7K-F312FQ-25 line card, a Nexus 7009 running 6.2(8a) reboots.

  Conditions: Exact conditions not known.

  Workaround: This issue is resolved.

- CSCus08720

  Symptom: The dhcp_snoop service crashes.

  Conditions: This has been seen on a Cisco Nexus 7000 running 6.2(10) code.

  Workaround: This issue is resolved.

- CSCus18893

  Symptom: A Nexus 7000 crashed due to a kernel panic. CPU 8 seemed to have the issue:

```
Logging time: Fri Dec 12 06:00:10 2014
1418382010:126475629 process sysinfo (7268), jiffies 0x13c4b55df
Oops

REGISTERS:
CPU: 8   Process: sysinfo (pid 7268)   Tainted: P        W
RIP: 0010:[<ffffffffa28e7fc6>] mts_sys_my_node_addr_get+0x36/0xa0 [klm_mts]
RSP: 0000:ffff88042c12dbd8  EFLAGS: 00010286
RAX: 0000000000000501  RBX: ffff880431030000  RCX: 0000000000000001
```

```
RDX: 00000000335edb01  RSI: 000000000000431c  RDI: ffffffffa2cf93c0
RBP: ffff88042c12dbe8  R08: ffff88042c069b28  R09: 0000000000000002
R10: 000000022c12dc28  R11: 0000000000000000  R12: 00000000e7410020
R13: 0000000080044d1e  R14: ffff880431030000  R15: 00000000e7410020
CALL TRACE:
 CPU 8  Process: sysinfo (7268)
 [<ffffffffa28e7fc6>] mts_sys_my_node_addr_get+0x36/0xa0 [klm_mts]
[<ffffffffa28ce2c9>] mts_syscall+0x419/0x8b0 [klm_mts] (64)
[<ffffffffa28ce7fb>] mts_compat_ioctl+0x7b/0x2890 [klm_mts] (736)
[<ffffffff802f19c5>] compat_sys_ioctl+0xfb/0x28e (112)
[<ffffffff802298e2>] ia32_syscall_done+0x0/0x5 (131927355826352)
```

Conditions: The conditions leading to the crash are still being investigated.

Workaround: None

- CSCus24030

  Symptom: CRC errors and packet drops on traffic received from 100G F3 interface.

  Conditions: Issue happens when traffic arrives on a different F3 card (10G) with a dot1q tag and is sent out go 100G F3 card untagged - access port.

  Workaround: This issue is resolved.

- CSCuq88032

  Symptom: HSRP VMAC is not programmed with the G flag in the software MAC table on vPC peer who is HSRP standby:

```
RTP-Agg2# sh hsrp brief
*:IPv6 group   #:group belongs to a bundle
                    P indicates configured to preempt.
                    |
 Interface   Grp  Prio P State    Active addr      Standby addr     Group addr
  Vlan2       2    0      Standby  2.2.2.2          local            2.2.2.3          (conf)



RTP-Agg2# sh mac ad ad 0000.0c07.ac02
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link,
        (T) - True, (F) - False
  VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 2        0000.0c07.ac02    static       -       F     F  vPC Peer-Link
```

Conditions: This is condition is seen when the HSRP configuration includes a track object and the decrement value will set the priority to less than or equal to zero when the track object is in a down state:

```
interface Vlan2
  hsrp 2
    priority 80
    ip 2.2.2.3
    track 10 decrement 80 <<<<<<< This decrement value is greater than or equal to the configured or default priority
```

Workaround: This issue is resolved.

- CSCus21250

  Symptom: When entering the **show ip bgp route-map/filter-list** command, a BGP process crash can occur.

  2014 Dec 17 16:54:42 tskanai %SYSMGR-2-SERVICE_CRASHED: Service "bgp" (PID 12405) hasn't caught signal 6 (core will be saved).

  Conditions: The Cisco Nexus 7000 has 500k - 900k routes with ~10K attributes. The issue might occur if the command returns no matching results.

  Workaround: This issue is resolved.

- CSCus50031

  Symptom: GOLD PortLoopback test may not run (both Ondemand and Health Monitoring) on F3 module

  When running "show diagnostic result module <x> test PortLoopback detail", one will see a status of "DIAG TEST UNTESTED" if you are hitting this issue (assuming the test is not manually disabled via configuration).

  Conditions: PortLoopback diag test is not executed on F3 10G ports with 6.2(10) code.

  Only 6.2(10) is affected. Earlier and later releases will not have this issue.

  Workaround: This issue is resolved.

Resolved Caveats—Cisco NX-OS Release 6.2(10)

- CSCul46401

  Symptom: ACL removal times out.

  Conditions: This symptom might be seen with a 20K IPv4 and 20K IPv6 ACL on an interface and an attempt is made to remove the IPv6 ACL after the first time it times out.

Workaround: This issue is resolved.

- CSCul20441

  Symptom: A syslog message occurs for an ARP packet with source IP address of 0.0.0.0 and a destination local virtual IP address. It is a just probe ARP.

  Conditions: This symptom might be seen when an ARP probe packet is received with a source IP address of 0.0.0.0 destined to one of our local IP addresses.

  Workaround: This issue is resolved.

- CSCul63987

  Symptom: On an F1 Series module, the q transmitted bytes and packets are not displayed for Layer 2 packets with CoS 4 to 7.

  Conditions: This symptom might be seen when a Layer 2 packet marked with CoS 4 to 7 is flowing through an F1 Series module.

  Workaround: This issue is resolved.

- CSCul54591

  Symptom: All vPC port channels corresponding to TCB_RID mentioned in the following syslog can have traffic loss or duplicates on vPC legs depending on whether the peer's leg is up or down after you see the following message:

  ```
  2013 Oct 16 14:18:50 switch  PIXM-3-PIXM_SYSLOG_MESSAGE_TYPE_ERR
  Process:VDC-2-pixm_vl.Please collect show tech pixm-all, show tech pixmc-all
  .TCB_RID:0x9e,TXN_Type:67,Status:0x421c0017,Error_Type:0x1c.
  ```

  Conditions: This symptom might be seen when you insert a crossbar or a module.

  Workaround: This issue is resolved.

- CSCui63735

  Symptom: After you enter the **show ip mroute** command, you may see a large number of routes with a "delete pending" flag.

  Conditions: You might see this symptom when you have a large number of mroutes installed in the MRIB and enter the **no feature pim** command for example, which disables the PIM protocol completely from the switch. In this case, the PIM protocol may halt without sending the acknowledgments to MRIB process leading to the symptom.

  Workaround: This issue is resolved.

- CSCuq00300

  Symptom: After a supervisor failover, you cannot use the Secure Shell (SSH) to access the switch. You must change the password (on the console) before you can access the switch again using SSH.

  Conditions: This issue might happen if you use the regular **tacacs-server key** command during a supervisor failover.

  Workaround: This issue is resolved.

- CSCuo11989

  Symptom: IPv6 address configuration is not supported on tunnels on the Cisco Nexus 7700 series chassis.

  Conditions: You might see this symptom when you are working on the Cisco Nexus 7700 series chassis.

  Workaround: This issue is resolved.

- CSCuo10992

  Symptom: When you are working with the N7K-F248XP-25E module, you might see CRC errors on the last 8 ports.

  Conditions: You might see this symptom when you are working with the N7K-F248XP-25E module with the ports at 1-Gigabit speed and with the CTS configuration.

  Workaround: This issue is resolved.

- CSCun20590

  Symptom: You might see the RSVP fail after an SSO test.

  Conditions: You might see this symptom after RP SSO, after you reload module 1,2, and the standby supervisor.

  Workaround: This issue is resolved.

- CSCuo13699

  Symptom: The vPC leg on the primary and or secondary vPC might be in the FAILURE state while the vPC leg is being brought down.

  Conditions: This symptom might be seen when you have private VLAN configured on the vPC leg, which is brought down by entering the command **shutdown** command on the peer link or the vPC leg.

  Workaround: This issue is resolved.

- CSCum89656

  Symptom: You might see a MAC address with no port an/or switch ID assigned to it.

  Conditions: This symptom might be seen in a FEX vPC+ setup, where you have entered the **vdc reload** command on one side and the MAC address is synchronized from the peer.

  Workaround: This issue is resolved.

- CSCul51350

  Symptom: Traffic might be lost on ingress PE with data MDT configuration.

  Conditions: This symptom might be seen after you are working with the data MDT configuration.

  Workaround: This issue is resolved.

- CSCun16295

Symptom: The MAC addresses may be learned in the incorrect BD.

Conditions: You might see this symptom after you perform an ISSU from Cisco NX-OS Release 6.2(6) to Release 6.2(6a) with OTV, followed by allowed VLAN additions or the ports going up and down.

Workaround: This issue is resolved.

- CSCuo13937

  Symptom: When a vPC is configured in private VLAN host mode and vPC is up, and the vPC leg is configured to be promiscuous, you might see the following error:

  ```
  Failed to configure hardware
  ```

  Conditions: This symptom might be seen when one of the vPC legs is still in some other mode than promiscuous, and the configuration is applied on the other vPC leg to make the mode promiscuous. The actual error is seen when the promiscuous mapping is configured.

  Workaround: This issue is resolved.

- CSCul44598

  Symptom: If you have hosts configured with SPT thresholds as infinity in a network with sparse mode hosts, you might see an intermittent traffic loss for hosts.

  Conditions: This symptom might be seen when the host with and SPT threshold of infinity and the sparse mode host share the common intermediate router, which is in the shared tree path for both the hosts and also in the (S, G, R) prune path from the sparse mode host while it sends joins to the source tree.

  Workaround: This issue is resolved.

- CSCul36036

  Symptom: An error is displayed when you delete and then add VLANs.

  ```
  Failed to run the commands. Please try again later.
  ```

  Conditions: This symptom might be seen when you perform sequential delete and add VLANs.

  Workaround: This issue is resolved.

- CSCum07107

  Symptom: ISIS adjacencies will not form for Overlay1 with the trigger provided in the Conditions section.

  Conditions: This symptom might be seen when you set the trigger as follows: Change the join interface for Overlay1 to the same join interface for Overlay2. This action allows two different overlays to share the same join interface.

  Workaround: This issue is resolved.

- CSCun68731

Symptom: You might see the FabricPath ISIS process go down while recovering the LSP database from its PSS for stateful recovery. This would occur when you switch over the active supervisor with FabricPath ISIS configured with multiple topologies.

Conditions: This symptom might be seen when you are running FabricPath ISIS with multiple topologies and perform a switchover.

Workaround: This issue is resolved.

- CSCuo53059

  Symptom: Because of a specific sequence of PVLAN events, the PIXM process might have undergone memory corruption, which can lead to of the following issues:

  - Port channel and its members are in the error-disabled state
  - Incorrect CBL programming
  - PIXM goes down.

  Conditions: This issue can happen **only** if you have performed an ISSU from Cisco NX-OS Release 6.2(6) or Release 6.2(6a) to Release 6.2(8).

  This symptom might be seen when you have made the following configurations running on Cisco NX-OS Release 6.2(6) or Release 6.2(6a):

  - Primary VLAN is associated to secondary VLAN
  - Port channel is a trunk port channel carrying both the primary and secondary VLANs
  - STP operates on either the primary or secondary VLAN on the trunk port channel

  Then, when you perform an ISSU to Release 6.2(8) and modify the port channel (such as adding or removing a port), you might see one of the symptoms listed above.

  Workaround: This issue is resolved.

- CSCun16347

  Symptom: CLI hangs. To get back CLI, you either need to power cycle or reload the router by telnet to the mgmt IP.

  Conditions: 1. This happens in 6.2.6 with "system admin-vdc migrate xxx" when user has banner motd configurations with "!" and "#" as delimiter.

  2. "#" and "!" work in all other releases except 6.2.8. The following characters are known to cause similar issue:

  [$ ^ . | ? \ >

  The following characters works fine

  { }]*+<

  Workaround: This issue is resolved.

- CSCul96145

  Symptom: STP sets the native VLAN to disable instead of blocking when you enter a **shut lan** command on the secondary vPC peer on a phy-port leg, immediately followed by entering a **shut lan** command on the same phy-port vPC leg on the primary switch.

Conditions: This symptom might be seen when you are working on a secondary vPC switch on a phy-port vPC leg. You might seen this if you enter a **shut lan** command on the secondary vPC peer on a phy-port leg, immediately followed by entering a **shut lan** command on the same phy-port vPC leg on the primary switch.

Workaround: This issue is resolved.

- CSCuq04394

  Symptom: Route is stuck in a cc pending state, as indicated in the following output:

  ```
  x.x.x.x/29, ubest/mbest: 1/0, attached, cc pending
  y.y.y.y/28, ubest/mbest: 1/0, attached, cc pending
  z.z.z.z/21, ubest/mbest: 1/0, attached, cc pending
  ```

  The output from the s**how forwarding inconsistency module all** command is as follows:

  ```
  # show forwarding inconsistency module all

      IPV4 Consistency check : table_id(0x1)
      Execution time : 2400001 ms (timedout)
      No inconsistent adjacencies.
      No inconsistent routes.
  ```

  Conditions: This issue might be seen if Release 6.2(6a) is running on your switch and the switch has a large RIB/FIB.

  Workaround: This issue is resolved.

- CSCum06321

  Symptom: The prefix insertion in TCAM fails but the system does not generate a TCAM resource exhaustion message.

  Conditions: This symptom might be seen if the TCAM resource is larger than the DRAM resource and the TCAM resource exhaustion is not reached before the DRAM resource exhaustion. You can use the **show forwarding internal errors** command to display the DRAM exhaustion message.

  Workaround: This issue is resolved.

- CSCuj12958

  Symptom: U6RIB structural errors might be seen during withdrawing and/or adding routes.

  Conditions: This symptom might be seen when there is a route withdraw or when new routes are advertised.

  Workaround: This issue is resolved.

- CSCui65661

  Symptom: After a configuration replay, virtual fabric interface (VFI) members that inherit a port profile have an incomplete configuration.

  Conditions: This symptom might be seen when encapsulation Multiprotocol Label Switching (MPLS) is configured for the port profile.

  Workaround: This issue is resolved.

- CSCup76149

  Symptom: The secondary VLAN traffic does not egress out of the switch from a private VLAN promiscuous mode vPC.

  Conditions: This issue might be seen if your switch is running Release 6.2(x) and one of the following is true:

  - The promiscuous vPC is configured and up and the **no vpc** and **vpc** *x* commands are configured on the VPC secondary leg port-channel.
  - The community host vPC is configured and up and the **no vpc** and **vpc** *x* commands are configured on the vPC secondary leg port-channel.
  - A trunk/pvlan trunk mode vPC is changed to a private-VLAN host/promiscuous mode vPC.

  Workaround: This issue is resolved.

- CSCuo91958

  Symptom: On F2/F3 cards, at the tunnel termination point, the frames are replicated further towards the trunk interfaces. As the bug describes, the default CoS value is not retained (what was set in tunnel head end) and instead would be observed to be 0.

  Conditions: The frames as observed on a downstream trunk port at the point of tunnel termination.

  Workaround: This issue is resolved.

- CSCur30073

  Symptom: On a VPC+ setup where Fabricpath multi-topology feature is being used and "no port-channel limit" is not being used, unicast traffic arriving at certain ports in the default topology, directed towards the VPC hosts might fail under certain race conditions.

  The failure might be encountered if in the system described above, if the first few VPCs are de-configured followed by moving ports of the VDC to another VDC and then moving them back in. Based on certain race conditions, this sequence might sometimes cause incorrect forwarding of unicast traffic arriving at those ports (that were moved) if the destination is behind the VPCs on the VLANS belonging to the default topologies.

  Conditions: All of the below conditions are required at a minimum before there is a possibility to hit this issue.

  1. VPC+ with multiple Fabricpath topologies in use.

  2. No port-channel limit configuration is not being used under VPC domain. If this configuration is being used the issue will not be encountered.

  3. The first few VPC legs that were brought up are de-configured.

  4. Certain ports of the VPC+ VDC are moved out to another VDC and moved back in.

  5. Unicast Traffic is arriving in the switch at any of those ports is directed towards the hosts behind the VPCs on the default Fabricpath topology may experience failure. Traffic on non-default topologies should not be affected.

  This issue can be encountered if the ports that are moved between VDCS are F2, F2E or F3.

  Workaround: This issue is resolved.

- CSCur01307

Symptom: Delayed receiver (3 minutes) traffic for newly joined receiver for multicast traffic that is already flowing.

Conditions: Occurs when there is a recent RP change and the new RP did not have an (S,G) state for the appropriate group.

Workaround: This issue is resolved.

- CSCul53494

  Symptom: The IP multicast ping over GRE tunnel does not work. The ping packets are dropped at the tunnel destination.

  Conditions: The first or few packets reach the supervisor with the proper SHIM header at the remote side until the (S,G) route is installed in the hardware. All further packets hit this (S,G) and punt to supervisor without the SHIM header, which causes these packets to be dropped by the packet manager. This is a special case where we have (S,G) with punt flag and no OIFs associated with the route. The problem is only seen for tunnel interfaces that need a decap at the tunnel end. A multicast ping over a nontunnel interface does not have this problem.

  Workaround: This issue is resolved.

- CSCun25245

  Symptom: Packets with unicast IP addresses and multicast MAC addresses are duplicated on the destination interface, which can cause performance issues to the application**.**

  Conditions: This symptom might be seen when you are working with MS NLB Option 2: Static ARP + MAC-based L2 Multicast Lookups + Static Joins + IP Multicast MAC.

  Workaround: This issue is resolved.

- CSCuj45625

  Symptom: Not all HSRP secondary addresses are displayed when many addresses are configured. This situation affects the output of the **show running configuration** command and the **show hsrp** command.

  Conditions: This symptom might be seen when over 126 HSRP secondary virtual addresses are used on a single group. Some of the secondary addresses might not be displayed in the output of the **show running configuration** command.

  If the switch reloads, some secondary addresses are lost from the configuration.

  In addition, the output of the **show hsrp** command is truncated when many secondary addresses are configured per group. This issue is largely cosmetic.

  Workaround: This issue is resolved.

- CSCuq98748

  Symptom: Cisco NX-OS contains a version of Bash that is affected by vulnerabilities.

  Common Vulnerability and Exposures (CVE) IDs:

  CVE-2014-6271

  CVE-2014-6277

  CVE-2014-7169

CVE-2014-6278

CVE-2014-7186

CVE-2014-7187

Conditions: Occurs when user logs in and authenticates via SSH to trigger this vulnerability.

Workaround: This issue is resolved.

- CSCur03111

  Symptom: MACSEC link goes down in SAP AUTHEN INCOMPLETE state.

  Conditions: Occurs when:

  - C is enabled.

  - Egress traffic rate exceeds the effective line rate for encrypted traffic. (At 6.5G of traffic, CRC errors appear on the peer port and link goes down.)

  Workaround: This issue is resolved.

- CSCuq54506

  Symptom: Host moves across an overlay and becomes not reachable. Devices in the original location still have the MAC address of old location.

  Conditions: Occurs when configured with OTVs and RARPs are lost.

  Workaround: This issue is resolved.

- CSCuq55749

  Symptom: New HSRP follow group becomes stuck in initial state with "No Master for Slave" status.

  When using HSRP MGO an attempt to configure a new HSRP follow group may result in this group being stuck in the Initial state with HSRP reporting a reason of "No Master for Slave".

  Conditions: Occurs when HSRP MGO is configured, some HSRP follow groups are deleted from the running configuration, and a supervisor switchover is performed.

  Workaround: This issue is resolved.

- CSCuq81826

  Symptom: High latency on traffic going through F3 module.

  Conditions: Occurs when:

  - There is an F3/F2 module in an F-series only chassis.

  - ISSU performed.

  - There is an incorrect COS-to-Queue mapping.

  Workaround: This issue is resolved.

- CSCul05217

  Symptom: Cisco NX-OS software contains a directory traversal vulnerability.

Conditions: Occurs when a user uses the "copy" command to duplicate the contents of arbitrary files on the device to a user writable area of the file system. The copied contents can be viewed in the user writable area of the file system.

Workaround: This issue is resolved.

- CSCup10049

  Symptom: A Nexus C7009 running NX-OS version 6.2(2a) may have a crash in aclqos with accompanied core file(s).

  %SYSMGR-SLOT3-2-SERVICE_CRASHED: Service "aclqos" (PID <snip>) hasn't caught signal 6 (core

  will be saved).

  Post crash, the below messages may also be seen:

  %IPQOSMGR-3-QOSMGR_PPF_ERROR: PPF library error: MTS

  Error 0x801c0010 (Device or resourc e busy) after 1706 retries

  Conditions: There were too many queries from ipqosmgr and COPP to aclqos.

  Workaround: This issue is resolved.

- CSCur12257

  Symptom: Unicast flooding for packets destined to HSRP VMAC.

  HSRP VMAC for some VLANs may not have G and R bit set. This will result in unicast flooding for the packets destined to HSRP VMAC.

  Conditions: Occurs when there is a switchover (ISSU or EPLD upgrade of supervisor) and followed by flapping of the peer-link or bringing up the peer-link. This results in some VLANs not having the G and R bit set for HSRP VMAC.

  Workaround: This issue is resolved.

- CSCuq50590

  Symptom: Multicast forwarding stops working resulting in lost multicast traffic.

  Conditions: Occurs when:

  - The unicast client adds a route in urib with nexthop N1.

  - The multicast client adds a route in urib with the same nexthop N1.

  - The unicast client deletes its route.

  When the unicast routes are deleted, it triggers the removal of bindings which exist between N1 and resolved routes.

  Workaround: This issue is resolved.

- CSCup85198

Symptom: Error message not received indicating FEX ID already in use.

Conditions: Occurs when a non-default VDC's FEX with an erroneously configured duplicate FEX ID tries to come online. Observed when:

- Configured FEX ID or VPC in a non-default VDC in a dual supervisor setup in Cisco NX-OS pre-release 6.2.8 version, reload the non-default VDC, and switchover after system is in full redundancy.

- ISSU from older version to version up to Cisco NX-OS release 6.2.8(including 6.2.8).

Workaround: This issue is resolved.

- CSCuh81367

  Symptom: Following an ISSU/D, updating the CoPP profile takes a long time to complete.

  Conditions: Always occurs.

  Workaround: This issue is resolved.

- CSCuh29201

  Symptom: After a reload or switchover, some multicast flows might experience complete packet loss.

  Conditions: Occurs for layer 3 interfaces or SVIs with multiple PIM related configurations.

  Workaround: This issue is resolved.

- CSCur11131

  Symptom: Interfaces are not recognized on switch in (privilege mode, config terminal prompt, or show interface status).

  Conditions: Occurs when the set up has 400+ VLANs and 500 layer 2 ports or port-channels.

  Workaround: This issue is resolved.

- CSCuq31720

  Symptom: ES common LID is 0x0 after switchover or after ISSU.

  Conditions: Occurs for M1/F1 linecards with VPC+ in multiple VDCs.

  Happens only to the VPC secondary.

  Workaround: This issue is resolved.

- CSCup93375

  Symptom: All interfaces were VDC unallocated as result of reload upgrade.

  Conditions: Occurs when performing reload upgrade from Cisco NX-OS release 5.2.5 to Cisco NX-OS pre-release of 6.2.10.

  Workaround: This issue is resolved.

- CSCuh74299

Symptom: When the standby bootflash is full, standby force reload fails to bootup.

Conditions: The standby image sync triggered during the standby bootup preventing other processes from being scheduled properly. This results in heartbeat failure to the processes and causes the standby reset.

Workaround: This issue is resolved.

- CSCuj88721

  Symptom: vPC forwarder election fails when same VLAN does not exist on peer vPC switch.

  Conditions: When a vPC VLAN is made into a non-vPC VLAN on a vPC switch and has the lowest numbered vlan-id among the vPC VLANs on the peer, then it causes the vPC forwarder election to fail if that same VLAN does not exist on the other vPC switch (peer).

  Workaround: This issue is resolved.

- CSCuo70393

  Symptom: PVLAN promiscuous trunk tags native VLANs.

  Conditions: Occurs when a PVLAN promiscuous trunk is configured with the “switchport private-VLAN trunk native VLAN X” command.

  Workaround: This issue is resolved.

- CSCup20959

  Symptom: M2 linecard might be reset by the supervisor

  Conditions: The M2 linecard might be reset by the supervisor as a result of EOBC heartbeat being missed by the linecard.

  Workaround: This issue is resolved.

- CSCup27711

  Symptom: Heavy packet loss from all traffic hitting a CPU/glean adj.

  Conditions: Occurs when the F2 line cards/deleting FCoE VDC and moving the FCoE allocated interfaces to the ethernet VDC and the default route (static route) is not installed in the hardware on an F2e module.

  Workaround: This issue is resolved.

- CSCuq34350

  Symptom: F3 linecard might be reset by supervisor.

  Conditions: Occurs when line card does not respond to heartbeat messages with no logs saved.

  Workaround: This issue is resolved.

- CSCuq53321

  Symptom: Traffic across peer-link is not sent correctly resulting in traffic loss because of control plane protocols like HSRP becoming out of sync.

Conditions: Occurs when flapping VPC links configured as access ports.

Workaround: This issue is resolved.

- CSCuq82625

  Symptom: EIGRP Router sends classic updates that are ignored by peer. Peer no longer adds routes from that peer.

  Conditions: Occurs when:

  - The EIGRP process is configured "metric version 64 bit".

  - There are invalid values for "classic / wide metric" peers from the "show eigrp interface" command (the number of classic peers is greater than the number of total peers).

  Workaround: This issue is resolved.

- CSCuq93325

  Symptom: Unable to configure new ports to inherit an existing port-profile resulting in the 'port-profile' service being terminated.

  This situation will be accompanied by the 'port-profile' service terminating unexpectedly,

  Once this has occurred, it will not be possible to set additional ports to inherit the profile, or to run other port-profile commands such as 'show running-config port-profile'.

  Conditions: Occurs on switch running Cisco NX-OS releases 6.2(2a) and 6.2(6b) when the interface description contains the word "shutdown" prior to inheriting the port-profile:

  Workaround: This issue is resolved.

- CSCur05393

  Symptom: Expected message is not received. Interpreting wrong message payload causes crash and message: %SYSMGR-2-SERVICE_CRASHED: Service "diag_port_lb" (PID 4775) hasn't caught signal 6 (core will be saved).

  Conditions: Occurs when a request-response message happens to have an MTS internal msgid of 0x80000000 and the sender uses the returned rr_token of 0 to receive the response. This results in the first message in the queue. The expected message is not received. Interpreting the wrong message payload causes the crash.

  Workaround: This issue is resolved.

- CSCur03111

  Symptom: When MACSEC is configured, traffic limited to 5 gigs of traffic.

  Conditions: Occurs when MACSEC is configured.

  Workaround: This issue is resolved.

- CSCur06193

  Symptom: With deny ACE feature enabled, packets take PBR route, not default routing through ACL.

  Conditions: Occurs when:

- The "hardware access-list allow deny ace" feature is configured.

- The route-map with deny sequence is enabled.

- The permit ACE in ACL is referenced in PBR.

Workaround: This issue is resolved.

- CSCud84480

  Symptom: As a result of excessive memory usage, the SNMPD process restarts.

  Conditions: This occurs because of SNMP polling.

  Workaround: This issue is resolved.

- CSCud88400

  Symptom: User inherits permissions on the default VDC that they would have had on their own VDC.

  Conditions: When a user from a non-default VDC is configured with an SSH key, the key can be used to login to the default VDC.

  Workaround: This issue is resolved.

- CSCug75245

  Symptom: ERROR: Inherit request received on offline interface

  Conditions: Occurs when removing a port-profile and the FEX interface is in an 'up/up' state.

  Workaround: This issue is resolved.

- CSCuh27510

  Symptom: Egress policies/ACLs on interfaces that belong to the line card that is being reloaded, are not applied in the software forwarding path after the line card reload

  Conditions: Occurs after the line card is reloaded.

  Workaround: This issue is resolved.

- CSCui36490

  Symptom: BGP outbound messages get stuck in certain conditions.

  Conditions: When BGP is updating its peers with all route updates, the sender thread might go to sleep. When this occurs, the residual outbound messages to the peers are stuck until another message (such as keepalive, update, etc.) gets queued to wake the sender thread.

  Workaround: This issue is resolved.

- CSCuj16682

  Symptom: After reload, the ACL applied on mgmt interfaces fail to act.

Conditions: When a reload completes, the ACL that is applied on mgmt interfaces does not work even though telnet to mgmt IP appears to be working correctly.

Workaround: This issue is resolved.

- CSCuj63314

  Symptom: SNMP response for fan down not displayed.

  Conditions: After removal of the fan trays (fab and system fans), the SNMP response for the removed fan trays is not displayed.

  Workaround: This issue is resolved.

- CSCul33530

  Symptom: The F2 FCOE license displayed as being installed even though it not.

  Conditions: Occurs when the TRANSPORT_SERVICES_PKG license is installed in the system.

  Workaround: This issue is resolved.

- CSCul57133

  Symptom: During reload ACLQOS crashes and F2E line card brought down.

  Conditions: Occurs when the nq-7e policy is configured on a system that was previously configured with the 8e policy.

  Workaround: This issue is resolved.

- CSCum20932

  Symptom: Configuring a QoS ingress service-policy is rejected with error message:

  ERROR: Unable to perform the action due to incompatibility.

  Conditions: While configuring a QoS ingress service-policy directly on an ethernet interface, where only classes exists which match an access-group and cos, is rejected with the error message:

  ERROR: Unable to perform the action due to incompatibility: Module 3 returned status Policies with classes containing combined 'match dscp', 'match cos', 'match precedence' or 'match qos-group' are not supported. Only the same match type is supported between classes.

  Workaround: This issue is resolved.

- CSCun34206

  Symptom: Switch crashes after switchport configuration.

  Conditions: Occurs while configuring the switchport on an interface of a non-default VDC.

  Workaround: This issue is resolved.

- CSCun82155

  Symptom: When primary switch is reloaded, the vPC secondary switch might flap the vPC legs.

  Conditions: Occurs when the management interface is used for vPC peer keepalive messages and the vPC peer-link is on F2E or F3 cards.

Workaround: This issue is resolved.

- CSCun99528

  Symptom: ARP not being resolved and communication not working.

  Conditions: Occurs when:

  - The VPC with peer-gateway enabled

  - The private-VLAN is configured on the switch and access switch.

  - The peer-link is built on F1 ports.

  - The affected traffic does not go directly to the access switch, but through the peer and peer-link.

  Workaround: This issue is resolved.

- CSCuo17554

  Symptom: Peer-link between two F2E modules remains at MTU 1500.

  Conditions: By default, the peer-link is supposed to have an MTU value of 9216. However, when creating a peer-link between two F2E modules, the peer-link remains at MTU 1500.

  When trying to ping across the peer-link with a larger MTU value, it fails (with or without df-bit set for the pings)

  Workaround: This issue is resolved.

- CSCuo36343

  Symptom: After upgrading image, ipqosmgr crashes.

  Conditions: Upgrading dual supervisors from Cisco NX-OS release 6.1.3 to Cisco NX-OS release 6.2.2a.

  Workaround: This issue is resolved.

- CSCuo41201

  Symptom: Netstack exhausts buffers.

  Conditions: When IPv6 features are configured, such as OSPFv3, DHCPv6 and HSRPv6, Netstack runs out of buffers.

  Workaround: This issue is resolved.

- CSCuo41243

  Symptom: Flooding of frames occurs with traffic ingressing an M2 series line card when no TCN, MAC timeout, or manual clear of the MAC table occurs.

  Conditions: Occurs when traffic ingresses an M1/M2 series line card to be switched, and the ARP/ND entry for the IP which is associated with the MAC address ages out.

  Workaround: This issue is resolved.

- CSCuo46541

  Symptom: cmp_install.log continually increments resulting from 'show version' command.

  Conditions: Occurs when upgrading to Cisco NX-OS release 6.2(6a) using ISSU.

  Workaround: This issue is resolved.

- CSCuo51069

  Symptom: When igmp is configured under cfs region, igmp state is not in sync between vPC pairs.

  Conditions: Occurs when igmp is configured under the cfs region.

  Workaround: This issue is resolved.

- CSCuo66929

  Symptom: ulib core file is generated.

  Conditions: Occurs when running the 'show mpls switching internal fec label' command.

  Workaround: This issue is resolved.

- CSCuo76209

  Symptom: While performing an ISSU, the "pixm_vl" or "pixm" process crashes.

  Conditions: The crash occurs during an ISSU upgrade to Cisco Nexus release 6.2(8) or Cisco Nexus release 6.2(8a).

  Workaround: This issue is resolved.

- CSCuo76512

  Symptom: Connected route from a named vrf cannot be leaked into the default vrf using static routes when the next-hop is not specified.

  Conditions: Occurs when the next-hop is not specified.

  Workaround: This issue is resolved.

- CSCuo81438

  Symptom: AH packets (protocol 51) with an invalid AH header or fragmented AH packets with an offset greater than zero might be dropped by F2/F2E modules.

  Conditions: Occurs on F2/F2E modules for AH packets and fragmented AH packets.

  Workaround: This issue is resolved.

- CSCup11309

  Symptom: Vulnerability in HSRP authentication could allow an unauthenticated, adjacent attacker to affect the state of HSRP group members and cause blackholing of traffic.

Conditions: Occurs on devices configured for TEXT or MD5 group authentication that accept malformed HSRP packets leading to bypass of authentication. A potential attacker can affect the state of HSRP group members, causing them to release ACTIVE/STANDBY roles and go back to SPEAK state.

Workaround: This issue is resolved.

- CSCup13566

  Symptom: Error Message Guide for 6.x lists only VSAN as a reason for error message ETHPORT-5-IF_DOWN_INACTIVE.

  Conditions: When a VLAN is suspended or deleted.

  Workaround: This issue is resolved.

- CSCup18643

  Symptom: XCVR-ALARM

  Conditions: Occurs when device goes from shut to no-shut,

  Workaround: This issue is resolved.

- CSCup19027

  Symptom: Switch powers up with no configuration.

  Conditions: When the "copyrunning-config startup-config" command was entered, it was terminated abnormally.

  Workaround: This issue is resolved.

- CSCup24634

  Symptom: In problem state, BGP sessions do not come up after reset sessions. Reset triggers are clear ip bgp * and interface flap.

  Conditions: Occurs when the password is configured under some neighbors on both sides and the the password is removed from at least one BGP neighbor on both sides. This can be triggered by issuing "no password" under neighbor or "no neighbor x.x.x.x".

  Workaround: This issue is resolved.

- CSCup30064

  Symptom: Queue dropped packets counters for the "show policy-map interface" command do not work for M1 and F1 cards.

  Conditions: Occurs on M1 and F1 cards.

  Workaround: This issue is resolved.

- CSCup44154

  Symptom: The %SYSMGR-2-SERVICE_CRASHED syslog error message might be triggered by the 'tunnel' process in an OTV unicast transport configuration.

The core file generated by the 'tunnel' process should be visible under 'show core' command output.

Conditions: Occurs when the OTV unicast transport configuration is configured on Cisco NX-OS release 6.2(6).

Workaround: This issue is resolved.

- CSCup65916

  Symptom: An F3 linecard might reload after a system switchover.

  Conditions:

  Workaround: This issue is resolved.

- CSCup66113

  Symptom: While in enable mode, not able to create network admin user account.

  Conditions: Occurs in enable mode.

  Workaround: This issue is resolved.

- CSCup67234

  Symptom: Replies to SNMP requests from a server not permitted in SNMP ACL.

  Conditions: Occurs when:

  1) The new SNMP ACL syntax is used

  - snmp-server community [community] use-ipv4acl [SNMP-ACL]

  - snmp-server community [community] use-ipv6acl [SNMP-ACL]

  2) A supervisor switch over has occurred.

  Workaround: This issue is resolved.

- CSCup70386

  Symptom: A syslog message is printed while the logflash is absent.

  Conditions:

  Workaround: This issue is resolved.

- CSCup75514

  Symptom: Breakout port is stuck in error disabled status

  Conditions: Occurs when:

  - The interface breakout is done in a non-default VDC.

  - The ' copy r s ' command is NOT done on the non-default VDC.

  - Reloading the non-default vdc.

  - When the VDC comes online, perform the 'interface breakout' on the same interface.

  Workaround: This issue is resolved.

- CSCup83034

  Symptom: MAC address on FEX HIF interfaces (access ports) learned on wrong VLAN.

  Conditions: Occurs after a reload on an OTV setup running Cisco NX-OS release 6.2.8a with FEX modules connected to the OTV VDC.

  Workaround: This issue is resolved.

- CSCup88921

  Symptom: The VLANs (a part of the timed-out LOGICAL_BRINGUP message) are in suspended state on the affected port.

  Conditions: Occurs on setups with several FEXs and multiple fabric port channels set as SPAN sources on an active (no shut) SPAN session when one of the following triggers are executed:

  1. Reload box with config copied.

  2. VDC reload.

  3. Reloading the module with multiple SPAN source fabric port channels.

  Workaround: This issue is resolved.

- CSCup89222

  Symptom: When FIB programming on the module does not match with the RIB, traffic is blackholed.

  Conditions: Occurs running Cisco NX-OS release 6.2.8a with PBR applied on the ingress SVIs and the F3 module was reloaded and in use.

  Workaround: This issue is resolved.

- CSCup90167

  Symptom: The PBR redirect adjacent index in TCAM might point to 0x1 or other wrong next-hop index, causing PBR to fail.

  Conditions: Occurs when the PBR TCAM entries programming and FIB adjacent entries programming are done at the same time when the interfaces are brought up.

  Workaround: This issue is resolved.

- CSCup95948

  Symptom: LISP core file generated.

  Conditions: Occurs while configuring with LISP Multihop, when the node acting as the first hop and two eid_notify gateways are configured.

  Workaround: This issue is resolved.

- CSCup96023

  Symptom: mgmt 0 cannot be configured as "no shutdown force".

  Conditions: Occurs when:

  - mgmt 0 is connected to the device and up status.

- Cable removed from mgmt 0 and shutdown mgmt 0.

- When reconnecting cable on mgmt 0, the neighbor port link is up, but mgmt 0 is in shutdown.

- Using the "no shutdown" command on mgmt 0 is not reflected in the configuration and mgmt 0 is still in "shutdown force" status.

Workaround: This issue is resolved.

- CSCuq14012

  Symptom: MFIB (Multicast forwarding information base) stops updating statistics for multicast routes causing the S,G corresponding to the multicast route to time out.

  Conditions: Occurs when:

  1) A new module added to a chassis (all interfaces unallocated).

  2) An ISSU is performed.

  3) After ISSU, allocate interfaces to VDC. At this point, multicast statistics for traffic begin hitting the forwarding engine for the newly allocated interface, but do not increment.

  Workaround: This issue is resolved.

- CSCuq14407

  Symptom: snmppoll of ifHCOutOctects OID for a FEX HIF might have a greater value than the actual byte count and/or counters on the FEX interfaces might appear to go backwards or have very large spikes.

  Conditions: Occurs on FEX interfaces.

  Workaround: This issue is resolved.

- CSCuq22831

  Symptom: The snmpd process crashes multiple times resulting in a HAP-Reset:

  Conditions: Occurs while polling interface counter statistics via SNMP.

  Workaround: This issue is resolved.

- CSCuq30850

  Symptom: SNMPD crashes at snmp_log_string resulting in supervisor reset.

  Conditions:

  Workaround: This issue is resolved.

- CSCuq32260

  Symptom: After ISSU, multiple ipfib process crashes occur.

  Conditions: Occurs in the MFIB when multicast is enabled during the ISSU.

  Workaround: This issue is resolved.

- CSCuq39731

  Symptom: Error messages of %IPFIB-SLOT6-4-FLN_FIB_ADJ_EXHAUSTED: Adjacency allocation failed on instance X

  Conditions: Occurs when there are a high number of topology changes or mrib changes.

  Workaround: This issue is resolved.

- CSCuq41416

  Symptom: High CPU utilization on F3 module.

  Conditions: Occurs when there are spurious MAC-learn notifications.

  Workaround: This issue is resolved.

- CSCuq70531

  Symptom: Non-fabric capable ports added to existing FEX port-channel.

  Conditions: If the above operation is tried twice in quickly. It results into some configuration commands being locked; which has resulted in all the later switchport command failure.

  Workaround: This issue is resolved.

- CSCuq73614

  Symptom: Netstack crashes while making IP route change with track object.

  Conditions: Occurs when the track object is part of the route change and the track object is in flapping/change state.

  Workaround: This issue is resolved.

- CSCuq79790

  Symptom: A route-map that matches a tag for an eigrp internal route does not work because the tag is not present on the route, although the tag is present in the topology entry.

  Conditions: Occurs for only for internal routes. External routes carry the tag and are not affected.

  Workaround: This issue is resolved.

- CSCuq79862

  Symptom: AM route for a moved away dynamic host exists along with the lisp-dyn eid Null0 route with the same route-preference.

  Conditions: Occurs in an ESM single-hop mobility case on an away xTR. lisp creates a Null0 route for an away host discovered on another datacenter and needs it to win in order to trigger SMRs. However, AM installs a route for the same /32 host pointing to the interface extended to another data-center. As a result, the traffic destined for that dynamic host will be tromboning over the dci link to the other data-center since the map-cache on the remote datacenter is not updated.

  Workaround: This issue is resolved.

- CSCuq81249

  Symptom: The Control Plane Policing (CoPP) process might unexpectedly crash and restart during initialization.

  Conditions: Occurs when PPF sends two responses for a single statistics request sent by the CoPP manager.

  Workaround: This issue is resolved.

- CSCuq81826

  Symptom: High latency on traffic going through and F3 module.

  Conditions: Occurs when there is an F3 module, ISSU, and an incorrect COS-to-Queue mapping.

  Workaround: This issue is resolved.

- CSCuq81832

  Symptom: Modules might experience l2mcast process crashes.

  Conditions: Occurs with v2 joins (*,G) and v3 joins (S,G) with OMF disabled on the fabric path VLAN 182.

  Workaround: This issue is resolved.

- CSCuq81871

  Symptom: Multicast traffic inside VRF is not forwarded over MDT group to remote PE.

  Conditions: Happens only for MVPN traffic in the situation where S,G pin Join is sent to the VPC secondary and there are no local receivers and multicast traffic is received over an orphan layer 3 interface.

  Workaround: This issue is resolved.

- CSCuq86293

  Symptom: The IGMP ORIB buffer might go down to zero.

  Conditions:

  Workaround: This issue is resolved.

- CSCuq79793

  Symptom: IPv6 neighbor discovery processes a neighbor advertisement with Link Layer address 0000.0000.0000 as valid.

  Conditions: The Cisco Nexus 7000 Series switch receives an IPv6 neighbor advertisement from a host with MAC address 0000.0000.0000 in the Link Layer address field.

  Workaround: This issue is resolved.

- CSCuq78160

  Symptom: Host machines are marking the HSRPv6 gateway as unreachable.

Conditions: The host sends unicast NS probes with a Link Local source address to the HSRPv6 IP address in a vPC or vPC+ scenario. These packets reach HSRP standby and need to go over the peer link to HSRP active.

Workaround: This issue is resolved.

- CSCuq77945

  Symptom: Syslog messages are truncated on "%AUTH" or "%AU."

  Conditions: You have entered the **logging timestamp microseconds** or **logging timestamp milliseconds** command.

  Workaround: This issue is resolved.

- CSCuq65354

  Symptom: In Cisco NX-OS Release 6.2.2a, a VDC with an F2e line card sends option templates with a sampling rate of configured rate * 100. Beginning with Cisco NX-OS Release 6.2.6, a VDC with an F2e line card is considered part of a mixed VDC, and the sampling rate is sent as configured to maintain consistent behavior for all the line cards in the VDC. As a result, an issue occurs while plotting the graph for Cisco NX-OS Release 6.2.8.

  Conditions: The graph is being plotted for the byte count for Cisco NX-OS Releases 6.2.6 and 6.2.8.

  Workaround: This issue is resolved.

- CSCuq54506

  Symptom: A host moves across an overlay and is not reachable. The devices in the original location (except the OTV AED) still have the MAC address pointing to the old location.

  Conditions: The Cisco Nexus 7000 Series switch is configured for OTV, and RARPs are lost.

  Workaround: This issue is resolved.

- CSCuq46582

  Symptom: The Cisco Nexus 7000 Series switch might fail to send an ARP reply on the affected interface after the same HSRP VIP configuration is moved to another interface.

  Conditions: The initial VIP config is left under the interface, and the same VIP is reused under another interface.

  Workaround: This issue is resolved.

- CSCuq40808

  Symptom: The Fabric Extender host interface (HIF) ports might not link up when connected with HP NC365T.

  Conditions: The Cisco Nexus 2248TP-E is connected as a Fabric Extender to a Cisco Nexus 7000 Series switch running Cisco NX-OS Release 6.1.4a.

  Workaround: This issue is resolved.

- CSCuq40487

Symptom: The multicast RPF could fail due to recursive route pointing to shut the interface.

Conditions: The default route is used for RPF, and the **show routing recursive-next-hop** command shows an incorrect next hop for the default route.

Workaround: This issue is resolved.

- CSCuq26934

  Symptom: One or more member ports on a port channel will still have the default queuing configuration when a custom queuing policy is applied on a port channel.

  Conditions: The default queuing configuration is applied to a port channel whose member interfaces belong to a module that was previously reloaded.

  Workaround: This issue is resolved.

- CSCuq25337

  Symptom: Broadcast traffic may not egress out from port channel members and/or physical ports.

  Conditions: Multiple ISSUs cause an incorrect setting of the port bitmaps as a result of PSS key corruption.

  Workaround: This issue is resolved.

- CSCuq13845

  Symptom: SNMP returns a "tooBig(1)" error instead of truncating an encapsulated message to equal or less than maximum message size.

  Conditions: Any of these methods or a combination of these methods can cause the SNMP "tooBig" error:

  - Changing the default setting for the **snmp-server packetize** command, especially when choosing smaller values
  - Using SNMP tool commands like the **-Cr** option that request large amounts of data to be returned
  - Using large ifmib ifalias descriptions and having many ports with extremely long descriptions

  Workaround: This issue is resolved.

- CSCuq05409

  Symptom: A high number of messages cause a low memory situation in MTS for the internal SPM process (SAP 487). This results in a crash of SPM and a switchover.

  Conditions: While collecting tac-pac when PPF statistics are taken, the PPF goes into an infinite loop that causes SPM to stop processing messages. This leads to an overload of MTS buffers and the consequent restart of SPM.

  Workaround: This issue is resolved

- CSCuq05465

  Symptom: When you run the **show system internal sdwrap buffers sap 885** command, one MTS message gets stuck in the MTS buffer on the line card and is not cleared.

Conditions: You run the **show system internal sdwrap buffers sap 885** command on a device with the M2 line card.

Workaround: This issue is resolved.

- CSCum91206

  Symptom: A TACACS+ accounting error can generate an "All servers failed to respond" syslog message.

  Conditions: A TACACS+ error is generated when you enter a large and complicated command, such as the following:

  ```
  echo po21,po22,po23,po24 | tr ',' '\n' | sed 's/^/show port-channel database interface
  /' | vsh | grep "Ethernet.*/.*up" | sed 's/^.*Ethe/show cdp neigh int Ethe/'| sed
  's/\[.*$//' | vsh | egrep "name" | cut -d '-' -f 1 | sort -u | egrep -c "abc123" |
  egrep -c "^1$" 0
  ```

  Workaround: This issue is resolved.

- CSCul08762

  Symptom: Snmpwalk of ospfNbrTable for a neighbor fails when the neighbor undergoes a router ID change.

  Conditions: OSPFv2 is configured, and you change the router ID of neighbor R2 (for example, from RTRID-R2 to RTDID-R2) either explicitly or implicitly (if inheriting from the IP address of the loopback or physical interface). Then you perform an snmpwalk of neighbor R2 on DUT (R1).

  Workaround: This issue is resolved.

- CSCuq60239

  Symptom: When the IPv6 host route is removed after an ISSU (for Cisco NX-OS Releases up to 6.2.8), the IPv6 adjacency does not get resolved because the subnet adjacency is misprogrammed, and the packet does not reach the supervisor to trigger an IPv6 neighbor discovery packet.

  Conditions: The IPv6 glean fast path feature is enabled and supported in the hardware, and you upgrade to any Cisco NX-OS software release up to 6.2.8.

  Workaround: This issue is resolved.

- CSCuq55140

  Symptom: After an ISSU or an l2mcast process kill on a device with F Series line cards, Layer 3 multicast traffic to FEX Layer 2 receivers will fail when the HIF interface flaps.

  Conditions: A post PSS restore of the l2mcast process occurs on F Series line cards.

  Workaround: This issue is resolved.

- CSCuq55039

  Symptom: A crash occurs in orib_show_internal_snmp_route.

  Conditions: The otv-isis process is stopped and restarted using the command line interface while the OTV MIB (1.3.6.1.4.1.9.9.810) is being polled.

  Workaround: This issue is resolved.

- CSCuq46076

  Symptom: The egress_vsl_drop for F2 Series line cards does not work as expected.

  Conditions: You convert vPC+ to vPC in an F2 VDC.

  Workaround: This issue is resolved.

- CSCup23990

  Symptom: The NetFlow configuration is not displayed in the output of the **show startup-config** command.

  Conditions: This issue might be seen if your switch is running Release 6.2(x) and you configured NetFlow using one username and issued the show command using a different username. Multiple usernames can access the **copy running-config startup-config** command. Entering the configuration using one username and saving it using a second username can cause the first username to not see the Netflow configuration. Everything except the feature disappears.

  Workaround: This issue is resolved.

- CSCup79010

  Symptom: The line card rehosting is not triggered when you change the VLAN configuration for the interface on a hosting line card.

  Conditions: This issue might be seen if the **switchport monitor** command was previously configured on that interface.

  Workaround: This issue is resolved.

- CSCup93710

  Symptom: LACP link flap occurs on multiple ports with lacp rate fast.

  Conditions: This issue might be seen if your Cisco Nexus 7018 switch is running Release 6.1(2E5) and you have a Sup1 and F2 modules,

  Workaround: This issue is resolved.

- CSCup98159

  Symptom: The MSDP process crashes after you remove the MSDP peer configuration.

  Conditions: This issue might be seen if Release 6.1(2) S42 is running on the Cisco Nexus 7000 Series switch and you configure the **no ip msdp peer** command.

  Workaround: This issue is resolved.

- CSCuq81960

  Symptom: The output of the **show hardware capacity forward** command for Single-Width Entries, Double-Width, and Quad-Width total resource does not add up with cshcProtocolFibTcamTotal mibwalk.

  Conditions: This issue might be seen during normal operations.

  Workaround: This issue is resolved.

- CSCuq27744

  Symptom: The output of the **show hardware internal forwarding table utilization** command is displaying misleading TCAM utilization percentages. For example:

  ```
  switch# show hardware internal forwarding table utilization

  Entry                       Logical         Physical        %Total
  Type                        Entries         Entries          TCAM Capacity
  ---------------------------------------------------------------------------
  Single-Width Entries:       99810           99810           76
    IPv4 Unicast              99810           99810           76
    L2VPN Peer                0               0               0
    MPLS                      0               0               0

  Dual-Width Entries:        15621           31242           11<==This should be ~22%
    IPv4 Multicast            4               8               0
    L2VPN IPv4 Mcast          0               0               0
    IPv6 Unicast              15617           31234           11

  Quad-Width Entries:         5               20              0
    L2VPN IPv6 Mcast          0               0               0
    IPv6 Multicast            5               20              0

  Free Single-Width Entries:  0               0               0
  ```

  Conditions: This issue might be seen because the TCAM capacity is not being calculated against the Physical Entries column. Instead it is being calculated against the not Logical Entries column.

  Workaround: This issue is resolved.

- CSCup22563

  Symptom: The following Cisco products are affected by the vulnerabilities identified by the Common Vulnerability and Exposures (CVE) IDs running Cisco NX-OS Release 7.1(0), NX-OS Release6.2(8), NX-OS Release6.2(7), and NX-OS Release5.2(8d):

  - Cisco Nexus 7000 Series switches
  - Cisco MDS 9000 Series switches

  Conditions: The symptom might be seen on these devices when the device is running LDAP in SSL mode and the following command is configured:

  ldap-server host *ipv4-address ipv6-address host-name* **enable-ssl**

  Workaround: This issue is resolved.

- CSCuo68846

  Symptom: An ACLQoS crash is seen on the line card due toa FU_MEM_hashtable_node_array_t continuous alloc error. The following errors are logged on the switch:

  - `%SYSMGR-SLOT10-2-SERVICE_CRASHED: Service "aclqos" (PID 12914) hasn't caughtsignal 6 (core will be saved).`

  - %SYSMGR-SLOT10-5-SUBPROC_TERMINATED: "System Manager (core-client)" (PID 17560) has finished with error code SYSMGR_EXITCODE_CORE_CLIENT_ERR (11).

  Conditions: This issue might be seen after 2 to 3 days of periodic statistics polling. If you monitor memory usage on line cards after the crash, the following object will show continuous allocations: FU_MEM_hashtable_node_array_t continuous.

  Workaround: This issue is resolved.

- CSCuo67870

  Symptom: The vPC VLANs on a trunk interface are incorrectly pruned by VTP on the vPC Operational Secondary Cisco Nexus 7000 Series switch. The interface can be configured as a vPC leg or as a vPC Orphan port.

  Conditions: This issue might be seen on a Cisco Nexus Series 7000 switch vPC setup running Release 6.2(6) or a later release and the switch is also configured to use VTP for VLAN pruning.

  The trigger for the issue is a flap of a vPC leg on the vPC Operational Primary switch. After the link flap, VTP does not successfully process join messages for the VLANs causing a timeout and the VLANs to be pruned from the link.

  Workaround: This issue is resolved.

- CSCuq03603

  Symptom: The output of the **show system internal aclmgr log** command in the standby supervisor displays an aclmgr leak in the command library. The size of leak is 196 bytes. per show command

  Conditions: This issue might be seen if you issue the **show system internal aclmgr log** command in the standby supervisor.

  Workaround: This issue is resolved.

- CSCup96876

  Symptom: Traffic to a subset of hosts is punted to the CPU and either forwarded in software or dropped.

  Conditions: This issue might be seen if you perform an ISSU from Release 6.2(8) to Release 6.2(8a) or from an earlier release to Release 6.2(8) or Release 6.2(8a) and you have M1-XL or M2 modules installed. This issue does not apply to F2 or F3 modules.

  This issue occurs because the FIB entry is programmed in the wrong SPL bank. Consequently, the longest prefix match is not hit and traffic is punted for an ARP glean.

  For example, a host exists in a /24 subnet. In the following output from the **show ip route** command to the host shows that a /32 installed is in the RIB. However, checking the longest prefix match in hardware will show the /24 route.

```
N7K-1# show system internal forwarding ipv4 route 172.22.31.173 module 7

Routes for table default/base

----+--------------------+----------+----------+------+----------

Dev | Prefix             | PfxIndex | AdjIndex | LIFB | LIF

----+--------------------+----------+----------+------+----------
```

```
1    172.22.31.0/24        0x423e     0xa2cb       0   0x1ff8b  <------hardware only
finds the /24 entry
```

Note The /32 entry is present in hardware but since it is programmed in the wrong bank it will always hit the /24 entry.

Workaround: This issue is resolved.

- CSCun60221

  Symptom: Aclqos crashes and the module goes to failure state.

  Conditions: This issue might be seen in the following scenarios:

  –A multi-instance card such as an F2 or F3 module, is installed and any redirection feature is configured on multiple instances.

  –There is not enough space in the TCAM on one of the instances where the policy will get programmed as well.

  –The atomic update goes through on one instance but fails on next.

  –The non atomic update is turned on and fails on the next instance.

  –During recovery path, aclqos crashes

  Workaround: This issue is resolved.

- CSCuq06354

  Symptom: Both of the ports on the Cisco Catalyst switch and Cisco Nexus switch stay down after the WS-C3750X-48T-S and N7K-F248XT-25E modules are connected and the F2 module is disconnected.

  Conditions: This issue might be seen if you connect the WS-C3750X-48T-S and N7K-F248XT-25E modules and then you disconnect or reconnect the F2 por.

  Workaround: This issue is resolved.

- CSCup95423

  Symptom: Cisco Nexus 7000 Series switches running Cisco NX-OS Release 6.2(6a) software with BFD causes BFD flaps for several minutes before staying up.

  Conditions: This issue is seen during BFD initialization when both end points send their respective “P” packets at the same time.T

  Workaround: This issue is resolved.

- CSCup95403

  Symptom: The following error message is seen when 64 FEX VDC ports are shutdown at the same time without performing a no-shut:

```
2014 Jul 17 10:46:03 N7k-F2 %$ VDC-2 %$ %VNTAG_MGR-2-VNTAG_SEQ_ERROR: Error ("sequence
timeout") while communicating with component MTS_SAP_HP_IFTMC for Opcode
MTS_OPC_VNTAG_ELTMC_SET_VLAN_CBL
```

Conditions: This symptom was seen when shutting down all 64 FEX FPC ports without doing a no-shut.

Workaround: This issue is resolved.

- CSCul79456

  Symptom: The DHCPv6 relay using source-int not working. The UDP6 stats displays "NO PORT" for the F3 module.

  Conditions: This issue might be seen if the **ipv6 dhcp relay source-interface loopback20** command is configured in global configuration mode and the client-side SVI does not have an IPv6 address.

  Workaround: This issue is resolved.

- CSCup70414

  Symptom: F3 Series GRE tunnels fail on a sixth VRF due to a hardware programming failure.

  Conditions: This symptom will be seen if you are creating more than 5 VRFs in a VDC.

  Workaround: This issue is resolved.

- CSCup93146

  Symptom: After applying a QoS configuration that moves all traffic to a default queue on the system admin-vdc, the ipqosmgr process fails.

  Conditions: This symptom is seen when applying a QoS configuration the moves all traffic to default queues.

  Workaround: This issue is resolved.

- CSCup65296

  Symptom: After perturbing FabricPath VLANs in the FP network, the interface VLAN component experiences a memory leak.

  Conditions: The VLAN perturbance involves VLAN addition/deletion, VLAN shut/no shut, VLAN mode change. Perturb the FabricPath VLANs, the memory starts to leak.

  Workaround: This issue is resolved.

- CSCup65294

  Symptom: After perturbing FabricPath VLANs in the FabricPath network, the VLAN leaks memory.

  Conditions: This issue might be seen if the VLAN perturbance involves adding or deleting a VLAN, shutting or no shutting the VLAN, or change the mode of the VLAN mode change.

  Workaround: This issue is resolved.

- CSCuq07308

  Symptom: The following error message is displayed when you use the **restart bgp** command to restart the Border Gateway Protocol (BGP):

```
Service "bgp" (PID 9158) hasn't caught signal 6 (core will be saved).
```

Conditions: This issue might be seen if BGP is configured with more than 1900 neighbors or peers.

Workaround: This issue is resolved.

- CSCuq24757

  Symptom: VTP resets.

  Conditions: This issue might be seen when a failed module is resetting.

  Workaround: This issue is resolved.

- CSCuq16692

  Symptom: The ACLQOS process crashes on all M series modules in the chassis.

  Conditions: This symptom might be seen if you perform an ISSU from Release 5.2(9) and to Release 6.2(x) without reloading the Cisco Nexus 7000 Series switch.

  Workaround: This issue is resolved.

- CSCup60557

  Symptom: A Cisco Nexus 7000 Series switch does not generate an ICMP unreachable message for packets with an IP length that is from 1501 bytes to 1516 bytes, with the DF bit set for path MTU discovery if the following is true:

  - Routing is between interfaces with F2/F2e modules where the ingress interface has an MTU of 9216 and the egress interface has an MTU of 1500.
  - No 802.1Q trunking or FabricPath is configured.
  - ICMP unreachable messaging is enabled on the ingress interface.

  Conditions: This issue might be seen when PMTU discovery is required between networks where jumbo MTU support is also required but the connecting device is unable to accept packets with an IP length that is greater than 1500 bytes.

  Workaround: This issue is resolved.

- CSCup55118

  Symptom: ORIB buffer exhaustion occurs when the system receives continuous IGMP join/leave messages.

```
Switch# show ip igmp snooping internal ribs
IGMP Snooping internal RIB information
  RIB name: M2RIB (type 0), ready: Yes No , xid 0x1f6cd
    Max. outstanding buffers: 4
    Current outstanding buffers: 0
    Max. OMF entries per buffer: 400
    Max. OMF route entries per buffer: 50
    Max. route entries per buffer: 400
    Fabricpath redist instance: 1
  Used buffer queue count: 0
  Free buffer queue count: 10
    Buffer: 0x0x98650ac, type: none, xid: 0x0, count: 0
    Buffer: 0x0x9869f04, type: none, xid: 0x0, count: 0
    Buffer: 0x0x986ed5c, type: none, xid: 0x0, count: 0
    Buffer: 0x0x9873bb4, type: none, xid: 0x0, count: 0
    Buffer: 0x0x985174c, type: none, xid: 0x0, count: 0
    Buffer: 0x0x984c8f4, type: none, xid: 0x0, count: 0
    Buffer: 0x0x9847a9c, type: none, xid: 0x0, count: 0
    Buffer: 0x0x98565a4, type: none, xid: 0x0, count: 0
    Buffer: 0x0x985b3fc, type: none, xid: 0x0, count: 0
    Buffer: 0x0x9860254, type: none, xid: 0x0, count: 0
  TXLIST member version: 0
RIB name: ORIB (type 1), ready: Yes No , xid 0x10009
  Used buffer queue count: 10
    Buffer: 0x0x97ff7ac, type: route, xid: 0x10000, count: 1
    Buffer: 0x0x9804604, type: route, xid: 0x10001, count: 1
    Buffer: 0x0x980945c, type: route, xid: 0x10002, count: 1
    Buffer: 0x0x980e2b4, type: route, xid: 0x10003, count: 1
    Buffer: 0x0x981310c, type: route, xid: 0x10004, count: 1
    Buffer: 0x0x97e6ff4, type: route, xid: 0x10005, count: 1
    Buffer: 0x0x97ebe4c, type: route, xid: 0x10006, count: 1
    Buffer: 0x0x97f0ca4, type: route, xid: 0x10007, count: 1
    Buffer: 0x0x97f5afc, type: route, xid: 0x10008, count: 1
 Buffer: 0x0x97fa954, type: route, xid: 0x10009, count: 1
  Free buffer queue count: 0
  TXLIST member version: 69545
```

Conditions: This issue might be seen in an OTV environment.

Workaround: This issue is resolved.

- CSCuo99830

  Symptom: The port-client core, ISSU failure, ports are suspended.

  Conditions: This issue might be seen during an ISSU from Release 6.1() to Release 6.2(8a). During the ISSU, the line card is upgraded once the SUP switchover is completed and the EthPm sends port commands to the F2module. Just before this occurs, if there are any process crashes, such as a vPC crash or an ungraceful VDC reload, the process crashes when the port commands from EthPm are processed by the port client. The crash might occur because an invalid pay-load size is sent to the port client.

  Workaround: This issue is resolved.

- CSCuo85450

  Symptom: The switch reloads and the following log is displayed in the output of the s**how system reset-reason** command:

```
----- reset reason for Supervisor-module 1 (from Supervisor in slot 1) ---
1) At 141764 usecs after Wed May 14 03:21:41 2014
    Reason: Reset triggered due to HA policy of Reset
    Service: arp hap reset
    Version: 7.0(1)N1(1)
```

Conditions: This issue might be seen when there are gratuitous ARP storms.

Workaround: This issue is resolved.

- CSCun26418

  Symptom: The replication ASIC that is responsible for SPAN, OTV encapsulation, and multicast replication stops replicating all traffic and each of these features also fail. This is particular to the SPAN/multicast replication pipeline so customer in an OTV environment will see that multicast traffic (including link-local multicast 224.0.0.0/24), along with broadcast traffic are not forwarded across the overlay.

  Conditions: This issue is might be seen in an OTV environment with a high rate of traffic egressing the OTV join-interface, with ERSPAN configured to span the join-interface. This is not typical.

  Workaround: This issue is resolved.

- CSCuo50598

  Symptom: A Cisco Nexus 7000 Series switch with an F3 module experiences link flaps on port-channels when a VDC is reloaded.

  Conditions: This issue might be seen if there is a port-channel configuration on an F3 module with active traffic.

  Workaround: This issue is resolved.

- CSCuo19840

  Symptom: Anycast HSRP is in an initial state after a supervisor failover.

  Conditions: This issue might be seen during a supervisor switchover and the output of the **show hsrp anycast** command is as follows:

```
Switch# sh hsrp anycast

Anycast bundle - 1 (IPv4)
  Admin Status: Up  Oper Status: Down
              Reason: Invalid switch-id Cfged, :
  Anycast Switch ID XXX

 %HSRP_ENGINE-3-BUNDLE_ASID_REG_FAIL: Switch-id:XXX DRAP registration failed for
bundle:1.the
```

  Workaround: This issue is resolved.

- CSCuo13444

  Symptom: Traffic breaks for all sub-interfaces under a parent interface when a subinterface under that parent interface is removed.

  Conditions: This issue might be seen if the subinterfaces configured on an F2 module.

  Workaround: This issue is resolved.

- CSCun78718

  Symptom: An incorrect rollback of a failed Port-Profile modification might be seen in some very specific cases on Cisco Nexus 7000 Series switches. As a result of this incorrect PPM rollback, a significant part of the running configuration might be negated on the device. Reloading the device recovers from the issue because the device is re-initialized from its startup configuration.

  Conditions: This is a scalability-related issue. The problem might be seen on heavily over-configured Cisco Nexus 7000 Series switches with many attached Fabric Extenders (FEXs) (~1,500 interfaces on the system, ~10K lines in the output from the **show run all** command, ~200 interfaces managed by a single Port-Profile). Under these specific conditions some port-profile operations can take several minutes to be completed or might be rejected.

  This problem might be triggered if two different port-profiles are being modified simultaneously from two different management sessions to the switch (VDC), and both of the two PPM modifications are failed. In this very specific case, two PPM rollback operations might be initiated simultaneously, breaking the normal PPM rollback operation and causing an incorrect completion of the PPM rollback operation.

  Workaround: This issue is resolved.

- CSCuo76751

  Symptom: vPC+ underlying links may remain in the suspended state indefinitely.

  Conditions: This issue was observed after reloading a VDC. When the VDC reload was complete, the links in the port channel remained in the suspended state.

  Workaround: This issue is resolved.

- CSCuq39832

  Symptom: A Cisco Nexus 7000 Series switch running a Cisco NX-OS Release 6.2.x prior to 6.2(10) may crash and produce core files in the Resource Manager Daemon.

  Conditions: With a relatively low likelihood, a buffer overflow might be triggered when you issue one of the following commands:

  - **show resource**
  - **show resource vlan** (or another resource name)
  - **show vdc resource detail**
  - **show vdc resource vlan** (or another resource name) **detail**
  - **show vdc vdc-name resource**
  - **show vdc vdc-name resource vlan** (or another resource name)

  When you frequently allocate or delete VLANs across multiple VDCs, the switch is more susceptible to this issue.

  Workaround: This issue is resolved.

- CSCuq03563

  Symptom: Multicast traffic is not getting forwarding over the non vPC vlan, source VLAN, and destination VLAN that are connected from the switch.

Conditions: The issue might be seen if the source and destination VLANs are not vPC VLANs.

Workaround: This issue is resolved.

- CSCuo63932

  Symptom: The following errror message is logged:

  ```
  2014 Apr 29 18:11:26 nxsfasc6a1-gw dcos-ping[12674]: IP-3-IP_FAILURE: Failed to
  Failed call urib_api_init()
  ```

  Conditions: This issue might be seen if you used the extended **ping** command from the Command Line Interface (CLI) of the your Cisco Nexus 7000 Series switch.

  Workaround: This issue is resolved.

- CSCuo71901

  Symptom: An interface might move to the error-disabled state with egress QoS policies applied.

  Conditions: This symptom might be seen when you reload several modules at once so that a large number of interfaces (more than 11 shots with 48 ports) are UP.

  Workaround: This issue is resolved.

- CSCuo22348

  Symptom: You might see one of the Layer 3 protocols flap when you enter the **show interface trunk** command.

  Conditions: This symptom might be seen when the system has multiple VDCs with M1 and F1 Series modules in the same VDC. With these conditions, entering this command can generate excessive traffic in the Tx direction from the CPU and may drop certain packets, causing Layer 3 instability.

  Workaround: This issue is resolved.

- CSCum61205

  Symptom: Broadcast and link local multicast traffic might be lost across the OTV circuit during OTV packet decapsulation because of a missing label in hardware.

  Conditions: This symptom might be seen when the overlay is up and extending to at least one VLAN, then you enter the **layer-2 multicast lookup mac** command on an extended VLAN, and the overlay goes up and down.

  Workaround: This issue is resolved.

- CSCuo69414

  Symptom: When you perform an ISSU from Cisco NX-OS Release 6.1(4a) to Release 6.2(6), the ALCMGR might go down.

  Conditions: This symptom might be seen when you have applied an IPv6 PACL on a port channel running Cisco NX-OS Release 6.1(4a), you remove a member of the port channel, and you enter the **ipv6 traffic-filter** *name* **in** command, and you add the port back to the port channel. Then, you perform an ISSU to Release 6.2(6).

  Workaround: This issue is resolved.

- CSCuo63254

  Symptom: After you reload an F3 Series module, multicast traffic might be dropped on the line card**.**

  Conditions: This symptom might be seen when you are working with an F3 Series module with the transit mode enabled and distributed channels are configured.

  Workaround: This issue is resolved.

- CSCuo51846

  Symptom: If you have upgraded to Cisco NX-OS Release 6.2(8) and enter the **service unsupported-transceiver** command to enable third-party transceiver modules, you might see these modules fail.

  Conditions: This symptom might be seen when you are using an F3 Series module in Cisco Nexus 7700 Series chassis and running Cisco NX-OS Release 6.2(8).

  Workaround: This issue is resolved.

- CSCuo71910

  Symptom: You cannot apply a QoS policy on Layer 2 ports on the egress direction. Attempts to do so may not be rejected properly.

  Conditions: This symptom might be seen when you attempt to apply a QoS policy on a Layer 2 port.

  Workaround: This issue is resolved.

- CSCuo63287

  Symptom: A packet might leak from one VDC to another if every interface on an M1 Series module is allocated to another VDC.

  Conditions: This symptom might be seen when you have a module in the admin VDC with some port channels configured and another module, which is not always used, in the admin VDC. Then if you simultaneously allocate every interface from the second module to another VDC, you might see leaking.

  Workaround: This issue is resolved.

- CSCuo64110

  Symptom: Some VLANs might not display in the running configuration, even though these VLANs are working correctly.

  Conditions: This symptom might be seen when you are working with some non-FabricPath VLANs that have not been configured with a name.

  Workaround: This issue is resolved.

- CSCup02927

  Symptom: The system may go down after you reload a module when you have more than 8 port channels in a system with regular or PVLAN mapping and VLAN translations on the port channels.

Conditions: This symptom might be seen when you reload the module when you have more than 8 port channels in a system with regular or PVLAN mapping and VLAN translations on the port channels.

Workaround: This issue is resolved.

- CSCuo78535

  Symptom: Some modules may power reset, and then be unable to start or send the redundancy timer**.**

  Conditions: This symptom might be seen when you are performing scale steady state traffic testing.

  Workaround: This issue is resolved.

- CSCuj22681

  Symptom: The default value for maximum-paths does not display in the output from the **show run bgp all** command**.**

  Conditions: The display for the **show run bgp all** command displays default maximum-paths value for different supported address-family groups when maximum-paths is not explicitly configured for all VRFs.

  Workaround: This issue is resolved.

- CSCun82279

  Symptom: The graceful restart for the Label Discovery Protocol (LDP) might take a few minutes to recover after a supervisor switchover. The switch returns the following error:

  ```
  2014 Mar 11 06:39:55.990 test-pe LDP-5-GR    GR session 192.168.1.1:0 (inst 4):
  interrupted--recovery pending
  ```

  Conditions: This symptom might be seen when the supervisor switches over after an HA or ISSU procedure.

  Workaround: This issue is resolved.

- CSCul91339

  Symptom: You might see traffic dropping when you perform back-to-back MAC address moves between two different sites.

  Conditions: This symptom might be seen when you are moving the local MAC addresses back-to-back across two sites, which results in nodes pointing to an incorrect route owner.

  Workaround: This issue is resolved.

- CSCuo86477

  Symptom: The LISP service may go down because of MTS buffer leaks in SAP associated with the LISP UFDM MTS queue.

  Conditions: This symptom might be seen when LISP is configured on the switch.

  Workaround: This issue is resolved.

- CSCup11341

Symptom: You might see the module fail after you power cycle the module, with the following message:

```
%VMM-2-VMM_SERVICE_ERR: VDC1: Service SAP Spm SAP for slot 3 returned error 0x41040001
(aclqos failure) in if_bind sequence
%IM-3-IM_RESP_ERROR: Component MTS_SAP_VMM opcode:MTS_OPC_IM_IF_VDC_BIND in vdc:1
returned error:aclqos failure
```

Conditions: This symptom might be seen when you power cycle a module.

Workaround: This issue is resolved.

- CSCup18687

  Symptom: On Cisco Nexus 7000 Series switches, IP to DGT mappings may not be honored following a module reload.

  Condition: This symptom may be seen after reloading a module.

  Workaround: This issue is resolved.

- CSCuj74494

  Symptom: The **show hardware queuing drops** command does not display the correct output.

  Conditions: This symptom might be seen with F3 Series modules.

  Workaround: This issue is resolved.

- CSCuo10029

  Symptom: When a vPC+ peer reloads in a mixed chassis scenario [M and F1/M and F2 Series modules], routed traffic ingressing on an M Series module will blackhole because the DI does not drive the peer switch switch ID, and does drive the ES switch ID. This DMAC was learned on the other peer and was not installed correctly on this peer upon reload.

  Conditions: This symptom might be seen when you are working with vPC+ in a mixed chassis with routed east-west traffic.

  Workaround: This issue is resolved.

- CSCul79472

  Symptom: The private VLAN port moves to Inactive on private VLAN promiscuous when you add an association in a non-PVLAN mode (access mode) and then change the port mode to VLAN promiscuous mode.

  Conditions: This symptom might be seen under the following conditions:

  - The port is up in PVLAN promiscuous mode with promiscuous mapping configured.
  - You delete both the VLANs from PVLAN association.
  - You change port mode to access mode by entering the **switchport mode access** command.
  - You re-create the PVLAN association between the primary and secondary VLANs.
  - You change the mode back by entering the **switchport mode private-vlan promiscuous** command.

  Workaround: This issue is resolved.

- CSCuo25489

  Symptom: You might not see the private VLAN list in the display after you reload the switch by entering the **copy r s** command, and then entering the **show vPC consistency-parameters vPC** command.

  Conditions: This symptom might be seen when you reload the switch. The vPC comes up and is formed correctly. But the consistency list does not show the private VLAN list, which consists of the private VLAN pairs for the current vPC leg port mode.

  Workaround: This issue is resolved.

- CSCun76395

  Symptom: TACACS authorization request packets are not created by the switch, so all TACACS-based authentications fail.

  Conditions: This symptom might be seen when you are working with Cisco NX-OS Release 6.(x) code when you load a nondefault VDC.

  Workaround: This issue is resolved.

- CSCuc32563

  Symptom: After reloading a line card, a silent multicast loss (no syslog) might occur. If this happens, the impacted multicast traffic will experience a complete, or near complete, packet loss. This is a result of the multicast distribution (MD) destination index (DI) not being programmed correctly.

  Conditions: This symptom occurs following line card reloads.

  Workaround: This issue is resolved.

- CSCuo80937

  Symptom: Spanning Tree Protocol (STP) topology changes and Bridge Protocol Data Units (BPDUs) are sent every 2 seconds for a long period of time after approximately 100 days of active supervisor uptime.

  Conditions: You might see this symptom if there are topology changes after you upgrade to Cisco NX-OS Release 6.2(6), 6.2(6a), or 6.2(8) on your Cisco Nexus 7000 or 7700 Series switches.

  Workaround: This issue is resolved.

- CSCua79354

  Symptom: An overlay flap immediately after a switchover can result in an OTV adjacency down.

  Conditions: This symptom might be seen during a supervisor switchover if a low memory situation occurs because of considerable changes in network. A loss of messages can lead to this situation.

  Workaround: This issue is resolved.

- CSCuo62938

Symptom:: Networks with Cisco Nexus 7000 Series gateway devices running Cisco NX-OS Releases 6.2(2) and 6.2(10) using F1 with M1 Series modules with FabricPath may have other nodes experience persistent MAC moves. The continuous MAC moves cause high CPU usage, which can lead to other problems for the control and data plane. The gateway MACs will change their association between the local switchID.0.SUP Lid on the gateway device and the local switchID.0.ffff.

Conditions: The symptom criteria is as follows:

- Cisco Nexus 7000 Series gateway devices must be using F1 with M1 Series modules with FabricPath.
- Cisco Nexus 7000 Series leaf devices must be using F1 or F1 with M1 Series modules.
- Gateway devices must be running NX-OS Release 6.2(2) and NX-OS Release 6.2(10).

Workaround: This issue is resolved.

- CSCuo68922

  Symptom: The **system vlan long-name** command cannot be configured.

  Conditions: This issue was introduced in NX-OS Release 6,2(8).

  Workaround: This issue is resolved.

- CSCun69113

  Symptom: The "ifmgr" process fails when unbinding an invalid interface from VDC.

  ??????????: Line cards may not be reporting the right number of ports on a card. For example, a 32-port card reports that it has 48 ports, as seen in the output of the **show vdc internal port-hash** command. The non-existent ports may be bound to VDC1 by default.

```
show vdc internal port-hash
Interface   If_index      Vdc_id
Eth1/33     XXXXXXXX      1. . .
Eth1/48     XXXXXXXX      1
```

  If the above output is seen and slot 1 is a 32-port card, this is not correct. The output should show a 48-port card.

  Workaround: This issue is resolved.

- CSCun69580

  Symptom: MPLS TE tunnel remains open for 30-40 seconds after link is shut down.

  Conditions: This symptom might be seen when you are working with MPLS TE tunnels.

  Workaround: This issue is resolved.

- CSCto65106

  Symptom: A connectivity loss for 5 to 10 seconds occurs when a vPC peer-link is brought back online.

  Conditions: This symptom may occur in one of the following situations:

  - When a vPC peer-link is brought down, the Cisco Nexus 7000 Series primary device continues forwarding traffic on all vPCs.

When a vPC peer-link is brought back up, the secondary Cisco Nexus 7000 Series device may experience a connectivity loss when traffic is present.

This issue may be seen when a large number of vPCs are configured.

Workaround: This issue is resolved.

- CSCuo71648

  Symptom: Traffic that ingresses on an F3 interface is blackholed.

  Conditions: This symptom occurs on all interfaces removed from a different VDC, followed by the removal of other interfaces from the affected VDC.

  Workaround: This issue is resolved.

- CSCuo63409

  Symptom: When a fabricpath spine node is configured in transit mode, the Cisco Nexus 7000 Series device may experience a traffic outage.

  Conditions: This symptom is more likely to occur in a scaled testbed with more VLANs, SVIs, or L2 multicast group entries. A reload or VDC restart may also trigger this issue.

  Workaround: This issue is resolved.

- CSCuo15015

  Symptom: A Cisco Nexus 7000 Series switch running NX-OS Release 6.2(2a) software may crash during the Unicast Routing Information Base (URIB) process.

  Conditions: This symptom has no known trigger.

  Workaround: This issue is resolved.

- CSCuo62072

  Symptom: A Cisco Nexus 7000 Series switch may experience a Web Cache Communication Protocol (WCCP) process crash.

  Conditions: This symptom only occurs if WCCP is configured on the Cisco Nexus 7000 Series device.

  Workaround: This issue is resolved.

- CSCuo57841

  Symptom: In a PVLAN configuration, IP routed traffic may be dropped on some of the ports.

  Conditions: This symptom results in a traffic drop.

  Workaround: This issue is resolved.

- CSCuo99528

  Symptom: In a vPC environment with a private-VLAN, traffic that is not directly accessing a Cisco Nexus 7000 Series switch, but through a peer and peer link, ARP replies are dropped.

  Conditions: This issue is due to the ARP tunneling mechanism implemented on F1 ports.

Workaround: This issue is resolved.

- CSCun86405

  Symptoms: Orphan ports connected over an FEX HIF may observe traffic being blackholed because the MAC address is not learned correctly in a vPC+ domain.

  Conditions: This symptom occurs because the MAC is learned from the emulated software identification (SWID) tag instead of the local SWID due to incorrect internal registers.

  Workaround: This issue is resolved.

- CSCuo55926

  Symptoms: The Multicast Routing Information Base (MRIB) fails when reloading VDC.

  Conditions: This symptom occurs when reloading VDC.

  Workaround: This issue is resolved.

- CSCuo20496

  Symptoms: While EIGRP debugs are enabled on a Cisco Nexus 7000 Series switch, the EIGRP service fails.

  Conditions: This symptom is seen when enabling EIGRP debugs.

  Workaround: This issue is resolved.

- CSCuj937222

  Symptoms: The **show system internal pktmgr internal event-history lcache-err** command causes inband path latency errors.

  Conditions: This issue occurs when running the **show tech** command.

  Workaround: This issue is resolved.

- CSCul35819

  Symptoms: BPDUs sent from a remote device do not activate BPDUGuard on an edge trunk port.

  Conditions: This symptom occurs when a remote device is not configured correctly, or when a VLAN on an access port is not included in the allowed VLAN list of the edge trunk.

  Workaround: This issue is resolved.

- CSCun50901

  Symptoms: FEX modules on a Cisco Nexus 7000 Series switch may fail without performing a save core dump.

  Conditions: This symptom occurs during normal operation of the Cisco Nexus 7000 Series switch.

  Workaround: This issue is resolved.

- CSCul71240

Symptoms: Netflow monitoring data is not collected on traffic ingressing an F2 line card.

Conditions: F2E modules and M-series modules (M1/M2) are in the same VDC, whereas the Netflow flow monitor is associated with a SVI interface on the F2E module. When the F2E module is reloaded, Netflow monitoring data is not collected on the M-series module.

Workaround: This issue is resolved.

- CSCuo20186

  Symptoms: The **show otv** command output displays the secondary IP address of a join interface that has already been removed.

  Conditions: This symptom can be seen when you:

  - Configure an IP or secondary IP on a port channel or sub-interface
  - Add that interface to an overlay interface as join-interface
  - Remove that interface using no int.
  - Reconfigure the same join interface with different secondary IP addresses.
  - Run the **show otv** command to show all IP addresses, as well as the deleted IP and secondary IP addresses.

  Workaround: This issue is resolved.

- CSCuo28456

  Symptoms: External routes in OSPF are installed over a NSSA area rather than the backbone area.

  Conditions: If there is reachability of OSPF external prefixes in both the NSSA area and the backbone area, OSPF will install the routes over those prefixes in the NSSA area.

  Workaround: This issue is resolved.

- CSCum01502

  Symptoms: A Cisco Nexus 7000 Series switch running Cisco NX-OS Release 6.1(1) reboots due to a netstack HAP reset.

  Conditions: This symptom has no known triggers.

  Workaround: This issue is resolved.

- CSCum20367

  Symptoms: The Cisco Nexus 7000 Series device experiences a supervisor crash due to an SNMP HAP reset

  Conditions: This symptom occurs when the Cisco Nexus 7000 Series switch is running Cisco NX-OS Release 6.2(2).

  Workaround: This issue is resolved.

- CSCun63523

  Symptoms: A monitorc crash occurs on an M2 line card due to a memory leak.

  Conditions: This symptom is seen only on M2 line cards.

Workaround: This issue is resolved.

- CSCum29184

  Symptoms: An snmpd process crash happens when polling the CommonUserTable and the following log message is displayed:

  ```
  %SYSMGR-3-HEARTBEAT_FAILURE: Service "snmpd" sent SIGABRT for not setting heartbeat
  for last 6 periods. Last heartbeat 64.85 secs ago.
  %SYSMGR-2-SERVICE_CRASHED: Service "snmpd" (PID 4329) hasn't caught signal 6 (core
  will be saved).
  ```

  Conditions: SNMP will not respond to CommonUserTable after 12 to 20 users log in using TACACS. This causes the SNMP process to fail.

  Workaround: This issue is resolved.

- CSCuo24939

  Symptoms: The ACL MGR process experiences a memory leak and displays the following message:

  ```
  Lib ID: 0x1
  Mem stats for UUID : Non mtrack users(0) Max types: 177
  --------------------------------------------------------------------------------
  TYPE NAME                                       ALLOCS                 BYTES
                                           CURR      MAX        CURR       MAX

     2 [r-xp]/isan/plugin/0/isan/bin/aclmgr 5285752 5285756  363171903  363181841
  ```

  The memory leak may lead to a process crash, resulting in a system reload due to a HAP Reset.

  Conditions: This symptom occurs when a user runs the configure session command to configure ACLs. This can also be seen if a user configures time-ranges.

  Workaround: This issue is resolved.

- CSCun94251

  Symptoms: RIPv2 is not propagating default-routes.

  Conditions: This symptom occurs on a Cisco Nexus 7000 Series switch running NX-OS Release 6.2.(6).

  Workaround: This issue is resolved.

- CSCum52399

  Symptoms: A Cisco Nexus 7000 Series device running NX-OS Release 6.2(2) sets the default EIGRP MTU to 16384, regardless of the interface MTU. If the system receives a large number of update packets, the packets may get fragmented. The CoPP on the Cisco device can drop these fragments as they get classified under the default-class of CoPP.

  Conditions: This symptom is caused by the EIGRP MTU being set to 16384.

  Workaround: This issue is resolved.

- CSCuo77566

  Symptoms: An OTV Overlay Interface goes into a stuck state showing “Cleanup in Progress.”

```
Overlay in delete holddown, reject command.
```

Conditions: This issue is seen on Cisco Nexus 7000 Series devices funning NX-OS Release 6.2(2). It also carries over through ISSU to NX-OS Release 6.2(6a).

Workaround: This issue is resolved.

- CSCun34667

  Symptoms: A Cisco Nexus 7000 Series switch experiences an unexpected reboot when running the **show vPC consistency** command.

  Conditions: This symptom is seen when entering the s**witchport mode private-vlan trunk promiscuous** command.

  Workaround: This issue is resolved.

- CSCum76187

  Symptoms: An STP interface state is seen incorrectly on CE interfaces,after performing a sup switchover *LOOP_Inc and FWD & *LOOP_Inc & BKN after a second switchover.

  Conditions: The VDC should have both CE and FP interfaces.

  Workaround: This issue is resolved.

- CSCuo79856

  Symptoms: On a Cisco Nexus 7000 Series switch in a CTS enforcement, not all policies for a configured security group tag may appear on the device after a port-channel becomes active.

  For example, in the following interface configuration:

  ```
  interface Ethernet1/1
    cts manual
      policy static sgt 0x64 trusted
    switchport mode fabricpath
    no snmp trap link-status
    logging event port link-status
    channel-group 11 mode active
  ```

  If this port channel flaps, and two member ports of the port channel come up back-to-back, not all policies configured on an ISE or ACS server with a destination group tag 0x64 may be seen in the **show cts role-based policy** output.

  Conditions: This symptom is seen on Cisco Nexus 7000 Series devices when configuring CTS.

  Workaround: This issue is resolved.

- CSCuo82093

  Symptoms: A customer edge (CE) multicast receiver receives duplicate multicast packets for a period of 1 - 10 seconds. The receivers are physically located behind one or more provider edge routers on a Cisco Nexus 7000 Series switch running Cisco NX-OS Release software. The CE multicast sources are physically located behind a pair of Cisco Nexus 7000 Series devices in a vPC domain.

Conditions: These symptoms occur on an MVPN with the Cisco Nexus 7000 Series devices positioned as provider edge routers, and the multicast sources are part of a vPC domain. Both Cisco Nexus 7000 Series devices are also configured for Rendezvous Points for the customer edge VRF.

Workaround: This issue is resolved.

- CSCuo93903

  Symptoms: Packets with an IP payload of between 1497 to 1500 are dropped on an F3 interface with a default MTU of 1500.

  Conditions: This symptom is seen on the F3 modules of a Cisco Nexus 7000 Series switch running NX-OS Release 6.2(6a).

  Workaround: This issue is resolved.

- CSCuo96097

  Symptoms: Layer 2 broadcast frames are encapsulated differently on an F3 OTV virtual device context (VDC).

  Conditions: This symptom is seen when an F3 OTV VDC uses the incorrect multicast MAC for L2 broadcast frames.

  Workaround: This issue is resolved.

- CSCup09291

  Symptoms: When Bidirectional Forwarding Detection (BFD) is configured for a port channel, a BFD port channel per-link session is configured. After the member links flap several times, some member sessions might get stuck in the down state.

  Conditions: This symptom occurs when a BFD port-channel per-link session is configured, resulting in interface flapping.

  Workaround: This issue is resolved.

- CSCup10643

  Symptoms: Enabling NetFlow on a Cisco Nexus 7000 Series device causes the line cards to fail.

  Conditions: This symptom occurs on Cisco Nexus 7000 Series devices with NetFlow enabled.

  Workaround: This issue is resolved.

- CSCup42943

  Symptoms: When you have a combination of VDC M2 and F2E modules running fabricpath, a frame with DMAC 0000.0000.0000 loops endlessly between the fabricpath peers.

  Conditions: This issue is seen in a chassis running a combination of M2 and F2E modules running fabricpath and applies to only NX-OS Releases 6.2(6) and 6.2(8).

  Workaround: This issue is resolved.

- CSCup46560

  Symptoms: The AAA daemon process crashes on Cisco Nexus 7000 Series switches.

Conditions: This symptom is seen on Cisco Nexus 7000 Series switches only.

Workaround: This issue is resolved.

- CSCup52842

  Symptoms: IBGP continuously flaps after an ISSU upgrade from NX-OS Release 6.2(8) to 6.2.9(3)

  Conditions: This symptom is seen only after an ISSU upgrade.

  Workaround: This issue is resolved.

- CSCup72846

  Symptoms: One or both supervisors on a Cisco Nexus 7000 Series chassis may reload unexpectedly after running the **monitor hap reset** command.

  Conditions: This might happen only with VLAN addition/deletion,

  AND when SPAN is configured,

  AND the SPAN source contains a FPC interface.

  Workaround: This issue is resolved.

- CSCum78918

  Symptoms: Syslog messages are not being sent to the server even though pinging the syslog-server shows the server is working correctly.

  Conditions: This issue is seen under normal operational conditions.

  Workaround: This issue is resolved.

- CSCum94152

  Symptoms: License grace period alerts are still sent after the MPLS grace period expires or the feature is disabled.

  Conditions: This symptom occurs after a grace period expires.

  Workaround: This issue is resolved.

- CSCum68908

  Symptoms: Performing the following sequence of SNMP walks on a switch that is fully loaded, configured in a snake topology, and has traffic running through the snake, causes a *Reason: (tooBig) Response message would have been too large* error.

  – If the bulkwalk is executed by itself, no issue is seen.

  – If the bulkwalk is preceded by 1 snmpwalk, no issue is seen (on a manual retry).

  – If the bulkwalk is preceded by 2 snmpwalks, exactly as below, the issue can be reproduced manually.

  Conditions: This symptom is seen when running SNMPv3.

  Workaround: This issue is resolved.

- CSCuo81646

  Symptoms:

  The link to ports 41 through 48 remain up even if the cable is disconnected or the partner device is admin down.

  Conditions: This symptom is seen when using a1Gig N77-F348XP-23 module.

  Workaround: This issue is resolved.

- CSCui18245

  Symptom: When forming an LACP port channel, the port channel does not come up, and the interface moves to suspended state.

  Conditions: You might see this symptom if you are running either of the following configurations:

  – **vlan dot1q tag native** globally

  – **switchport trunk native vlan** *x* on an interface

  Workaround: This issue is resolved.

- CSCuo37471

  Symptom: You issued either the **shutdown** or **no hsrp** command and the HSRP VIP is down, but the route is still installed.

  Conditions: This symptom might be seen when a new SVI is created and HSRP is configured while you are running Cisco NX-OS Release 6.2(6).

  Workaround: This issue is resolved.

- CSCtu12501

  Symptoms: A keepalive frame received at HIF is returned to the HIF by the forwarding engine of Cisco Nexus 7000 Series switch, which makes the switch show a loopback was detected.

  Conditions: This symptom occurs when connecting a Layer 2 switch to HIF.

  Workaround: This issue is resolved.

- CSCup05138

  Symptoms: CPU generated traffic such as HSRP and ARP is getting dropped and does not sent to the forwarding egress line card. When capturing the HSRP or ARP packets via ethanalyzer,, the ethanalyze has set the NXOS DEST INDEX: 0 as follows:

```
METMXC2(config)# ethanalyzer local interface inband decode-internal capture-fi "host
10.47.25.4 and host 224.0.0.102" det Capturing on inband NXOS Protocol
    NXOS VLAN: 247
    NXOS SOURCE INDEX: 1054
    NXOS DEST INDEX: 0
```

  Conditions: This issue occurs when the HSRP state is Unknown and ARP requests are incomplete.

  Workaround: This issue is resolved.

- CSCup04851

Symptoms: On a Cisco Nexus 7000 Series switch, the Bidirectional Forwarding Detection (BFD) protocol reports that a session is supposed to be up, but OSPF states that BFD is still down.

Conditions: This symptom was reported on a Cisco Nexus 7000 Series switch running NX_OS Release 6.2(6),

Workaround: This issue is resolved.

- CSCuo63486

  Symptoms: Even when DCBX TLV is disabled on a Cisco Nexus 7000 Series switch, the port is still error disabled after acknowledging100 PDUs.

  Conditions: This symptom occurs when reloading the switch.

  Workaround: This issue is resolved.

- CSCug12791

  Symptoms: Performing the following steps on a Cisco Nexus 7000 Series switch might cause an ACL QOS process crash:

  1. Enable access-list resource pooling.
  2. Load a large number of access-list entries.
  3. Purchase and load the scalable services license to expand the available TCAM space.
  4. Disable access-list resource pooling.

  Conditions: This symptom is caused by removing resource pooling from a Cisco Nexus 7000 Series switch.

  Workaround: This issue is resolved.

- CSCup15693

  Symptoms: When using Cisco NX-OS Release 6.2, you cannot configure SNMP traps to successfully broadcast IP addresses using the **snmp-server host x.x.x.255 traps** command.

  Conditions: This symptom impacts Cisco Nexus 7000 Series switches that are running NX-OS 6.2 and later.

  Workaround: This issue is resolved.

- CSCuh72694

  Symptoms: A PTP clock does not synchronize with a master clock connected via ether-channel on an F2 line card.

  Conditions: This symptom is due to the PTP server being upstream connected via an ether-channel on an F2 line card

  Workaround: This issue is resolved.

- CSCuh89791

  Symptoms: Jobs that are run on a Cisco Nexus 7000 Series switch using Scheduler may regularly fail with the following error messages:

  ```
  Error: opening file: /tmp/cmd sched.job_nameschedule_name
  ```

```
Syntax error while parsing '<i><user configured command></i>''
```

Conditions: This symptom occurs when jobs are configured on multiple VDCs to be run at the same time using the same job and schedule names.

Workaround: This issue is resolved.

- CSCuj13356

  Symptoms: You cannot remove a QoS global configuration using %IPQOSMGR-4-QOSMGR_PPF_WARNING.

  Conditions: This symptom is seen when you configure a policy to a SPAN destination port and the process fails with the following error:

  ```
  "ERROR: Unable to perform the action due to incompatibility:  Module 3, 8, 13 returned
  status "Egress policy on an L2 interface is not supported"
  ```

  Workaround: This issue is resolved.

- CSCup24183

  Symptoms: Inserting a CPAK LR4 into a device may cause the link to not come up.

  Conditions: This symptom might be seen when you are working with the CPAK LR4 transceiver.

  Workaround: This issue is resolved.

- CSCul73502

  Symptom: Configuring a QoS ingress service-policy and not mixing/matching type qos class-maps under the same policy may be rejected with the following error message:

  ```
  ERROR: Unable to perform the action due to incompatibility:  Module 1, 2 returned
  status "Policies with classes containing combined 'match dscp', 'match cos', 'match
  precedence' or 'match qos-group' are not supported. Only the same match type is
  supported between classes."
  ```

  Conditions: This symptom may be seen when configuring an ingress marking service-policy on a L2 switchport port-channel interface.

  Workaround: This issue is resolved.

- CSCul93014

  Symptom: The CoPP process crashes frequently due to a memory leak in the process. The memory leak can be verified using the following command:

  show sys int copp mem-stats detail | grep *COPP_MEM_match_t*

  Conditions: This symptom is seen if the first number in the output is increasing every 5 to 10 minutes.

  Workaround: This issue is resolved.

- CSCum06429

Symptom: On a Cisco Nexus 7000 Series switch with N7K-F248XP-25E and N7K-SUP2 modules with PBR statistics added, the line card fails after a hard restart.

Conditions: This symptom is seen when a line card fails because PBR statistics were included in the PBR statement.

Workaround: This issue is resolved.

- CSCum22217

  Symptom: Incoming DSCP values are not saved when an egress policy setting includes MPLS Experimental Bits and other ingress policies. Use ingress and egress policies to set MPLS experimental bits only.

  Conditions: This symptom occurs on a Cisco Nexus 7000 Series switch running NX-OS Release 5.2(9),

  Workaround: This issue is resolved.

- CSCup25953

  Symptom: Input discard occurs when an egress 1-Gigabit port is down due to a remote system being down.

  Conditions: This issue happens only with 1-Gigabit egress ports on an F3 module.

  Workaround: This issue is resolved.

- CSCum47367

  Symptom: A vulnerability in the TACACS command authorization code of a Cisco Nexus 7000 Series switch might allow an authenticated, local attacker to execute certain commands without authorization from the TACACS server.

  Conditions: This symptom is seen when the switch is configured for TACACS command authorization.

  Workaround: This issue is resolved.

- CSCum53340

  Symptom: A BGP failure occurs on a Cisco Nexus 7000 Series switch when the device experiences a heartbeat failure.

  Conditions: This symptom is seen on Cisco Nexus 7000 Series switches when BGP is using regular CLI expressions.

  Workaround: This issue is resolved.

- CSCum58820

  Symptom: The ifHCoutUcastPkts value decrements on interfaces.

  For example:

  ```
  first poll: 1.3.6.1.2.1.31.1.1.1.11.436256768, Counter64    , 11187
  after 20 min: 1.3.6.1.2.1.31.1.1.1.11.436256768, Counter64    , 11112
  ```

Conditions: This symptom is seen on Cisco Nexus 7000 Series switchs running NX-OS Release 6.2(2a) on multiple interfaces.

Workaround: This issue is resolved.

- CSCum77476

  Symptom: The **show sockets internal event-history proto** command shows passwords in clear text. The password should not be shown in clear text, even in an internal event-history.

  ```
  show sockets event-history display bgp neighbor passwords
  382) Event:E_DEBUG, length:52, at 703540 usecs after Fri Jan 24 03:18:42 2014
      [162] [6512]: Password:cisco123 and address:1010101
  ```

  Conditions: This symptom is seen on Cisco Nexus 7000 Series switches.

  Workaround: This issue is resolved.

- CSCum99773

  Symptom: In an EIGRP topology, routes remain in “State is Active”.

  Conditions:

  Workaround: This issue is resolved.

- CSCuo39308

  Symptom: An LLDP peer on a Cisco Nexus 7000 Series switch is including an optional TLV in the LLDP frame with a length of 0. This optional TLV does not have a specified minimum length per the LLDP standard (802.1AB).

  Conditions: This symptom is seen on Cisco Nexus 7000 Series switches running NX-OS Release 6.2(8) and earlier.

  Workaround: This issue is resolved.

- CSCun34585

  Symptom: When upgrading from Cisco NX-OS Release 5.2(1) to NX-OS Release 6.2(2a) using a multi-step ISSU, the following error message is displayed:

  ```
  ACLQOS-SLOT1-2-PPF_FAILED  Database failure: aclqos_dst_find failed while reading
  stats
  ```

  Conditions: This symptom is seen during an upgrade procedure.

  Workaround: This issue is resolved.

- CCSCup58719

  Symptom: The SystemMgmtBus test fails when a faulty fan is replaced. Even though the fan has been replaced, it still fails testing.

  Conditions: This symptom occurs when performing a SystemMgmtBus test.

  Workaround: This issue is resolved.

- CSCun68179

  Symptom: In an OTV environment, once FHRP isolation is configured it cannot be removed.

  Conditions: This symptom is seen during FHRP configuring.

  Workaround: This issue is resolved.

- CSCun31865

  Symptom: The following error messages may be printed and is not very useful:

  ```
  %NETSTACK-2-MPULLUP:  netstack [4628]  udp_input: m_pullup failed for UDP, error
  Operation not permitted
  ```

  Conditions: This symptom does not have a known cause.

  Workaround: This issue is resolved.

- CSCun88767

  Symptom: After copying and pasting the **copy <file> running-config** command, an empty line is added next to the **banner motd** command.

  Conditions: This symptom occurs when the copy running-config.

  command contains a return code CRLF.

  Workaround: This issue is resolved.

- CSCun97048

  Symptom: After a disruptive upgrade to Cisco NX-OS Release 6.2(6), ACLs may no longer be programmed correctly.

  Conditions: This issue is seen on XL line cards only.

  Workaround: This issue is resolved.

- CSCuo24370

  Symptom: The **crypto** command is in a different location in the startup and running configurations, which results in a failure when entering the **show running-config diff** command

  Conditions: This symptom might be seen when running the show **running-config diff** command using the CLI.

  Workaround: This issue is resolved.

- CSCuo54566

  Symptom: TACACS authentication fails on a Cisco Nexus 7000 Series device with the following error in the TACACS debugs.

  ```
  tplus_make_author_request: enterin2014 Mar 25 09:32:28 ALT-CORE-SW3
  %USER-3-SYSTEM_MSG: Unable to create temporary user x655501. Error 0xffffffff (0) -
  sshd
  ```

  Conditions: This symptom does not have a known cause.

  Workaround: This issue is resolved.

- CSCuo77974

  Symptom: Continuously adding or deleting VLAN/SVI pairs at the same time, as in both vPC peer switches, LIFs on the vPC peer and vPC port are leaked.

  Conditions: This symptom is seen on Cisco Nexus 7000 Series NX-OS Release 6.2(8) devices.

  Workaround: This issue is resolved.

- CSCuo55621

  Symptom: The /var/tmp directory is filled with npacl.debug files after enabling an IP access list.

  Conditions: This symptom occurs when entering the **disable ip access-list detailed** command.

  Workaround: This issue is resolved.

- CSCuo57613

  Symptom: On a Cisco Nexus 7000 Series switch, when the r**ate-mode dedicated force** command is enabled on a Layer 3 10 Gigabit interface, the MACsec interface does not come up. However, when the bandwidth is configured to be shared, the MACsec tunnel comes up correctly.

  Conditions: This symptom is seen on a 10 Gigabit interface (N7K-M148GS-11) with the **rate-mode dedicated force** command enabled.

  Workaround: This issue is resolved.

- CSCuo24944

  Symptom: After connecting a fabric extender (FEX) to a standard F2 module on a Cisco Nexus 7000 Series switch with FET optics, and then changing to standard optics, the port does not work correctly. Traffic is not sent to the CPU so protocols such as ARP and CDP do not work.

  Conditions: This symptom occurs when you reconfigure a port that was an FEX fabric port as a normal port.

  Workaround: This issue is resolved.

- CSCuo35291

  Symptom: You might see routing to a local subnet fails after you remove the LISP configuration.

  Conditions: This symptom might be seen when the LISP mobility configuration is present before the upgrade or when the LISP mobility configuration is removed when you issue any of the following commands:

  - **no feature lisp**

  OR

  - **interface** *number*
  - **no lisp mobility**

  Workaround: This issue is resolved.

- CSCuo83543

Symptom: On a Cisco Nexus 7000 Series switch, after a VDC reload, it takes about 5 minutes for the peer keepalive link (L3) to come up.

Conditions: This symptom is seen during initial system bootup or a VDC reload.

Workaround: This issue is resolved.

- CSCuo83756

  Symptom: A Cisco Nexus 7000 Series switch may experience an MTS buffer leak when a syslog server is configured by a hostname.

  Conditions: This symptom is only seen when the hostname is not configured correctly.

  Workaround: This issue is resolved.

- CSCuo84009

  Symptom: Ports on F2 and F2E modules may fail to learn MAC addresses. The output for the **show hardware internal error** command shows that the following count is incrementing:

  ```
  Ingress redirect due to DNL check
  ```

  Conditions: This symptom is seen on F2 and F2E modules only and applies to NX-OS Release 6.2(8) and earlier.

  Workaround: This issue is resolved.

- C5SCup93375

  Symptom: Because ISSU Cisco NX-OS Release 5.2(5) to pre-release of NX-OS Release 6.2(10) is not supported a reload upgrade was performed. When the reload was completed, all interfaces were VDC unallocated.

  Conditions: This symptom is seen when a reload upgrade from NX-OS Release 5.2(5) to a pre-release of Cisco NX-OS Release 6.2(10).

  Workaround: This issue is resolved.

- CSCui72592

  Symptom: The F3 module may reload unexpectedly due to a kernel panic at InstructionTLBError47x. Any or all of the following symptoms might be present:

  – Syslogs indicate that the module was unresponsive, and was reset.

  ```
  Example:
  2014 Aug 21 20:20:08 NEXUS7K %MODULE-2-MOD_NOT_ALIVE: Module 2 not responding...
  resetting (Serial
   number: XXXXXXXXXXX)
  2014 Aug 21 20:20:19 NEXUS7K %PLATFORM-2-MOD_DETECT: Module 2 detected (Serial
  number XXXXXXXXXXX)
   Module-Type 10/40 Gbps Ethernet Module Model N77-F324FQ-25
  2014 Aug 21 20:20:19 NEXUS7K %PLATFORM-2-MOD_PWRUP: Module 2 powered up (Serial
  number XXXXXXXXXXX
  )
  2014 Aug 21 20:20:19 NEXUS7K %PLATFORM-5-MOD_STATUS: Module 2 current-status is
  MOD_STATUS_POWERED
  _UP
  2014 Aug 21 20:21:57 NEXUS7K %BIOS_DAEMON-SLOT2-5-BIOS_DAEMON_LC_PRI_BOOT:  System
  booted from Pri
  ```

```
mary BIOS Flash
2014 Aug 21 20:24:03 NEXUS7K %VDC_MGR-5-VDC_STATE_CHANGE: vdc 1 state changed to
updating
2014 Aug 21 20:24:03 NEXUS7K %VDC_MGR-5-VDC_STATE_CHANGE: vdc 1 state changed to
active
2014 Aug 21 20:24:32 NEXUS7K %PLATFORM-5-MOD_STATUS: Module 2 current-status is
MOD_STATUS_ONLINE/
OK
```

- The output of **show logging onboard module <x> internal reset-reason** command indicates that the reload reason was due to a kernel panic.

```
Reset Reason for this card:
        Image Version : 6.2(8a)
        Reset Reason (LCM): Line card not responding (60) at time Thu Aug 21
20:21:59 2014
        Reset Reason (SW): Kernel Panic (19) at time Thu Aug 21 20:19:28 2014
          Service (Additional Info): Kernel Panic
        Reset Reason (HW): System reset by active sup (by writing to PMFPGA regs)
(100) at time Thu Aug 2
1 20:21:59 2014
        Last log in OBFL was written at time Thu Aug 21 19:32:39 2014
```

- The output of **show logging onboard module <x> stack-trace** command displays the kernel stack trace.

Note The top function on this stack is "InstructionTLBError47x." The running process and the rest of the stack are not relevant.

Example:

```
---------------------------------
 Module: 2 stack-trace
---------------------------------

------------------   LOG 01    -------------------

Logging time: Thu Aug 21 20:19:28 2014
1408670368:500000000 process xxxxxxxx (1219), jiffies 0xb79c1f6
 Machine Check in kernel mode : Process xxxxxxxxx (1219)
 MCSR: 0x0 MCAR: 0x0  MCSSR0: 0xc040100c MCSSR1: 0x21000


STACK

 CPU 1  Call Trace:

 [<c040100c>]InstructionTLBError47x+0x6c/0xa8 <------------- HERE
```

Conditions: You might see this issue on an F3 Series module with any current Cisco NX-OS Release.

Workaround: This issue is resolved.

- CSCuo93631

Symptom: After you perform an ISSU, your switch does not learn new MAC addresses in the software. The **show hardware mac address-table** *module-number* command, where *module-number* is the ingress module in question, includes the MAC entry in the output and the **show mac address-table** command does not include the same MAC entry.

Conditions: You might see this issue if you perform an ISSU on a Cisco Nexus 7000 Series switch, or in combination with an ISSU, a new learned MAC address is introduced by a new MAC, or a flush of the MAC table and relearn of existing MAC addresses occurs while the system is undergoing an ISSU.

Workaround: This issue is resolved.

- CSCuh24768

  Symptom: The VLAN translation table entries might not be correct for vPC leg port channels with PVLAN mode configured. This might cause incorrect VLAN translation in egress direction for PVLAN vPC legs.

  Conditions: VLAN translation tables might contain the translations for both the previous PVLAN mode and for the current PVLAN, or non PVLAN, mode if the following conditions are true:

  - The vPC leg is in PVLAN mode (PVLAN host, PVLAN promiscuous, PVLAN trunk isolated, PVLAN trunk promiscuous) and PVLAN port mapping is configured.
  - The mode is changed to a different mode, PVLAN or non PVLAN, and new port mappings are added on the vPC leg PC.

  Typically, the Translation table for the vPC leg includes programming specific to the current port-mode only.

  Workaround: This issue is resolved.

- CSCuo14607

  Symptom: During an ISSU, lldpRemIndex will still return 0 instead of valid index id in the range: (Integer32(1..2147483647)) as per the LLDP MIB RFC. This condition is not present in a fresh boot of Cisco NX-OS Release 6.2(8). This condition is only present during an ISSU from Cisco NX-OS Release 6.2(x) to 6.2(8).

  Conditions: This might be seen during an ISSU from Cisco NX-OS Release 6.2(xc) to Cisco NX-OS Release 6.2(8) if LLDP is enabled in the Cisco NX-OS 6.2(x) release on the switch and there are one or more LLDP neighbors of the switch.

  Example:

  ```
  [root@sse-26 ~]# snmpwalk -v 2c -c public 10.104.238.120 1.0.8802.1.1.2.1.4.1.1.8
  LLDP-MIB::lldpRemPortDesc.0.440451072.0 = STRING: Ethernet7/13
  ```

  Workaround: This issue is resolved.

- CSCul56872

  Symptom: vlan_mgr crash is seen on a Cisco Nexus switch in one or more vdc's.

  Conditions: Reserve VLAN configuration - 'system vlan < > reserve' followed by "show vlan internal usage" command in non-default vdc or when trying to save the configuration right after the changes.

  Workaround: This issue is resolved.

- CSCun60847

  Symptom: This is a modification on the product to adopt new secure code best practices to enhance the security posture and resiliency of the product.

  Conditions: Device configured with default configuration.

  Workaround: This issue is resolved.

- CSCun90491

  Symptom: Traffic loss will be observed on IGMPV3 (G,S) entries for untouched VLAN's on F3 modules.

  Conditions: 1) IGMP snooping enabled and working in IGMPv3 mode (and)

  2) Delete of VLAN (and)

  3) Deleted VLAN's HW_BD should be same as the MAC table HW BD field value of any other VLAN's S entry.

  Under the above conditions traffic loss on IGMPV3 (G,S) entries will be seen. The traffic loss is not applicable when IGMPV2 (G,*) entries are used. This issue happens very rarely because of condition 3.

  Workaround: This issue is resolved.

- CSCup10049

  Symptom: An ACLQoS crash occurs with the accompanying core file(s): %SYSMGR-SLOT3-2-SERVICE_CRASHED: Service "aclqos" (PID <snip>) hasn't caught signal 6 (corewill be saved). The following errors are logged on the switch after the crash:

  - `014 May 29 08:30:12.708 snrchi5pcore1 %IPQOSMGR-3-QOSMGR_PPF_ERROR: PPF library error: MTS Error 0x801c0010 (Device or resourc e busy) after 1706 retries .`
  - `2014 May 29 08:31:47.563 snrchi5pcore1 %IPQOSMGR-3-QOSMGR_PPF_ERROR: PPF library error: MTS Error 0x801c0010 (Device or resourc e busy) after 1755 retries .`

  Conditions: The switch is running Release 6.2(2a).

  Workaround: This issue is resolved.

- CSCup47983

  Symptom: %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet6/29 is down (Link failure)

  %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet6/30 is down (Link failure)

  %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet6/31 is down (Link failure)

  Port on same SOC may flap during while removing 10G SFP and using 1G SFP instead.

  Conditions: When an SFP is inserted into an interface after module is UP and no shut is executed on that interface, the port of the same speed and in the same port-group might flap.

  Workaround: This issue is resolved.

- CSCuq23295

Symptom: Crash seen with cli 'show system internal clk_mgr get reference-port-ts interface eth7/15'.

Conditions: clk_mgr crash was seen with cli 'show system internal clk_mgr get reference-port-ts interface eth7/15'.

Workaround: This issue is resolved.

- CSCuq55749

  Symptom: When using HSRP MGO an attempt to configure a new HSRP follow group may result in this group being stuck in the Initial state with HSRP reporting a reason of "No Master for Slave".

  "No Master for Slave" condition results from group being stuck in initial state.

  Conditions: Occurs after a supervisor switchover when the HSRP MGO is configured and HSRP follow groups are deleted from the running config.

  Workaround: This issue is resolved.

- CSCuq85943

  Symptom: Getting TCAM error while configure L2 feature (PACL/QOS) followed by egress RACL with bank chaining. The same is working fine in reverse order - egress racl followed by L2 feature.

  Conditions: Only for the egress.

  Workaround: This issue is resolved.

- CSCuq89979

  Symptom: Two Cisco Nexus 7000 are being used as VPLS peers,

  When the access switch connecting to the Cisco Nexus 7000 floods the packet to all interfaces except the one on which it has received, the Cisco Nexus 7000 also gets a copy of the packet and resends the same packet back to the source MAC of the packet to be learned on the interface going to the Cisco Nexus 7000s on access switch which causes mac flap and traffic loss until the devices sends packet again and corrects the learning.

  Conditions: Cisco Nexus 7000 connecting to the same VLAN.

  Workaround: This issue is resolved.

- CSCur05565

  Symptom: Traffic is blackholed when source is intermittent.

  Conditions: Occurs when:

  - Data MDT is configured.

  - Source starts/ stops for few minutes and starts again.

  Workaround: This issue is resolved.

- CSCur12912

  Symptom: Following module bringup after failure, vmm timeout seen causing ports not getting allocated to vdc. RPM cored post that and VDC came up

Conditions: None

Workaround: None

- CSCur14220

  Symptom: Netstack process crash.

  Conditions: Crash can happen after removing and consequently adding interface VLANs with ACL/QOS/PBR configuration - this happen usually when config is copied from file to running configuration.

  Workaround: This issue is resolved.

- CSCur19732

  Symptom: "ETHPM timed out to bring up all portch" error message will be seen on the console when the "port-channel load-balance hash-modulo" command is applied.

  Conditions: "ETHPM timed out to bring up all portch" error message will be seen on the console when the "port-channel load-balance hash-modulo" command is applied.

  Workaround: This issue is resolved.

- CSCur19960

  Symptom: IPv6 multicast transient traffic drops during ISSU.

  Conditions: During ISSU [LC upgrade time frame] PIM6 might age out S,G entries, since packet statistics are not available during that time. This would result in a transient loss.

  Workaround: This issue is resolved.

- CSCur20515

  Symptom: The Cisco Nexus 7000 with F1 modules may modify CoS markings through the switch.

  Conditions: The Cisco Nexus 7000 utilizing the F1 line card might maintain the default CoS mappings for the module and not program the default mapping derived from the system QoS policy. This problem can happen during the following scenarios:

  - Switch reload
  - VDC reload
  - Module reload

  Workaround: This issue is resolved.

- CSCuh85338

  Symptom: Cisco Nexus 7000 object-group displays a minus sequence number if the sequence number is larger than 2147483647.

  N7K-1# sh object-group test

  IPv4 address object-group test

```
    2147483647 host 1.1.1.1
    2147483648 host 2.2.2.2
    4294967294 host 3.3.3.3
    4294967295 host 4.4.4.4
```

But from "show run", we can see minus sequence number like -2147483648, -2 and -1 like the following output

```
N7K-1# sh run | b object-group
object-group ip address test
  2147483647 host 1.1.1.1
  -2147483648 host 2.2.2.2  <-----------
  -2 host 3.3.3.3 <-----------
  -1 host 4.4.4.4 <-----------
```

This is just a cosmetic issue

Conditions: If sequence number is larger than 2147483647

Workaround: This issue is resolved.

- CSCun16236

  Symptom: The OSPF interfaces might be down after you perform an ISSU.

  Conditions: You might see this symptom after you perform an ISSU.

  Workaround: This issue is resolved.

- CSCuh24768

  Symptom: When you are working with vPCs and private VLANS, you might see incorrect mapping.

  Conditions: This symptom might be seen when you are working with vPCs and private VLANS.

  Workaround: This issue is resolved.

- CSCul40023

  Symptom: ACLQoS crashes when a RACL is applied on a tunnel interface on an F3 Series module.

  Conditions: This symptom might be seen when you are running Cisco NX-OS.

  Workaround: This issue is resolved.

- CSCup99777

  Symptom: When adding a port to a port channel the member port inherits the MTU of the port channel.

However when this member is removed in config the MTU it inherited from the port channel remains in config. But in hardware it still retains the MTU value it had previously before joining the port channel. hence there is a mismatch.

Conditions: removal of port from portchannel

Workaround: This issue is resolved.

- CSCue91483

  Symptom: It takes 30 seconds to bring up a port channel configured with LACP.

  Conditions: If a port channel has nine or more member ports and is configured as a trunk port.

  Workaround: This issue is resolved.

- CSCun74440

  Symptom: MAC in MAC header (e.g. Fabricpath header) is preserved on traffic when sent to SPAN destination. No way to disable this for customers who do not want this functionality. This bug requests a new command to remove fabricpath header on traffic going to a span destinations.

  Conditions: SPAN of FabricPath traffic to any destination.

  Workaround: This issue is resolved.

- CSCup47983

  Symptom: %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet6/29 is down (Link failure)

  %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet6/30 is down (Link failure)

  %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet6/31 is down (Link failure)

  Port on same SOC may flap during while removing 10G SFP and using 1G SFP instead.

  Conditions: When an SFP is inserted into an interface after module is UP and no shut is executed on that interface, the port of the same speed and in the same port-group might flap.

  Workaround: This issue is resolved.

- CSCuq21501

  Symptom: Symbol-Error counted up for the interfaces in F3 modules as follows:

```
STLD1-630-02.10-N7K-RU37# sh int e1/47 counters errors


--------------------------------------------------------------------------------
Port        Align-Err    FCS-Err   Xmit-Err    Rcv-Err  UnderSize OutDiscards
--------------------------------------------------------------------------------
Eth1/47            31          0          0         31          0          0


--------------------------------------------------------------------------------
```

```
Port        Single-Col  Multi-Col   Late-Col  Exces-Col  Carri-Sen       Runts
-------------------------------------------------------------------------------
Eth1/47             0          0          0          0          0          31


-------------------------------------------------------------------------------
Port          Giants SQETest-Err Deferred-Tx IntMacTx-Er IntMacRx-Er Symbol-Err
-------------------------------------------------------------------------------
Eth1/47           0          --           0           0           0     115089    <<==


-------------------------------------------------------------------------------
Port        InDiscards
-------------------------------------------------------------------------------
Eth1/47             0
```

Conditions: This occurs when no shut is done for the interface and link is down. And it was confirmed for the interfaces ranging from 41 to 48 ports in F3 module.

Workaround: This issue is resolved.

- CSCul91443

  Symptom: After configuring OTV on a Cisco Nexus 7000 Series switch, adjacency does not come up between sites.

  Conditions: This symptom is seen when the value of the MTU in the core is less than 1500.

  Workaround: This issue is resolved.

- CSCty67801

  Symptom: This is a feature request for SVI, where SVI creation has to fail if VFI is configured under a VLAN, and vice-versa, VFI configuration under a VLAN has to fail if corresponding SVI is created.

  Conditions: If both SVI and VFI are configured for a VLAN at the same time.

  Workaround: This issue is resolved.

- CSCub21497

  Symptom: There is a programming failure on a port channel and the following error message appears:

  ```
  %WCCP-1-SBADDFAIL: Unable to add WCCP subblock on
  interface Vlan200: Error string: Verify failed in LC
  ```

  Conditions: This symptom might be seen when the redirect-list is attached to WCCP groups, when a policy is attached to a port-channel interface, or when a VLAN has a WCCP policy attached to a port-channel interface.

  Workaround: To work around this issue, restart the feature by entering the **no feature wccp** command and the **feature wccp** command.

- CSCud60005

  Symptom: Multiple ingress queues show the same nonzero drop counter when the **show policy-map interface** *ethx/y* command is entered.

  Conditions: This symptom might be seen under the following conditions:

  - The CoS is moved between different queues in such a way that it has the same Independent VLAN Learning (IVL) value for different queues. The output of the **show queuing interface** *ethx/y* command shows the IVL value of the ingress queues.
  - If the template of the system is changed to one of the template types with multiple queues that have the same IVL, such as 8e-4q4q in a Cisco Nexus 7000 Series switch and 8e-4q8q in a Cisco Nexus 7700 Series switch.

  Workaround: This issue is resolved.

- CSCue00645

  Symptom: A rollback fails when as part of the patch, the current trunk-allowed VLAN list is the default 1000 to 4000 and there is a **switchport trunk allowed vlan** *range* command after that in the configuration.

  Conditions: This symptom might be seen on the 32-port 10-Gigabit Ethernet I/O module XL (N7K-F132XP-15).

  Workaround: This issue is resolved.

- CSCug93147

  Symptom: A VDC reload can cause the following syslog in a FEX scale setup:

  ```
  013 Might 15 13:09:22 F2-FEX %VNTAG_MGR-2-VNTAG_SEQ_ERROR: Error ("sequence timeout")
  while communicating with component MTS_SAP_HP_I
  FTMC for Opcode MTS_OPC_VNTAG_ELTMC_SET_VLAN_CBL
  2013 Might 15 13:10:28 F2-FEX %ETHPORT-5-IF_DOWN_ERROR_DISABLED: Interface
  Ethernet154/1/48 is down
  ```

  Conditions: This symptom might occur in a FEX scale setup with 16 FEXes connected to one line card and a VDC reload is triggered for the VDC.

  Workaround: This issue is resolved.

- CSCuh18007

  Symptom: After an ISSU to Cisco NX-OS Release 6.2(2), BFD sessions that are booted on IPv6 interfaces stay in a down state.

  Conditions: This symptom might be seen on Cisco Nexus 7000 Series devices. It applies only to BFD sessions that are booted on IPv6 interfaces after an ISSU to Cisco NX-OS Release 6.2(2) from an earlier release.

  This issue does not impact existing and newly booted BFD sessions that use IPv4 interfaces.

  Workaround: This issue is resolved.

- CSCuh44476

  Symptom: Some neighbors are not discovered and the virtual circuits do not come up.

Conditions: This symptom might be seen when Virtual Flow Interfaces (VFIs) with autodiscovery Border Gateway Protocol (BGP) are configured, and some of the provider edges (PEs) are VDCs on the same Cisco Nexus 7000 Series device.

Workaround: This issue is resolved.

- CSCuh51343

  Symptom: In a VDC with ERSPAN destination sessions that are operationally up, port error messages might appear after certain events such as a VDC reload, module reload, or sequence timeout. Some of the affected ports might move to an error disabled state. The error message might look like the following:

  ```
  ETHPORT-2-IF_SEQ_ERROR: Error ("sequence timeout") communicating with <none: Internal
  Error> for opcode <None> (RID_PORT: Ethernet3/1)
  ```

  Conditions: This symptom might be seen under the following conditions:

  - More than one operationally up ERSPAN destination session is present in the same VDC.
  - The VDC reloads.
  - The module with one or more ERSPAN destination ports reloads.

  Workaround: This issue is resolved.

- CSCuh51701

  Symptom: The WCCP policies are not applied and errors are displayed.

  Conditions: This symptom might appear when any type of error goes to the application. After that, no policies are applied.

  Workaround: This issue is resolved.

- CSCuh73709

  Symptom: The command-line interface (CLI) does not return the switch prompt.

  Conditions: This symptom might be seen in Release 6.2(2) when the pseudowire interface is entered for a VLAN that it does not belong to. In this case, the **show mac addr tbl vlan** *x* **interface** *pw y* command was entered, and pw y belonged to VLAN z.

  Workaround: This issue is resolved.

- CSCui02179

  Symptom: In a large scale setup running rapid Per VLAN Spanning Tree (PVST) on a root bridge change, STP disputes can occur.

  Conditions: This symptom might be seen on a root bridge change. Multiple sequences of events are triggered starting from Layer 2 to the Layer 3. The number of these sequences is proportional to the number of VLANs in the system. This situation creates contention for the CPU that does not allow enough CPU for STP to send and receive BPDUs to be able to perform the root bridge for all VLANs.

  Workaround: This issue is resolved.

- CSCui07565

Symptom: Probe packets sent by the Gold StandbyFabric Loopback test and RewriteEngine test are dropped and the current run of the tests are treated as failed. The HIGH_NULL_POE_DROP_CNT and RO4 TDS timeout interrupt counters also start incrementing for the fabric module.

Conditions: This symptom might be seen on a Cisco Nexus 7700 switch when the fabric module is power cycled or the fabric ejectors are quickly opened and closed.

Workaround: This issue is resolved.

- CSCui07806

  Symptom: IPv6 Unicast Reverse Path Forwarding (URPF)/Redirects is configured on an interface and the **switchport** command is configured on the interface. In this case, IPv6 URPF/Redirects get disabled, and the Layer 3 configuration associated with the interface is removed from that interface.

  The interface is again moved to a Layer 3 interface when the **no switchport** command is entered on it and an IPv6 address is configured on the interface. In this case, even though the IPv6 URPF/Redirects configuration is not available on the interface, IPv6 URPF/Redirects continues to be enabled on that particular interface because of stale information.

  Conditions: This symptom might be seen when an interface that has IPv6 URPF/Redirects enabled is moved to a switchport, the **no switchport** command is entered on that interface, and an IPv6 address is configured. There is no issue (URPF/Redirects is disabled) if the interface is moved to a switchport and the interface continues to be in a switchport state. The issue occurs only after the **no switchport** command is entered and an IPv6 address is configured on the interface.

  Workaround: This issue is resolved.

- CSCui20080

  Symptom: When multiple overlays are defined with VLAN mapping configurations, attempts to roll back fail.

  Conditions: This symptom might be seen when you define multiple OTV overlays, and the rolled back configuration requires a redefinition of VLAN mappings for both overlays, or you add or remove VLAN mappings for both overlays.

  Workaround: This issue is resolved.

- CSCui25984

  Symptom: MPLS LDP is missing local labels, which can result in LDP not installing labeled routes and not advertising labels to peer LSRs. To confirm this issue, enter the **show mpls ldp internal event err | inc wbool_set** command to check for the presence of LDP event traces:

  ```
  2013 Jul 18 10:06:01.365948 ldp [9075]:ldpx_sched_wbool_set: invalid arg
  ```

  Conditions: This symptom might be seen when unconfiguring or reconfiguring MPLS LDP. It is possible that LDP does not allocate local labels for routes.

  Workaround: This issue is resolved.

- CSCui26012

  Symptom: The CISCO-VLAN-MEMBERSHIP-MIB does not support Layer 2 VPN interfaces.

  Conditions: This symptom might be seen because pseudowire and VFI membership are not supported in the CISCO-VLAN-MEMBERSHIP-MIB.

Workaround: This issue is resolved.

- CSCui26886

  Symptom: When a VDC is reloaded, the reload fails. A "Failure at interface manager" error might appear after the failure happens.

  Conditions: This symptom might be seen when multiple features like NetFlow, QoS, and various types of ACLs such as VLAN ACLs, router ACLs, and port ACLs are configured on interfaces that belong to a VDC.

  Workaround: This issue is resolved.

- CSCui28043

  Symptom: FabricPath capable ports that are coming up can become error disabled with a "sequence timeout" error when another module is reloaded.

  Conditions: This symptom might be seen when a FabricPath capable port is coming up and the EthPM process sends a message to the port client in all of the line cards. If the line card goes offline at the same time as the message is sent, this issue occurs. The line card going offline exactly at the same time is rare.

  Workaround: This issue is resolved.

- CSCui33560

  Symptom: The initial walk of onep_routing_rib_add_route_state_listener for IPv4 does not return all the routes. Subsequent route state events return all the routes correctly.

  Conditions: This symptom might be seen if the initial walk is without any filter for the specific owner type.

  Workaround: This issue is resolved.

- CSui37100

  Symptom: ACL configurations become inactive, which means they are present in the configuration but not on the module.

  Conditions: This symptom might be seen when you perform a non-ISSU upgrade to Cisco NX-OS Release 6.2(2) from an earlier release, and ACL configurations that were active before the upgrade become inactive.

  This symptom is not seen when you perform an ISSU on switches with dual supervisor modules.

  This symptom occurs when you do the following:

  - Enter the **copy running-config startup-config** command on a Cisco Nexus 7000 Series device running a release earlier than Release 6.2(2).
  - Change the boot variables to Release 6.2(2).
  - Reload the device with Release 6.2(2).

  This issue also occurs if you perform an ISSU on a switch with a single supervisor module.

  Workaround: This issue is resolved.

- CSCui43107

  Symptom: Intermittent Intelligent Service Card Manager (ISCM) failures occur on a Cisco Nexus 7000 Series switch when there are three or more NAM cards in a single VDC.

  Conditions: This symptom might be seen when there are multiple (three or more) NAM cards in a single VDC.

  Workaround: This issue is resolved.

- CSCui53128

  Symptom: Layer 2 policies are not applied if you attempt to apply an ASCII configuration file from a release earlier than Cisco NX-OS Release 6.2(2) to Release 6.2(2). The policies that are affected include:

  mac port access-group *name*

  Conditions: This symptom might be seen if you generate an ASCII configuration file in a release earlier than Cisco NX-OS Release 6.2(2) and apply it to Release 6.2(2). Layer 2 policies will not be present in the running configuration.

  Workaround: This issue is resolved.

- CSCui60557

  Symptom: A LISP Tunnel Router (xTR) is unable to register database entries in setups that have only default routes.

  Conditions: This symptom might be seen when a LISP xTR that has only default routes is not able to perform lookups and forward LISP control plane messages. This situation results in the LISP xTR being unable to register database entries or send map requests.

  Workaround: This issue is resolved.

- CSCui61230

  Symptom: When a new pseudowire (PW) is added for a Virtual Fabric Interface (VFI) context, it does not come up.

  Conditions: This symptom might be seen for manual PWs, such as those configurations of the “member 1.2.3.4 encapsulation mpls” type in the VPLS context.

  Workaround: This issue is resolved.

Resolved Caveats—Cisco NX-OS Release 6.2(8b)

- CSCuq98748

  Symptom: The Nexus 7000 includes a version of bash that is affected by the vulnerabilities

  identified by the Common Vulnerability and Exposures (CVE) IDs:

  CVE-2014-6271

  CVE-2014-6277

CVE-2014-7169

CVE-2014-6278

CVE-2014-7186

CVE-2014-7187

This bug has been opened to address the potential impact on this product.

All current versions of NX-OS on this platform are affected unless otherwise stated

.

Exposure is not configuration dependent.

Authentication is required to exploit this vulnerability.

This bug is fixed in NX-OS versions specified below:

5.2(9a)

6.1(5a)

6.2(8b)

6.2(10) and above

Condition: A user must first successfully log in and authenticate via SSH to trigger this vulnerability.

Workaround: This issue is resolved.

Resolved Caveats—Cisco NX-OS Release 6.2(8a)

- CSCuo80937

  Symptom: TCNs are displayed from the interface every 35 seconds. Connected devices do not see the TCNs.

  Condition: This symptom might be seen on a Cisco Nexus 7000 Series device running Cisco NX-OS Release 6.2(6).

  Workaround: This issue is resolved.

- CSCuh55184

  Symptom: When you have a combination of M1/M2 modules with F1/F2e modules, looping of packets when the destination MAC address is all zeroes might cause a network traffic issue.

  Conditions: You might see this symptom if you are running Cisco NX-OS Release 6.2(6), 6.2(6a) or 6.2(8) on your Cisco Nexus 7000 or 7700 Series switches.

  Workaround: This issue is resolved.

- CSCuo73774

Symptom: While you are configuring an IPv6 access control list (ACL) on an F3 line card, you might experience the following unexpected behavior regardless of the ACL configuration: the traffic is dropped, redirected, permitted, or QoS policed.

Conditions: You might see this symptom when you are configuring an IPv6 ACL on an F3 line card and you are running Cisco NX-OS Release 6.2(6), 6.2(6a), or 6.2(8) on your Cisco Nexus 7000 or 7700 Series switches.

Workaround: This issue is resolved.

- CSCuo50598

  Symptom: A Cisco Nexus 7000 Series switch with an F3 module may experience link flaps on port channels when reloading a VDC.

  Conditions: Configure the port channel on an F3 module with active traffic.

  Workaround: This issue is resolved.

Resolved Caveats—Cisco NX-OS Release 6.2(8)

- CSCum40651

  Symptom: You cannot enter a CLI **config** command or a regular **show** CLI command with more than 64 characters (including white spaces and full CLIs). Cisco NX-OS will print an AAA authorization error.

  Conditions: This symptom might be seen when you have a TACACS+ server with CLI authorization, restricted CLI access for users, and a CLI string greater than 64 characters. (The adjacent memory location for the buffer where the CLI string will be stored should have nonzero characters to simulate a non-null terminated string.)

  Workaround: This issue is resolved.

- CSCum12906

  Symptom: The switch does not send traffic destined to an Ethernet port shared with the storage VDC.

  Conditions: This symptom might be seen under these conditions:

  - The shared port is on an F2e Series module. (F2 Series modules do not show this problem.)

  - The port is a shared port with an FCoE storage VDC.

  Workaround: This issue is resolved.

- CSCul48242

  Symptom: The switch might go down when you perform an ISSU from Cisco NX-OS Release 6.1(4) to Release 6.2(2a) because of the VLAN manager. When the standby supervisor is loading a new version of the Cisco NX-OS software through ISSU, the standby supervisor might crash and subsequently reset all modules in the chassis.

  Conditions: This symptom might be seen when you have vPC+ configured on the switch.

  Workaround: This issue is resolved.

- CSCum38422

  Symptom: The system might go down with the following reasons displayed: Reset triggered due to HA policy of Reset (sysmgr stateful recovery) or Reset triggered due to HA policy of Reset (AAA Daemon hap reset).

  Conditions: This symptom might be seen during a switchover and when you have more than one VDC.

  Workaround: This issue is resolved.

- CSCum05295

  Symptom: You might see the following error message in the syslog and IGP-redistributed BGP routes might fail:

  ```
  2013 Dec 13 01:21:58 PTAINTS410 %OSPF-3-RPM_LIB_API_FAILED: bgp_lookup_ext_attr() -
  failed in rpm_acquire_bgp_shmem_lock()
  ```

  Conditions: You might see this symptom when the following conditions exist:

  - IGP (for example, OSPF) redistributes BGP routes.
  - The redistribution uses a route map to evaluate the community that is associated with the routes.
  - The **maximum-paths** command is configured.
  - BGP receives paths with only attribute (for example, AS-PATH) change.

  Workaround: This issue is resolved.

- CSCul52450

  Symptom: The system might have may have a prefix installed in the unicast routing table but the prefix is unreachable. The prefix is using the default route recursive next hop, and the route in the URIB but has not been pushed down to the FIB after a RIB change (routing update).

  Conditions: You might see this symptom when the recursive next hop (RNH) is the default route (0.0.0.0/0) and there should be a routing update (device or link failure, which means that re-routing has occurred).

  Workaround: This issue is resolved.

- CSCum33567

  Symptom: The UDLD process might go down because of memory corruption.

  Conditions: You might see this symptom associated with mismatched link detection errors.

  Workaround: This issue is resolved.

- CSCum30306

  Symptom: The security process might go down, which results in a HAP Reset that brings down the supervisor.

  Conditions: You might see this symptom when you are configuring SSH authentication.

  Workaround: This issue is resolved.

- CSCuj17443

  Symptom: The **inherit** command is not working when TACACS authorization is enabled.

  ```
  switch(config)# interface port-channel1006
  switch(config-if)# inherit port-profile Blade-Servers
  ERROR: Failed to write VSH commands
  ```

  Conditions: This symptom might be seen with the Cisco Nexus7000 (9-slot) chassis ("Supervisor Module-1X") with Cisco NX-OS Releases 6.0(4) to 6.2(1) and Cisco ACS 5.3 patch 6

  Workaround: This issue is resolved.

- CSCul68883

  Symptom: After you delete an HSRP bundle, other Anycast HSRP bundles might have an incorrect state.

  Conditions: This symptom might be seen when you are working with Anycast HSRP.

  Workaround: This issue is resolved.

- CSCum18989

  Symptom: Multicast traffic starts flowing on the link that is supposed to be pruned for multicast.

  Conditions: This symptom might be seen when a Cisco Nexus 7000 Series device is running Cisco NS-OS Release 6.2.2a and a Cisco Catalyst 6500 device is connected both upstream and downstream with a router-on-a-stick scenario.

  Workaround: This issue is resolved.

- CSCum14547

  Symptom: BGP VPNV4 might not synchronized properly.When a route in the source VRF is removed or added, the target VRF is either keeping a state route or not adding the route back when it is removed or added.

  Conditions: This symptom might be seen when a route in the source VRF is added or removed.

  Workaround: This issue is resolved.

- CSCul11180

  Symptom: BGP might go up and down on a Cisco Nexus 7000 Series device that is running the Cisco NX-OS Release 6.2(2) because of an FD read error

  Conditions: This symptom might be seen when you are running Cisco NX-OS Release 6.2.

  Workaround: This issue is resolved.

- CSCum74698

  Symptom: When you are running Cisco NX-OS Release 6.2(6) configured for detailed IP ACL logging, the system fills up the /var/tmp part of the disk, which affects the CLI. You will see the following error message:

  ```
  No space left on device
  ```

Conditions: This symptom might be seen when you are running Cisco NX-OS Release 6.2 (6) configured for detailed IP ACL logging.

Workaround: This issue is resolved.

- CSCun24082

  Symptom: You might see an ARP entry as Storm status without updating with the new MAC (failovered NIC MAC) with the default configuration, when the sever NIC has failed.

  Conditions: This symptom might be seen when you are running Cisco NX-OS Release 6.2 (2) configured with the default GARP storm configuration.

  Workaround: This issue is resolved.

- CSCun32932

  Symptom: When you are running Cisco NX-OS Release 6.2 (2) and you disable the GARP storm by entering the **no ip arp garp-storm** command, this command is not displayed when you enter the **show run** or **show run all** commands.

  Conditions: This symptom might be seen when you are running NX-OS Release 6.2 (2) configured with the default GARP storm configuration.

  Workaround: This issue is resolved.

- CSCun37067

  Symptom: Some dynamic MAC address entries might be missing for certain port groups on the F2 and F2e Series modules after you reload the device or the device experiences a link-flap.

  Conditions: A timing issue might cause the MAC address to be removed from the hardware Layer 2 forwarding table, which leads to unknown unicast flooding.

  Workaround: This issue is resolved.

- CSCul36724

  Symptom: You might see the following error message when you are running Cisco NX-OS Release 6.1 (3):

```
dc-224-gw1# sh forwarding ipv6 inconsistency mod 1
IPV6 Consistency check (in progress): table_id(0x80000001) slot(1)
Elapsed time : 10157 ms
Inconsistent adjacencies:
  2. slot(1), vrf(default), ipaddr(fe80::226:51ff:fecb:e2c1), ifindex(Vlan112),
Adjacency missing FIB Hardware.
  4. slot(1), vrf(default), ipaddr(fe80::226:51ff:fecb:e2c1), ifindex(Vlan104),
Adjacency missing FIB Hardware.
  6. slot(1), vrf(default), ipaddr(fe80::250:56ff:feba:1a8e), ifindex(Vlan614),
Adjacency missing FIB Hardware.
```

  Conditions: You might see this symptom when you are running NX-OS Release 6.1 (3) with vPC and vPC+ configured.

  Workaround: This issue is resolved.

- CSCul48500

  Symptom: When you shut down a (Fabric Extender) FEX fabric port channel and bring it back up, the system might send linkDown and cieLinkDown traps again, rather than the expected linkUp.

  Conditions: You might see this symptom when you are running Cisco NX-OS Release 6.1 (3) on Cisco Nexus 2000 FEXs are connected to F2 Series modules. This issue has been seen with Cisco Nexus 2232 and 2248 module FEXs.

  Workaround: This issue is resolved.

- CSCul53824

  Symptom: You might see MAC addresses out of synchronization between FEs. The interfaces from the FE that is out of synchronization are those members of port channels configured as access ports.

  Conditions: You might see this symptom when you are running port channel configurations.

  Workaround: This issue is resolved.

- CSCul53909

  Symptom: You might see a port status as down for a LACP hot standby port when you perform an SNMP walk.

  Conditions: You might see this symptom when you are running Cisco NX-OS Release 6.1 (4).

  Workaround: This issue is resolved.

- CSCul57328

  Symptom: You might see the neighbor 1000BASE-T interface of an F2 or F2e Series module flapped when you run the Port Loopback test.

  Conditions: You might see this symptom when you have an F2 or F2e Series module with GLC-T (1000BASE-T module) and the port is configured as shut when it is connected to a Cisco Catalyst interface that is configured as no shutdown. The interface on the Catalyst switch might go up and down.

  Workaround: This issue is resolved.

- CSCul57444

  Symptom: The HSRP delay minimum might not work for IPv6 HSRP.

  Conditions: You might see this symptom when you have used the interface configuration mode to configure HSRP delay minimum.

  Workaround: This issue is resolved.

- CSCul57895

  Symptom: The switch does not return an RFC compliance value for lldpRemIndex. Whenever lldpRemIndex is used as an index in an LLDP MIB table, the switch does not respond with a valid index ID in the range (Integer32(1..2147483647) as per the LLDP MIB RFC and just uses the value of 0 for all lldpRemIndex values.

  Conditions: This symptom might be seen when lldpRemIndex is used as an index in an LLDP MIB table.

Workaround: This issue is resolved.

- CSCul69832

  Symptom: FEX modules temporarily go offline after a supervisor switchover.

  Conditions: This symptom might be seen when a switchover is triggered by a kernel panic.

  Workaround: This issue is resolved.

- CSCul71874

  Symptom: The system might drop packets above 1492 bytes and low the MTU if the link is configured for CTS static SGT.

  Conditions: This symptom might be seen when you are working with F2 and F2e Series modules.

  Workaround: This issue is resolved.

- CSCul80719

  Symptom: When you enable multicast ECMP, the mcastfwd process might go down.

  Conditions: This symptom might be seen when you enable multicast ECMP.

  Workaround: This issue is resolved.

- CSCul83215

  Symptom: You might see connectivity issues with end hosts because of ARP.

  Conditions: This symptom might be seen when you configure GLBP on an SVI with vPC configured.

  Workaround: This issue is resolved.

- CSCul85287

  Symptom: If you removed the active supervisor of two FabricPath spine switches and you are running HSRP anycast, you might see HSRP go down.

  Conditions: This symptom might be seen when you remove the active supervisor of both spine switches. A simple switchover does not trigger this symptom.

  Workaround: This issue is resolved.

- CSCum04595

  Symptom: In a vPC setup, if an NLB packet comes in through N7K1 and then to N7K2 on the peer-link, the packet is flooded on N7K2 (and vice versa).

  Conditions: This symptom might be seen if you are working with the Cisco NX-OS Release 6.1(4).

  Workaround: This issue is resolved.

- CSCum07047

  Symptom: You might see the system go down.

Conditions: This symptom might be seen if you are working with a multi-VDC system with the number of total processes and threads more than 5120 system wide.

Workaround: This issue is resolved.

- CSCum08129

  Symptom: You might see the snmpd go down.

  Conditions: This symptom might be seen if you are using SNMP to copy or archive the configuration for CiscoWorks LMS.

  Workaround: This issue is resolved.

- CSCum09912

  Symptom: If you are working on an M1 Series module and enter the **switchport monitor** command after you changed the primary port rate mode dedicated, ping fails on the other ports that belong to that same port group.

  Conditions: This symptom might be seen if you are working with port groups on an M1 Series module.

  Workaround: This issue is resolved.

- CSCum10610

  Symptom: You might see the TACACS+ daemon exiting with a message that the system could not initialize as IPC-TACACSd.

  Conditions: This symptom might be seen if you are working with TACACS+.

  Workaround: This issue is resolved.

- CSCum18490

  Symptom: The system is unable to turn PFC on an interface on an F2 Series module VDC when PONG is enabled in the VDC. The error message indicates that PFC cannot be enabled while PONG or PTP is enabled.

  Conditions: This symptom might be seen when PONG is already enabled and you attempt to configure PFC on the interface of an F2 Series module.

  Workaround: This issue is resolved.

- CSCug75586

  Symptom: Storm control does not function on uplinks of FEX on the Cisco Nexus 7000 Series devices.

  Conditions: This symptom might be seen when you configure a FEX port with storm control on a Cisco Nexus 7000 Series device.

  Workaround: This issue is resolved.

- CSCum47956

Symptom: Undersize (illegal) frames might be sent from the N7K-M108X2-12L module when a port on that module is configured for CTS/SGT. These small packets are not padded to the minimum required 64 bytes and might or might not be dropped on the peer depending on which minimum size the peer accepts when configured for CTS.

Conditions: This symptom might be seen when you configure the device for CTS.

Workaround: This issue is resolved.

- CSCuh64744

  Symptom: When you enter the **system jumbomtu 9216** command, you might see the following message:

  ```
  2013 Jun 11 22:07:24 nx7010-1-r212d  ETHPORT-2-IF_CRITICAL_FAILURE  (Debug
  syslog)Critical failure: qosmgr_dce_gldb_get_all_vl_params returned error: , no such
  pss key
  2013 Jun 11 22:07:24 nx7010-1-r212d  ETHPORT-5-IF_SEQ_ERROR  Error ("no such pss key")
  communicating with <None:Internal Error> for opcode <None> (RID_MODULE: 254)
  ```

  Conditions: This symptom might be seen after you perform an ISSU from Cisco NX-OS Release 6.0(3) to 6.1(3).

  Workaround: This issue is resolved.

- CSCum51886

  Symptom: Connectivity might fail during reporting on a system configured with the Smart Call Home (SCH) feature if an explicit class for either the HTTPS method or the SMTP method is not defined in the control-plane policing and there are continual violations occur in the copp class-default class.

  Conditions: This symptom might be seen when the configured destination from SCH is known inband.

  Workaround: This issue is resolved.

- CSCum52602

  Symptom: IPv6 neighbor discovery might go down when you configure and unconfigure a secondary address with the same subnet mask as the primary address; the solicited address mc group is deleted under the IPv6 interface.

  Conditions: This symptom might be seen when you configure more than one IPv6 address on an interface and each address has the same 24 least significant bits. Because the last 24 bits for both addresses are the same, global addresses configured on the interface have the same solicited node address.

  Workaround: This issue is resolved.

- CSCum54502

  Symptom: Although RFC1213 says the MIB OID for sysName is 255 characters, the Cisco Nexus 700 Series devices do not return the entire configured string for SysName. The SysName cannot be more than 32 characters.

  Conditions: This symptom might be seen when you are working with the SysName.

  Workaround: This issue is resolved.

- CSCum57545

  Symptom: After the system receives a corrupt BPDU, you might see VLANs go down because of a PVID mismatch and inconsistent vPC peer link. If the problem is because of a bad link, the VLANs and vPCs remain down even after the bad link has been shut down.

  Conditions: This symptom might be seen when the system receives a bad BPDU and then immediately receives a good BPDU for the same VLAN.

  Workaround: This issue is resolved.

- CSCum78696

  Symptom: After you upgrade to Cisco NX-OS Release 6.2(2a), you might receive system message related to the aaad and securityd processes.

  Conditions: This symptom might be seen when the remote logging server severity level is set to 7-debug and the default level is set to 5-notif.

  Workaround: This issue is resolved.

- CSCui92577

  Symptom: You might see the following error:

  ```
  2013 Aug 14 22:54:40 fc03.frc1 %U6RIB-3-U6RIB_ASSERT_ERROR:
  ../routing-sw/routing/u6rib/u6rib.c:4573: Assertion "*cand_best_count" failed.
  2013 Aug 14 22:54:40 fc03.frc1 %U6RIB-3-ASSERT_ERROR: -Traceback: 0x809dce5 0x805c89a
  0x806fe35 0x807003a 0x8071227 0x8073904 0x8075dbc 0x807b44d 0x807cbef
  librsw.so+0xa73ff libpthread.so.0+0x6140 libc.so.6+0xca8ce
  2013 Aug 14 22:55:43 fc03.frc1 %U6RIB-3-U6RIB_ASSERT_ERROR:
  ../routing-sw/routing/u6rib/u6rib.c:4573: Assertion "*cand_best_count" failed.
  ```

  Conditions: This symptom might be seen when there is a routing loop present in the network.

  Workaround: This issue is resolved.

- CSCum87512

  Symptom: You might see the following error after you enter a **show** command:

  ```
  mmap: Cannot allocate memory.
  ```

  Conditions: This symptom might be seen when you are working with **show** commands.

  Workaround: This issue is resolved.

- CSCun03801

  Symptom: When you are working on an F2 Series module with port security enabled and a port is not shut down and receives more than a 10 MAC addresses violation, that module might incorrectly forward the frames to other ports.

  Conditions: This symptom might be seen when you are working on an F2 or F2e Series module with Cisco NX-OS Release 6.2(6).

  Workaround: This issue is resolved.

- CSCuj56186

  Symptom: You might see that the PBR next hop is incorrectly installed into TCAM. If the system attempts to push PBR into TCAM before adjacency is resolved. the traffic might not be forwarded properly.

  Conditions: This symptom might be seen when you are working on an F2 Series module with all PBR-related interfaces configured on the same module.

  Workaround: This issue is resolved.

- CSCun09294

  Symptom: After you upgrade from Cisco NX-OS Release 6.1(3) to Release 6.2(2a), you might see broadcast traffic dropped because of storm control drops on some other unrelated ports. This issue causes a connectivity issue/outage because packets, such as ARP, are getting dropped.

  Conditions: This symptom might be seen when you are working on an F2e Series module after you upgrade from Cisco NX-OS Release 6.1(3) to Release 6.2(2a).

  Workaround: This issue is resolved.

- CSCuj66760

  Symptom: The device might keep responding to the OIDs in the MIB after you enter the **no snmp-server load-mib dot1dbridgesnmp** command, even when the corresponding **show** command displays it as unloaded.

  Conditions: This symptom might be seen when you are using SNMP to manage the device.

  Workaround: This issue is resolved.

- CSCun32383

  Symptom: You might see some ports on F3 Series modules drop all data plane traffic on FabricPath ports. These ports fail to learn any MAC addresses from FabricPath even though the FabricPath neighbor is established. When you enter the **show hardware internal error module** command, the display shows "Ingress redirect due to non-MIM pkt on core port" is incrementing.

  Conditions: This symptom might be seen when you are running Cisco NX-OS Release 6.2(6) on an F3 Series module with FabricPath enabled on the port.

  Workaround: This issue is resolved.

- CSCul15177

  Symptom: If you configure the same area range that is configured on an ABR for multiple areas, the component routes in one area (the common area range) are advertised into other areas even though they needed to be suppressed.

  Conditions: This symptom might be seen when you configure the same area range configured on an ABR for multiple areas.

  Workaround: This issue is resolved.

- CSCun11449

Symptom: When you are working on the M2 Series module and enter the **show hardware flow ip module** command, you might see the following error message:

```
Service not responding
```

Conditions: This symptom might be seen when you are working on an M2 Series module with active flows present on a forwarding instance.

Workaround: This issue is resolved.

- CSCum76354

  Symptom: You might see the following error message when you are working with Cisco NX-OS Release 6.2(6) and issue an SNMP Get request to the cefcFanTrayOperStatus with an invalid index:

  ```
  %SYSMGR-2-SERVICE_CRASHED: Service "Platform Manager"
  ```

  Conditions: This symptom might be seen when you are performing an SNMP get request for an invalid index.

  Workaround: This issue is resolved.

- CSCuj12578

  Symptom: You cannot use the **import hashlib** command in Python mode,

  Conditions: This symptom might be seen when you are working in Python mode.

  Workaround: This issue is resolved.

- CSCtx01036

  Symptom: The IP DHCP Relay Information Trusted feature is not supported in the interface configuration mode.

  Conditions: This symptom might be seen when you are working in interface configuration mode.

  Workaround: This issue is resolved.

- CSCul90391

  Symptom: When you are working with some interface transceivers and you enter the **show interface ethernet x/y transceiver details** command, the display does not show Transmit (Tx) or Receiver (Rx) Power information.

  Conditions: This symptom might be seen when you are working with the following interface transceivers: N7K-M202CF-22L/CFP-100G-SR10; N77-F324FQ-25, 24 port QSFP; N77-F312CK-26, 12 ports CPAK; N77-F348XP-23E, 48 Ports SFP+; and N7K-F312FQ-25, 12 Ports QSFP.

  Workaround: This issue is resolved.

- CSCum07531

  Symptom: You might see an ARP fail to be created when you are routing using SVI as next hop when ARP entries exist on an interface even after the interface is shut down.

  Conditions: This symptom might be seen when the system is routing using an SVI as the next hop.

Workaround: This issue is resolved.

- CSCud63152

  Symptom: Traffic destined to CPU is flooded instead of being punted.

  Conditions: You might see this symptom when a specific instance on an F2 Series module does not have the gateway MAC address of the CPU programmed (example instance 10):

```
module-5# sh mac address-table address 4055.3907.10c3 vlan 784 vdc DIST
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, (d) - dec
        Age - seconds since last seen,,+ - primary entry using vPC Peer-Link
        h - hex, d - decimal

VDC = 3

  FE  VLAN   MAC Address    Type     Age   Secure  NTFY   Ports/SWID .SSID.LID (d)
-----+-----+--------------+--------+------+-------+------+----------------------
***************truncated************************
G  9   784  4055.3907.10c3   static    -       F     F    sup-eth1    (R)
G 11   784  4055.3907.10c3   static    -       F     F    sup-eth1    (R)
```

  Workaround: This issue is resolved.
  Additional Notes on this issue: After you complete an ISSU to a fixed release (that is, Cisco NX-OS Release 6.1(5) and above), you must reload the F2 and F2e Series modules to make the fix applicable. If you do not reload the F2 and F2e Series modules, you may continue to see this problem.

- CSCuj64873

  Symptom: You might see the VDC fail to be created and the module fail to come up.

  Conditions: This symptom might be seen when you are working with at least five F2 or F2e Series modules (48 ports) with QOS policy configurations on ports or port channels and you reload the switch. The fifth module may fail to come up.

  Workaround: This issue is resolved.

- CSCun60965

  Symptom: You might see the ISSU from Cisco NX-OS Release 6.2(2a) to Release 6.2(6a) fail with the following error message returned:

```
Verifying image bootflash:/n7000-s2-dk9-npe.6.2.6a.bin for boot variable "system".
[#                   ]   0% -- FAIL.
Return code 0x40930077 (Install is not supported between NPE and non-NPE system
image).
Pre-upgrade check failed. Return code 0x40930011 (Image verification failed).
```

  Conditions: This symptom might be seen when you are performing an SSU from Cisco NX-OS Release 6.2(2a) to Release 6.2(6a).

  Workaround: This issue is resolved.

- CSCul24462

  Symptom: You might see a FabricPath interface incorrectly programmed by PIXM when you are working with an F2 Series module running Cisco NX-OS Release 6.1(4a) and Release 6.1(4).

Conditions: This symptom might be seen if there is a directly attached Cisco Nexus 5000 device running vPC+, and one of the vPC peers is being reloaded.

Workaround: This issue is resolved.

- CSCum78755

  Symptom: You might see improved switchover times in Cisco Release NX-OS 6.x.

  Conditions: This enhancement might be seen when you are running Cisco Release NX-OS 6.x.

  Workaround: This issue is resolved.

- CSCua15172

  Symptom: Some flows are received on the Netflow collector with an approximative 10+ minutes delay after the flows in question ended, although the flow active/inactive timeout is configured to a much lower value.

  Conditions: This symptom might be seen after you perform an ISSU from Cisco Release NX-OS 6.1(1) to Release 6.1(3) on an M1 Series module with Netflow configured.

  Workaround: This issue is resolved.

- CSCun53797

  Symptom: The switch might go down in the "ipqosmgr" process when SNMP tries to poll the interface QoS statistics. This is likely tied to polling the cbqos-mib in particular, since this MIB specifically causes the Nexus to access data related to its interface QoS statistics, which is the action that causes this crash.

  Conditions: This symptom might be seen when you are working with SNMP polling cbqos-mib.

  Workaround: This issue is resolved.

- CSCuo12477

  Symptom: BGP might go down after removing aggregate-addresses with attribute map defined.

  Conditions: This symptom might be seen when BGP aggregates are removed from the configuration and an attribute-map is applied. This symptom occurs only when you using an attribute-map with aggregates.

  Workaround: This issue is resolved.

- CSCuo00001

  Symptom: All mappings for the port are removed when the port channel is enabled with static sgt port mapping configured. New mappings for that port are not added as long as port-channel is configured.

  Conditions: This symptom might be seen when you are working with the Cisco NX-OS Release 6.2(6) and you have static sgt port mappings configured.

  Workaround: This issue is resolved.

- CSCun74218

Symptom: Fiber ports on the F2e Series module might be identified as a copper port. with copper breakout cable.

Conditions: This symptom might be seen when you are working with the F2e Series modules with fiber ports and copper breakout cables.

Workaround: This issue is resolved.

- CSCun73067

  Symptom: When you try to delete a WCCP policy, the action fails. If you attempt to add more policies, you might see the following error:

  ```
  Invalid operation
  ```

  Conditions: You might see this symptom when you have stopped a WCCP configuration previously.

  Workaround: This issue is resolved.

- CSCun84708

  Symptom: Run the debug process memory usage process might crash VSH.

  Conditions: You might see this symptom when you have are working with Cisco NX-OS Release 6.2(6).

  Workaround: This issue is resolved.

- CSCun99720

  Symptom: The system may unexpectedly perform and ISSD.

  Conditions: You might see this symptom when you are working with a wildcard configuration.

  Workaround: This issue is resolved.

- CSCun77235

  Symptom: All supervisors a might reload unexpectedly after an ISSU upgrade, with the reset reason show as **monitor hap reset**.

  Conditions: You might see this symptom when you perform an ISSU from Cisco NX-OS Release 6.2(2) to Release 6.2(6a) with SPAN running.

  Workaround: This issue is resolved.

- CSCun89683

  Symptom: Routed traffic on an F2 Series module might not reach the destination

  Conditions: You might see this symptom when you are working with an F2 Series module.

  Workaround: This issue is resolved.

- CSCuj77407

  Symptom: Unable to modify access list. The system gives the error message: service not responding. Unable to modify route map. The system gives the error message: SAP did not respond in expected time frame. Unable to save the configuration. the system gives the error message: Config lock in progress.

Conditions: This symptom might be seen when the NetStack process goes down in the middle of processing an ACL modification using a PPF session.

Workaround: This issue is resolved.

- CSCul58596

  Symptom: In a vPC environment, all the VLANs allowed over the vPC peer-link are suspended momentarily with (Reason: Vlan is not allowed on Peer-link) on the vPC primary switch when we remove an RSPAN VLAN that is also allowed over the peer-link. This happens for both the peer-link as well as all the vPC port channels.

  Conditions: This symptom might be seen when you have configured VTP server mode.

  Workaround: This issue is resolved.

- CSCuo05318

  Symptom: The vPC leg on the primary vPC Primary might be in STP Disable (DIS) state while this vPC leg is being brought up.

  Conditions: This symptom might be seen when you are working with a vPC leg on the primary vPC and there are one ore more member link flaps while the vPC leg is being brought up.

  Workaround: This issue is resolved.

- CSCum96491

  Symptom: After you switch over the supervisor, BFD goes down.

  Conditions: This symptom might be seen when you switch over the supervisor.

  Workaround: This issue is resolved.

- CSCum08181

  Symptom: When you are running Cisco NX-OS Release 6.2(2), Release 6.2(2a), Release 6.2(6), or Release 6.2(6a), you might see a shut down trunk promiscuous vPC port channel cause BPDU timeouts on vPC peer link in all VLANs.

  Conditions: This symptom might be seen when you are When you are running Cisco NX-OS Release 6.2(2), Release 6.2(2a), Release 6.2(6), or Release 6.2(6a) with port-channel vPCs and you enter the **shutdown** command for one of the trunk promiscuous vPC port channels.

  Workaround: This issue is resolved.

- CSCuo73479

  Symptom: When you are running Cisco NX-OS Release 6.2(8) using the F348XP module with a copper SFP, the interface remains in the "up/up"state even when the opposite end is admin down or the cable is removed from the SFP.

  Conditions: This symptom might be seen when you you are running Cisco NX-OS Release 6.2(8) using the F348XP module with an SFP with speed hardcoded on the interface to 1 Gigabit.

  Workaround: This issue is resolved.

- CSCul70166

  Symptom: If you enable the MVRP feature after you have enabled the dot1x feature, you might see a premature dropping of control packets during congestion drop windows.

  Conditions: This symptom might be seen when you enable the MVRP feature on the F1 Series module after enabling dot1x.

  Workaround: This issue is resolved.

Resolved Caveats—Cisco NX-OS Release 6.2(6b)

- CSCuo80937

  Symptom: Spanning-Tree Protocol (STP) TC Bridge protocol Data Units (BPDUs) are sent every 2 seconds for a long period of time after approximately 100 days of active supervisor uptime.

  Conditions: You might see this symptom if there are topology changes (TCs) after you upgrade to Cisco NX-OS Release 6.2(6), 6.2(6a) or 6.2(8) on your Cisco Nexus 7000 or 7700 Series switches.

  Workaround: This issue is resolved. In order to circumvent this issue until an upgrade to 6.2(8a) can be performed, execute the appropriate workaround, depending on whether you have a dual-supervisor or single-supervisor configuration, before each 90 days of uptime.

  Use to **show system uptime** command to display the number of running days for the active supervisor.

  ```
  Switch# show system uptime
   System start time:          Fri Oct 25 09:40:58 2013
   System uptime:              236 days, 8 hours, 56 minutes, 59 seconds
   Kernel uptime:              110 days, 23 hours, 7 minutes, 49 seconds
   Active supervisor uptime:   110 days, 23 hours, 2 minutes, 23 seconds
  ```

  For a dual-supervisor configuration:

  1. Reload the standby supervisor.
  2. Use the **show module** command to confirm that the standby supervisor is up and in the HA-standby mode.
  3. Use the **system switchover** command to switch to the standby supervisor.

  For a single-supervisor configuration:

  1. Upgrade to Cisco NX-OS 6.2(6b).
  2. Reload the switch.

- CSCup02927

  Symptom: When you bring up or reload a line card, the ELTM process crashes and a Stateful Switchover (SSO) occurs.

  Conditions: You might see this symptom if you are running Cisco NX-OS Release 6.2(6), 6.2(6a), or 6.2(8) on your Cisco Nexus 7000 or 7700 Series switches.

  Workaround: This issue is resolved.

- CSCuh55184

Symptom: When you have a combination of M1/M2 modules with F1/F2e modules, looping of packets when the destination MAC address is all zeroes might cause a network traffic issue.

Conditions: You might see this symptom if you are running Cisco NX-OS Release 6.2(6), 6.2(6a) or 6.2(8) on your Cisco Nexus 7000 or 7700 Series switches.

Workaround: This issue is resolved.

- CSCuo73774

  Symptom: While you are configuring an IPv6 access control list (ACL) on an F3 line card, you might experience the following unexpected behavior regardless of the ACL configuration: the traffic is dropped, redirected, permitted, or QoS policed.

  Conditions: You might see this symptom when you are configuring an IPv6 ACL on an F3 line card and you are running Cisco NX-OS Release 6.2(6), 6.2(6a), or 6.2(8) on your Cisco Nexus 7000 or 7700 Series switches.

  Workaround: This issue is resolved.

Resolved Caveats—Cisco NX-OS Release 6.2(6a)

- CSCum46336

  Symptom: Cos2q mapping changes after an ISSU from Cisco NX-OS Release 6.1(4a) to 6.2(6) with a single supervisor.

  Conditions: This symptom might be seen when you are working with the 8e-4q4q template.

  Workaround: This issue is resolved.

- CSCum52344

  Symptom: The auto-recovery disable configuration and auto-recovery timeout value are non persistent.

  Conditions: This symptom might be seen when you upgrade to Cisco NX-OS Release 6.2(x).

  Workaround: This issue is resolved.

- CSCul47903

  Symptom: You might see the following message:

  ```
  ERROR: Error in Timer Group library
  ```

  Conditions: This symptom might be seen when the queuing policy is attached to all the interfaces or most of the interfaces. In this case, the dscp2q map is pushed to interfaces one at a time instead of pushing for all interfaces at once.

  Workaround: This issue is resolved.

- CSCuj77407

  Symptom: Unable to modify access list. The system gives the error message: service not responding. Unable to modify route map. The system gives the error message: SAP did not respond in expected time frame. Unable to save the configuration. the system gives the error message: config lock in progress.

Conditions: This symptom might be seen when the Netstack process goes down in the middle of processing an ACL modification using a PPF session.

Workaround: This issue is resolved.

- CSCul44262

  Symptom: A ping between the Cisco Nexus 6000 Series Dynamic Fabric Automation (DFA) leaf and the Cisco Nexus Series 7000 BGP route-reflector spine nodes does not work. The BGP session will not be established and the Nexus 6000 switch DFA node will be isolated from the network if it does not have a redundant BGP RR session.

  Conditions: This symptom might be seen when you reload on either the Nexus 6000 switch or the Nexus 7000 switch, or when an interface over fabric control-segment goes up and down. The Nexus 7000 switch must be running Cisco NX-OS Release 6.2(6) code, when the BGP RR spine feature was first introduced, for you to see this symptom.

  Workaround: This issue is resolved.

- CSCun34460

  Symptom: You cannot enable the feature set for MPLS or for MPLS Layer 3 VPNs on the N77-F348XP-23 module.

  Conditions: This symptom might be seen when attempt to enable MPLS functionality on the N77-F348XP-23 module.

  Workaround: This issue is resolved.

- CSCug87825

  Symptom: You might see a route map perform action only on some of the matching prefixes when you are running BGP on the Cisco Nexus 7000 Series switch.

  Conditions: This symptom might be seen when the BGP has an outgoing route map to a peer, with 1500 prefixes or more from various peers, including the peer with the outgoing route map. When the BGP adjacency to a third peer goes up and down, the outgoing route-map matches the prefixes defined in a Permit Sequence but rather than just matching the prefix-list and performing the action on the sequence, it performs the action only on some of the matching prefixes.

  Workaround: This issue is resolved.

Resolved Caveats—Cisco NX-OS Release 6.2(6)

- CSCth03474

  Symptom: The Cisco Nexus 7000 Series switch generic online diagnostics (GOLD) do not report the failed module in some failure scenarios as part of the syslog.

  Conditions: This symptom might be seen if a failure is encountered with one of the crossbar ASICs.GOLD can incorrectly report the failed module or might be unable to isolate the exact module. For example, the Cisco Nexus 7000 Series switch active supervisor engine might report RewriteEngineLoopback or PortLoopback (or some other) test failed for all (or several) ports in all (or several) modules present in the switch.

  Workaround: This issue is resolved.

- CSCui35747

  Symptom: The Netstack process and other processes, such as SNMP and BGP might crash with the following error message:

  ```
  %SYSMGR-2-SERVICE_CRASHED: Service "snmpd" (PID 26066) hasn't caught signal 6 (core
  will be saved).
  ```

  Conditions: This symptom might be seen during high utilization when heartbeat messages begin failing.

  Workaround: This issue is resolved.

- CSCtq57444

  Symptom: STP shows the VLAN in PVID_Inc state on the trunk port between two N7K-C7010 with N7K-M132XP-12 modules.

  Conditions: This symptom might be seen after one port in an EtherChannel is brought down and up to recover from a UDLD err-disabled state.

  Workaround: This issue is resolved.

- CSCui48355

  Symptom: When a VDC ID is greater than 7, the ERSPAN source session configuration on the M2 Series modules might be disregarded by the ASIC driver. This issue will result in the SPAN source on an M2 Series module not being spanned.

  Conditions: This symptom might be seen when you have an Admin VCD present and the VDC ID can be 1 to 9.

  Workaround: This issue is resolved.

- CSCui72164

  Symptom: If you have a hardware failure on a vPC peer link with the Cisco Nexus 7000 Series switch, the switch might receive corrupted frames. If these corrupted frames continue to flow, a VDC or switch reload might occur.

  Conditions: This symptom might be seen when the following conditions are true:

  - The VDC is part of a vPC pair.
  - Corrupted frames are sent by the peer because of a hardware fault.

  Workaround: This issue is resolved.

- CSCui74777

  Symptom: The switch might ultimately crash in the VTP process and produce a core file, which has been traced back to a memory leak in VTP.

  Conditions: This symptom might be seen when the VTP feature is enabled and you enter a **show run** command.

  Workaround: This issue is resolved.

- CSCui88768

  Symptom: The HSRP virtual MAC address points to a local supervisor instead of a newly active HSRP router after an HSRP state change.

  Conditions: This symptom might be seen when you have an HSRP state change.

  Workaround: This issue is resolved.

- CSCue46437

  Symptom: GOLD might fail for a module during high-traffic congestion.

  Conditions: This symptom might be seen during high-traffic congestion.

  Workaround: This issue is resolved.

- CSCuj86229

  Symptom: FEX modules go down and up because of a watchdog timeout reset.

  Conditions: This symptom might be seen when you are running Release 6.2(6) with multiple FEX modules.

  Workaround: This issue is resolved.

- CSCug43851

  Symptom: The F2 Series and F2e Series modules might not filter packets when you are running the capture filter feature.

  Conditions: This symptom might be seen when you are running the capture filter feature.

  Workaround: This issue is resolved.

- CSCul25039

  Symptom: After making changes to an ACL that has been applied on an interface, the ACL definition is unexpected although the commit is successful Only the remark ACLs remain.

  Conditions: This symptom might be seen after you change an ACL that is applied to an interface.

  Workaround: This issue is resolved.

- CSCug98353

  Symptom: During a switchover, the crossbar attempts to gain access to the other crossbars in the system. If this process fails, the syslogs might contain an error message such as: XBAR 4 Fabric 0 on switchover initialization failed xbm_fabric_soft_init_on_swovr 1372. If the software does not cause the crossbar to fail, the modules reloading might fail when they try to come online.

  Conditions: This symptom might be seen after a switchover.

  Workaround: This issue is resolved.

- CSCuh71356

  Symptom: Connectivity within the same VLAN fails during the migration procedure from vPC to FabricPath.

Conditions: This symptom might be seen when you enable the FabricPath feature set by entering the **feature-set fabricpath** command.

Workaround: This issue is resolved.

- CSCuh97059

  Symptom: The snmpd process might terminate with the following message:

  ```
  %SYSMGR-2-SERVICE_CRASHED: Service "snmpd" (PID XXXX) hasn't caught signal 11 (core
  will be saved).
  ```

  Conditions: This symptom might be seen when polling using SNMP.

  Workaround: This issue is resolved.

- CSCui11913

  Symptom: When you disable STP on certain VLANs on a vPC peer and when the vPC is down on the primary peer, STP on the vPC secondary peer disables those VLANs.

  Conditions: This symptom might be seen when STP is disabled on the VLANs of a vPC peer.

  Workaround: This issue is resolved.

- CSCui15694

  Symptom: After you perform an ISSU, the counter for the SNMP Dot3Stats error counters are wrong.

  Conditions: This symptom might be seen after you perform an ISSU.

  Workaround: This issue is resolved.

- CSCui20860

  Symptom: An OSPF route might be missing after ports facing ASBR go up and down.

  Conditions: This symptom might be seen when ports facing ASBR go up and down.

  Workaround: This issue is resolved.

- CSCui86494

  Symptom: The switch is pointing a remote dynamically learned HSRP vMAC address to an incorrect interface and also shows the entry as static in the MAC-address table. After you add a static MAC entry to ensure the switch points the MAC address to the right interface, the switch points the MAC address to the correct interface. If the static MAC address is removed, the entry still shows up in the switch as static.

  Conditions: This symptom might be seen when you configure HSRP on both vPC peers but the SVI is in the shutdown state, which was up in the past, and the switch is not using HSRP on these devices.

  Workaround: This issue is resolved.

- CSCua69620

Symptom: MAC addresses in a vPC secondary switch are not deleted when a vPC+ peer link is brought down.

Conditions: This symptom might be seen when the MAC addresses in a vPC secondary switch are not deleted when the vPC+ peer link is brought down.

Workaround: This issue is resolved.

- CSCui05696

  Symptom: If you issue an unlimited ping from a Telnet session and that Telnet session is stopped, the ping process remains and takes up much of the CPU.

  Conditions: This symptom might be seen when you issue an unlimited ping.

  Workaround: This issue is resolved.

- CSCuc52448

  Symptom: The **show policy-map interface** *interface* command sometimes does not display the correct counter.

  Conditions: This symptom might be seen if you are using F2 Series or F2e Series modules and you enter the **show policy-map interface** command.

  Workaround: This issue is resolved.

- CSCue79883

  Symptom: You might see many retransmissions.

  Conditions: This symptom might be seen in a fully loaded setup with many logical/physical ports and periodically, polled statistics result in large buffers that are transmitted synchronously from the modules to the supervisor. The ports can be down and still be polled.

  Workaround: This issue is resolved.

- CSCug88508

  Symptom: A module cannot program existing ACL for QoS classification in TCAM, if reprogramming is triggered, and you see the following message in the log:

  ```
  ACLQOS-SLOT7-4-ACLQOS_OVER_THRESHOLD  Tcam 1 Bank 1's usage has reached its threshold.
  ```

  Conditions: This symptom might be seen when the TCAM utilization is close to 50% with an atomic update or 50% or more without an atomic update prior to upgrading to Cisco NX-OS Release 5.2(9) using an ISSU. Upgrading triggers a TCAM reprogramming. For example, this issue can be the result of a module that is reloading or if you are modifying an ACL that is used in QoS classification and reapplied to VLAN.

  Workaround: This issue is resolved.

- CSCuh71308

  Symptom: Disabling debug logging and deleting the syslogd-debugs does not release the capacity of the log.

  Conditions: This symptom might be seen when you delete the syslogd-debugs.

Workaround: This issue is resolved.

- CSCuj80663

  Symptom: If you are running Release 6.2(2), the **show interface** *interface number* **transceiver detail** command might show incorrect values for DOM.

  Conditions: This symptom might be seen when you are running Release 6.2(2) and using SR/LR SFPs.

  Workaround: This issue is resolved.

- CSCui26788

  Symptom: When you are using F2 Series modules, the not connected port link sometimes goes up and down.

  Conditions: This symptom might be seen when you administratively open the port and insert the GLC-T or other 1G SFP without cabling.

  Workaround: This issue is resolved.

- CSCui79503

  Symptom: The SPAN destination ports might be err-disabled because of a sequence timeout. You will see the following message when you try to bring the port up:

  ```
  ETHPORT-5-IF_SEQ_ERROR  Error ("sequence timeout") communicating with MTS_SAP_ETH_SPAN
  for opcode MTS_OPC_ETHPM_PORT_PRE_CFG (RID_PORT: EthernetX/Y);
  ETHPORT-5-IF_DOWN_ERROR_DISABLED  Interface EthernetX/Y is down (Error disabled.
  Reason:sequence timeout)
  ```

  Conditions: The conditions for this error are unknown.

  Workaround: This issue is resolved.

- CSCue74142

  Symptom: Provide a PCIe health monitoring test on the supervisor module for Cisco Nexus 7700 Series switches.

  Conditions: This is an enhancement request.

  Workaround: This enhancement request is resolved.

- CSCuj80949

  Symptom: You might see high CPU usage on the pktmgr process and a consequent drop of ARP packets as well as slow or no responses for ARP requests. This issue starts a timeout on local devices, and traffic is not sent because of incomplete ARP requests.

  Conditions: This symptom might be seen when you are using OTV unicast mode and there is a constant rate or constant burst rate of broadcast ARP traffic sent from remote sites to the ADJ server (approximately 1000 ARP requests per second).

  Workaround: This issue is resolved.

- CSCue7179

Symptom: VDC creation status is pending during an automatic upgrade of EPLD images.

Conditions: This symptom might be seen when you are upgrading EPLD images for I/O modules.

Workaround: This issue is resolved.

- CSCue46437

  Symptom: GOLD failures appear for modules that are under high traffic congestion.

  Conditions: This symptom might be seen when you have traffic congestion.

  Workaround: This issue is resolved.

- CSCul53679

  Symptom: Port set goes into error-disabled state with the following error seen in the default VDC:

  ```
  %MODULE-2-MOD_SOMEPORTS_FAILED: Module x (Serial number: xxxxxxxx) reported failure on
  ports Ethernet x/y-z (Ethernet) due to error in device DEV_CLP_FWD (device error
  0xca802600)
  ```

  Conditions: This symptom might be seen when the default VDC accounting log reports an L2LU IG_SA result merge fifo full interrupt.

  Workaround: This issue is resolved.

- CSCuh57710

  Symptom: The MAC address is installed to vPC+ SWID.LID instead of to a local vPC+ interface.

  Conditions: This symptom might be seen in the vPC+/FP environment. The problem is triggered after an TCN that causes flash in vPC+ domain is received after entries are flushed. Some entries are not correctly installed toward the local vPC+ channels but are installed on the SWID.SWID.LID of local vPC+ channel.

  Workaround: This issue is resolved.

- CSCuj35561

  Symptom: The BFD state displays ADMIN DOWN on BFD-enabled interfaces while other protocols on the same interface are up and working. When in this state, the BFD will not report connectivity as expected. Once deadlocked in this manner, it is possible that a restart message from the module could restart BFD causing link and protocol flaps.

  Conditions: This symptom might be seen when you are working with BFD.

  Workaround: This issue is resolved.

- CSCuj40617

  Symptom: When the configuration for an F2e Series module is lost, and you cannot modify it, the following error logs appear:

  ```
  2013 Nov 29 16:18:03 NX7KD-S1F5R9 %$ VDC-2 %$ %VMM-2-VMM_SERVICE_ERR: VDC2: Service
  SAP Qosmgr SAP for slot 4 returned error 0x41170014 (Operation timed out) in if_bind
  sequence
  2013 Nov 29 19:30:32 F340.12.06-7000-1 %IM-3-IM_RESP_ERROR: Component MTS_SAP_VMM
  opcode:MTS_OPC_IM_IF_VDC_BIND in vdc:2 returned error:Operation timed out
  ```

```
2013 Nov 29 19:30:32 F340.12.06-7000-1 %VDC_MGR-3-VDC_ERROR: vdc_mgr: Error for port
Ethernet4/44. Port is currently in vdc NX7KD-S1F5R9 [2]. GIM returned 41170014
[Operation timed out]. Please run the command "allocate interface Ethernet4/44 force"
to try again
```

Conditions: This symptom might be seen when the switch is rebooted and an F2e Series module is allocated to a non-admin VDC.

Workaround: This issue is resolved.

- CSCuj42378

  Symptom: The port-profile manager (PPM) might go down.

  Conditions: This symptom might be seen when you are working on the device using Cisco NX-OS 6.2(2).

  Workaround: This issue is resolved.

- CSCuj50567

  Symptom: You might see the iftmc process go down on an F2 Series module.

  Conditions: This symptom might be seen during an ISSU upgrade.

  Workaround: This issue is resolved.

- CSCuh8353

  Symptom: The ARP entry for the server in the HSRP standby switch is not correct and is pointing the server IP to the SVI VMAC instead of the host's own MAC.

  Conditions: This symptom might be seen when you enable the **local-proxy-arp** command on the SVI and the HSRP standby switch sends out an ARP request to some host in that VLAN.

  Workaround: This issue is resolved.

- CSCuh99618

  Symptom: BFD packets sent from a host might crash the incoming module on the switch.

  Conditions: This symptom might be seen when a UDP stream with port 3784 (BFD) is forwarded by the switch.

  Workaround: This issue is resolved.

- CSCuj57803

  Symptom: You might see the dcos-telnetd process go down.

  Conditions: This symptom might be seen when you are running large commands such as **show tech**.

  Workaround: This issue is resolved.

- CSCui02155

  Symptom: Traffic Engineering Fast Re-Routing is not getting triggered when there is a BFD session failure. A corresponding interface down event does trigger FRR successfully.

Conditions: This symptom might be seen when FRR is configured and the protected link has BFD enabled. A BFD neighbor flap does not trigger an FRR event unless the physical interface itself goes down.

Workaround: This issue is resolved.

- CSCuj66487

  Symptom: The IPv6 /127 point-to-point subnet mask is not working.

  Conditions: This symptom might be seen when you configure a /127 subnet on a point-to-point link using Cisco NX-OS Release 6.(2).

  Workaround: This issue is resolved.

- CSCui09637

  Symptom: With PIM SSM at the switch's VRF, the RPF neighbor might be listed as 0.0.0.0 although the correct entry for the source is present in the RIB and PIM session with peer is up. You might also see RPF neighbor listed as A.B.C.D although there is no route in RIB to source or RIB's next-hop is D.C.E.F.

  Conditions: This symptom might be seen only when PIM SSM groups are defined with /32 mask in the SSM range configuration. After unicast convergence, their RPF neighbor might not get updated to be inline along with new the unicast routing topology.

  Workaround: This issue is resolved.

- CSCuj72919

  Symptom: PING traffic loss might be seen due to no ARP or broadcast leaving Layer 2 trunk ports.

  Conditions: This symptom might be seen when you perform an ISSU or ISSD with a system switchover.

  Workaround: This issue is resolved.

- CSCuj78025

  Symptom: You might see low throughput for COS4 traffic traversing from an M1 Series module to an F1 Series module.

  Conditions: This symptom might be seen when you are running QoS and the traffic is running from an M1 Series module to an F1 Series module.

  Workaround: This issue is resolved.

- CSCuj82708

  Symptom: Multicast packets from CE to FP might have the ftag=0 set, which brings about unpredictable forwarding at the next FP switch. Some packets are flooded while some packets are discarded at the next switch.

  Conditions: This symptom might be seen when you have peer-link and vPC legs in different modules, you bring down the peer-link and then bring down the keepalive link.

  Workaround: This issue is resolved.

- CSCuj98135

  Symptom: Proxy Layer 3 routing might be affected after configuring FabricPath proxy Layer 2 learning and the dynamic MAC entries are flushed from the MAC table. Unicast traffic sent from the supervisor or from an M Series module towards a FabricPath core port might be dropped because the FabricPath outer destination address is misprogrammed during encapsulation.

  Conditions: This symptom might be seen when you are running Cisco NX-OS Release 6.2(2) in a mixed chassis VDC where M1/M2 Series and F2e Series modules are used with proxy Layer 2 learning.

  Workaround: This issue is resolved.

- CSCui36584

  Symptom: TCP-dependent applications do not work as expected.

  Conditions: This symptom might be seen when multiple sockets are created and closed at the same time for clients that are multi-threaded such as BGP and MSDP.

  Workaround: This issue is resolved.

- CSCul00568

  Symptom: The M2 Series module might not boot up after reload when the QoS policy is attached to more than 512 VLANs.

  Conditions: This symptom might be seen when a QoS configuration is attached to more than 512 VLANs and either the switch is reloaded or the M2 Series module is reloaded. As the module boots up, the QoS configuration is loaded onto the module and the module times out while programming the hardware.

  Workaround: This issue is resolved.

- CSCul04756

  Symptom: HSRP peers do not see each other following a FabricPath (FP) topology configuration.

  Conditions: This symptom might be seen when you perform an ISSU from Cisco NX-OS Release 6.1 to Cisco NX-OS Release 6.2 with multiple topologies configured.

  Workaround: This issue is resolved.

- CSCul14107

  Symptom: You might see sequence timeouts and error-disabled ports if you are suspending one VDC while making configuration changes in another VDC.

  Conditions: This symptom might be seen when one VDC is being suspended, and in another session a configuration change (such as removing VLANs) is occurring in another VDC that shares the same module.

  Workaround: This issue is resolved.

- CSCul14407

  Symptom: The switch does not generate an IP unreachable message when packets are dropped due to an MTU fail check.

Conditions: This symptom might be seen when the hardware rate limiter is not correctly programmed on the module.

Workaround: This issue is resolved.

- CSCul22062

  Symptom: The convergence time for the OTV is highly impacted due to lacking RNH tracking, and RNH tracking, does not happen between pairs of VDC OTV.

  Conditions: This symptom might be seen when you reload VDC OTV or when the join-interface goes up and down on the VDC OTV.

  Workaround: This issue is resolved.

- CSCuj05809

  Symptom: The switch generates area aggregation routes with the wrong cost.

  Conditions: This symptom might be seen when the component routes for area aggregation are delivered with same the LSID, but with different costs from different devices. The switch applies the cost from the device with the lower router ID, not the bestpath cost.

  Workaround: This issue is resolved.

- CSCuj56217

  Symptom: The ARP packet is still redirected to the CPU and processed without creating an ARP cache even if you issue the **no otv suppress-arp-nd** command under the overlay interface.

  Conditions: This symptom might be seen when **no otv suppress-arp-nd** is configured, and at the same time at least one extended VLAN is not forwarding on any port in OTV VDC.

  Workaround: This issue is resolved.

- CSCua53069

  Symptom: The origin AS in FIB and RIB are different, which might result in a wrong NF export value.

  Conditions: This symptom might be seen when some BGP routes have the same next hop.

  Workaround: This issue is resolved.

- CSCub34109

  Symptom: You might see an error message or timeout when several commands are executed.

  Conditions: This symptom might be seen with any command used to check the NetFlow statistics.

  Workaround: This issue is resolved.

- CSCuj22757

  Symptom: The broadcast reply packet to a relay agent is dropped if the giaddr does not match the relay agent IP address.

Conditions: This symptom might be seen when the broadcast flag is set and reply (offer and ack) packets from server only. The first relay agent floods the packet correctly because the giaddr is its own. The packet is dropped when it is received by the second relay agent.

Workaround: This issue is resolved.

- CSCuj05880

  Symptom: You might see an unexpected module reboot with following message:

  ```
  %SYSMGR-SLOT3-4-SYSMGR_CORE_TRUNCATED: Core seems to be truncated on generation. 79448
  / 145136 KB. PID: 2122
  %SYSMGR-SLOT3-2-SERVICE_CRASHED: Service "val_usd" (PID 2122) hasn't caught signal 11
  (core will be saved).
  %MODULE-2-MOD_DIAG_FAIL: Module 3 (Serial number: JAF1717AQCJ) reported failure due to
  Service on linecard had a hap-reset in device DEV_SYSMGR (device error 0x2da)
  ```

  Conditions: This symptom might be seen when you are running the Cisco NX-OS Release 6.1.

  Workaround: This issue is resolved.

- CSCuj59174

  Symptom: The FEX module might still be sending an all 0 source MAC address after you perform an ISSU upgrade to Release 6.2(2) or Release 6.2(2a).

  Conditions: This symptom might be seen when you perform an ISSU upgrade to Release 6.2(2) or Release 6.2(2a) from a previous image.

  Workaround: This issue is resolved.

- CSCuj59684

  Symptom: No value is returned to the NMS when you are polling for the FEX interface connected to the switch for duplex settings.

  Conditions: This symptom might be seen when you are using the MIB dot3StatsDuplexStatus.

  Workaround: This issue is resolved.

- CSCug62922

  Symptom: When LLDP is disabled on a per-port basis, PFC stays open after the port goes up and down when LLDP is disabled.

  Conditions: This symptom might be seen when DCBX has negotiated PFC and the PFC mode is configured as auto.

  Workaround: This issue is resolved.

- CSCug63304

  Symptom: Invalid router LSAs installed in the OSPF database and flooded to neighbors.

  Conditions: This symptom might be seen when an OSPF neighbor is sending an invalid LSA.

  Workaround: This issue is resolved.

- CSCug93264

  Symptom: You might see a high traffic rate on the inters-witch links connecting two switches; one link is for regular CE VLANs and the other is for FabricPath.

  Conditions: This symptom might be seen when one of the existing CE VLANs is converted into a FabricPath VLAN on the switch that is the STP root.

  Workaround: This issue is resolved.

- CSCuh07613

  Symptom: You issued either the **shutdown** or **no hsrp** command and the HSRP VIP is down, but the route is still installed.

  Conditions: This symptom might be seen when a new SVI is created and HSRP is configured while you are running Cisco NX-OS Release 6.2(2) or 6.2(2a).

  Workaround: This issue is resolved.

- CSCuh12757

  Symptom: The route-map **set** commands can be applied incorrectly during export map evaluation.

  Conditions: This symptom might be seen if an export map is configured without an inbound route map. A previously evaluated export map **set** commands then can be erroneously applied to the prefix currently being evaluated.

  Workaround: This issue is resolved.

- CSCuh30369

  Symptom: You might see the switch PTP clock off from the GM by 35 seconds.

  Conditions: This symptom might be seen when the Announce Message sent from the switch always has the wrong flag bits.

  Workaround: This issue is resolved.

- CSCuh36180

  Symptom: The output of the **show hardware flow utilization module** command shows only the statistics for instance 1.

  Conditions: This symptom might be seen when you enter the **show hardware flow utilization module** command to collect statistics for an instance other than 1.

  Workaround: This issue is resolved.

- CSCuh44586

  Symptom: You might be unable to add new members to an existing port-channel interface.

  Conditions: This symptom might be seen after you perform an ISSU to Cisco NX-OS Release 6.1(3).

  Workaround: This issue is resolved.

- CSCuh44908

  Symptom: You might see the NFM application unexpectedly go down and write out a core file when the switch is being configured with an interface configuration.

  Conditions: This symptom might be seen during a dynamic configuration of an Ethernet interface.

  Workaround: This issue is resolved.

- CSCuh50150

  Symptom: You might see a downstream Nexus 5000 switch with a Layer 3 daughter card unable to ping GLBP VIP.

  Conditions: This symptom might be seen when the Nexus 7000 switches are in a vPC configuration with GLBP as the FHRP and FabricPath is enabled with downstream leaf devices. This situation occurs only if the Nexus 5000 switch has an installed Layer 3 daughter card.

  Workaround: This issue is resolved.

- CSCui94802

  Symptom: You might see the absolute timeout and logout warning configuration listed before the line configuration in the running config.

  Conditions: This symptom might be seen at any point.

  Workaround: This issue is resolved.

- CSCuh61950

  Symptom: You might see ports as a hardware failure when you issue the **show interface** command when you add ports from an M2 Series module from the default VDC to another VDC.

  Conditions: This symptom might be seen you are working with M2 Series modules and new VDCs.

  Workaround: This issue is resolved.

- CSCul45084

  Symptom: You might see Netstack go down or become unstable when you issue the **show ip int brief include-secondary** command.

  Conditions: This symptom might be seen you are using ip unnumbered interfaces, such as multicast tunnel interfaces.

  Workaround: This issue is resolved.

- CSCui05605

  Symptom: For GLC-T SFPs in a switch with N7K-M148GS-11 LC, the interface might remain up when the peer side interface is admin shut from the CLI or the copper cable was removed. This happens when the speed is hard-coded for ?speed 1000."

  Conditions: This symptom might be seen when the interface is configured for ?speed 1000."

  Workaround: This issue is resolved.

- CSCuj73049

  Symptom: A module might go down due to a memory leak, and a core file is generated after the crash. The following error messages are observed every 3 minutes after the module goes down and comes back online:

  ```
  SYSMGR-SLOT3-4-VAR_SYSMGR_FULL System core file storage usage is unexpectedly high at
  100%. This might cause corruption of core files
  SYSMGR-SLOT3-2-CORE_SAVE_FAILED core_client_main: PID 27674 with message sysmgr_logmgr
  script fails
  SYSMGR-SLOT3-5-SUBPROC_TERMINATED  "System Manager (core-client)" (PID 27674) has
  finished with error code SYSMGR_EXITCODE_CORE_CLIENT_ERR (11).
  ```

  Conditions: This symptom might be seen on any of the modules.

  Workaround: This issue is resolved.

- CSCui26744

  Symptom: Static port security MAC address might be programmed as a drop-in MAC address table if the switch receives traffic with a source MAC address from the static port security MAC address, and it is received on other ports.

  Conditions: This symptom might be seen in Cisco NX-OS Release 6.1(3).

  Workaround: This issue is resolved.

- CSCuj30289

  Symptom: When you enter-the **sh run port-profile** *name* command to display the configuration of a port profile, this command appears to be case insensitive.

  Conditions: This symptom might be seen with port-profile names with identical characters that are differentiated only by the use of capital letters.

  Workaround: This issue is resolved.

- CSCui67227

  Symptom: Syslog reports the wrong hostname after a switchover.

  Conditions: This symptom might be seen when a standby supervisor that is becoming active was just moved from a different chassis.

  Workaround: This issue is resolved.

- CSCui90226

  Symptom: When a dynamic ARP inspection is enabled on a particular VLAN, all the of the ARP-inspected packets are hitting copp-system-p-class-normal instead of copp-system-p-class-redirect.

  Conditions: This symptom might be seen when you enable dynamic ARP inspection.

  Workaround: This issue is resolved.

- CSCuj80898

  Symptom: The EPLD upgrade might fail on module N7K-M148GS-11L.

Conditions: This symptom might be seen when the hardware version of the N7K-M148GS-11L module is 2.0 or higher.

Workaround: This issue is resolved.

- CSCui96609

  Symptom: If you authenticate using TACACS and the server permits all commands for your role, you might see the following error message when attempting to create a new user:

  ```
  Secure password mode is enabled. Please use "change-passwd" CLI to change the
  password.
  ```

  Conditions: This symptom might be seen when the user has been authenticated with a TACACS server, the server permits all commands for this user, and the user is in the role of vdc-operator.

  Workaround: This issue is resolved.

- CSCuj95140

  Symptom: You might see an SSH vulnerability on the switch.

  Conditions: This symptom might be seen when you are working with a Cisco NX-OS release prior to Release 6.2(5).

  Workaround: This issue is resolved.

- CSCul00757

  Symptom: Each switch in a subnet appears as IGMP active querier although only the lowest IP address in the subnet is supposed to perform this role. High CPU might occur. A large number of IGMP Queries might be sent and/or received by the switch in one or more VLANs.

  Conditions: This symptom might be seen if you are working with four switches (two vPC pairs) connected to one another through Layer 2.

  Workaround: This issue is resolved.

- CSCul02492

  Symptom: The output from the **show run port-profile** command might be truncated after you add a VLAN to a port profile.

  Conditions: This symptom might be seen if you are adding a VLAN to a port profile.

  Workaround: This issue is resolved.

- CSCul04790

  Symptom: When you are working with Cisco NX-OS Release 6.1(4a), the BGP session on the active BGP session will not be torn down until BGP session reset/restart if the password is mismatched on both ends in either of the following scenarios (this will cause an inconsistent state):

  - One side has no passwd, another side has a password.

  - The password on both ends are not a match.

  Conditions: This symptom might be seen if you are working with Cisco NX-OS Release 6.1.(2).

  Workaround: This issue is resolved.

- CSCuj52700

  Symptom: An additional next hop appears for the default route when using a /0 subnet mask.

  Conditions: This symptom might be seen if you have incorrect configurations.

  Workaround: This issue is resolved.

- CSCul09672

  Symptom: Although the switches currently do not support SNMPv3 informs, the CLI allows users to configure it.

  Conditions: This symptom might be seen if you are configuring SNMPv3 informs.

  Workaround: This issue is resolved.

- CSCuj53818

  Symptom: When changes are made to a range of ports inheriting a port-profile configuration, the range operation does not work as expected.

  Conditions: This symptom might be seen when you are using interface range option.

  Workaround: This issue is resolved.

- CSCtr07662

  Symptom: The service attribute in Cisco NX-OS exec shell accounting requests and command accounting requests are set to TAC_PLUS_AUTHEN_SVC_NONE. (Set it to login instead to match what Cisco IOS sends.)

  Conditions: This symptom might be seen when you are running Cisco NX-OS.

  Workaround: This enhancement request is resolved.

- CSCuh95271

  Symptom: The **show cdp neighbor** command for interfaces in a port channel might indicate that the neighbor is not IGMP capable.

  Conditions: This symptom might be seen when you are working with interfaces in a port channel.

  Workaround: This issue is resolved.

- CSCuh97596

  Symptom: Hardware failures that cause a traffic impact should be sent to the logfile, but these failures are displayed only on the exception log.

  Conditions: This symptom might be seen if you have a hardware failure.

  Workaround: This issue is resolved.

- CSCui90093

  Symptom: The telnet session disconnects when you issue following commands:

- **show fex x version**

- **show system reset-reason fex x**

- **attach fex x** and then **show version**

Conditions: This symptom might be seen after replacing FEX.

Workaround: This issue is resolved.

- CSCui60491

  Symptom: You cannot save the running configuration.

  Conditions: This symptom might be seen when the /mnt/pss/ directory is at 100 percent capacity on either supervisor.

  Workaround: This issue is resolved.

- CSCul26607

  Symptom: The CoPP logging, which generally logs once when the said violation threshold is violated, is repeated every 5 minutes even if there are no new violations for the class after upgrading to Cisco NX-OS Release 6.2(2) using ISSU.

  Conditions: This symptom might be seen after you perform an ISSU upgrade to Cisco NX-OS Release 6.2(2).

  Workaround: This issue is resolved.

- CSCuj56956

  Symptom: The switch sends fast PDUs when you issue the **lacp rate fast** command. Transmissions should occur at a rate determined by the partner.

  Conditions: This symptom might be seen after you issue the **lacp rate fast** command.

  Workaround: This issue is resolved.

- CSCuc06762

  Symptom: When you issue the **show diagnostic content module** *module #* command, the switch shows tests running that are not supported on that module.

  Conditions: This symptom might be seen after you issue the command **show diagnostic content module** *module #* command.

  Workaround: This issue is resolved.

- CSCuj58189

  Symptom: The FEX type N2248TP-E is changed to N2348GTP in the running configuration.

  Conditions: This symptom might be seen after you upgrade to Cisco NX-OS Release 6.2(2).

  Workaround: This issue is resolved.

- CSCuj82155

Symptom: When configuring the port security on a switch port, the SNMP polling and CLI return an incorrect status after the interface is error-disabled due to a security violation (a different MAC than the configuration MAC being learned).

Conditions: This symptom might be seen when an interface is error-disabled because of a security violation.

Workaround: This issue is resolved.

- CSCue69943

  Symptom: You cannot configure the MAC address aging time to a 5-second timer value. This is possible on the Catalyst 6000 switch and the Nexus 5000 switch.

  Conditions: This symptom might be seen when you are working with the MAC address aging time.

  Workaround: This enhancement request is resolved.

- CSCuf90519

  Symptom: The switch cannot handle RADIUS packets larger than 4k.

  Conditions: This symptom might be seen when you are working with RADIUS.

  Workaround: This issue is resolved.

- CSCuf94072

  Symptom: When you are working with COPP, there is no command that allows the user to troubleshoot COPP drops quickly.

  Conditions: This symptom might be seen when you are working with COPP.

  Workaround: This enhancement request is resolved.

- CSCtc97478

  Symptom: The switch does not refresh policies periodically, and it should comply with the 'Download SGACL lists' timer sent from ACS or ISE.

  Conditions: This symptom might be seen when CTS policies have been downloaded to the switch.

  Workaround: This issue is resolved.

- CSCte69784

  Symptom: The switch does not indicate if the logged ACL was a permit or deny entry, which is different from other Cisco platforms.

  Conditions: This symptom might be seen when you are working with ACLs.

  Workaround: This enhancement request is resolved.

- CSCul80399

  Symptom: Netstack goes down if there are changes to the PBR.

  Conditions: This symptom might be seen when you are working with PBR.

  Workaround: This enhancement request is resolved.

- CSCum13080

  Symptom: Packets coming to the supervisor from a F3 line card that has FabricPath edge ports configured might get dropped.

  Conditions: This symptom might be seen if an F3 port initially configured as a FabricPath edge port is reconfigured as a core port. The Cisco NX-OS software does not recognize this change and might drop incoming packets from this port to the supervisor.

  Workaround: Restart the supervisor.

- CSCtx11656

  Symptom: Route redistribution fails and the following syslog message appears:

  ```
  %EIGRP-3-RPM_LIB_API_FAILED: bgp_lookup_ext_attr() - failed in
  rpm_acquire_bgp_shmem_lock()
  ```

  Conditions: This symptom might be seen when route redistribution from BGP to EIGRP is configured using community lists.

  Workaround: None.

Resolved Caveats—Cisco NX-OS Release 6.2(2a)

- CSCue05555

  Symptom: The satctrl service might fail on a FEX after several switchovers.

  Conditions: This symptom might be seen during a switchover. Some timers might not be correctly cleaned up during the switchover. As a result, a bad timer can be released, which might cause the FEX to fail.

  Workaround: This issue is resolved.

- CSCuh76946

  Symptom: MAC addresses on Cisco Nexus 7000 Series devices are sometimes not consistent in hardware and software, which causes a firmware flush.

  Conditions: This symptom might be seen after one of the following events:

  - MAC address moves
  - Online insertion and removal (OIR) of switch modules
  - In-service software upgrade or in-service software downgrade on the switch

  Workaround: This issue is resolved.

- CSCui30261

  Symptom: The pltfm_config software module on Cisco Nexus 7000 Series devices sometimes fails when you enter the **show running-conf** command on the default virtual device context (VDC).

  Conditions: This symptom might be seen after multiple in-service software upgrades (ISSUs) to Release 6.1(x).

  Workaround: This issue is resolved.

- CSCui33523

  Symptom: A secure shell (SSH) connection is established with a logical interface (port channel) that is down.

  Conditions: This symptom might be seen when the SSH packet causes an ICMP redirect message to be sent, and the incoming and outgoing port are the same.

  Workaround: This issue is resolved.

- CSCui39061

  Symptom: On a Cisco Nexus 7000 Series device that is running Cisco NX-OS Release 5.2(5), the supervisor sometimes restarts when the ipqosmgr process fails.

  Conditions: This symptom might be seen on Cisco Nexus 7000 Series devices running Cisco NX-OS Release 5.2(5).

  Workaround: This issue is resolved.

- CSCui58446

  Symptom: In a mixed-chassis setup on Cisco Nexus 7000 Series devices, with an M Series module and an F2e module in the same virtual device context (VDC), the vPC does not start when using any VLAN greater than 4040, and the switch virtual interface (SVI) remains down.

  Conditions: This symptom might be seen in a mixed-chassis setup on Cisco Nexus 7000 Series devices with an M Series module and an F2e module in the same VDC.

  Workaround: This issue is resolved.

- CSCui63317

  Symptom: Cisco Nexus 7000 Series devices report the following errors, and the VSH process fails:

  ```
  2013 Aug 10 14:13:30 N7k vsh[27390]: CLI-4-WARN_OUT_OF_MEMORY: Out of memory
  ../feature/vsh/cli/cli_common/cli_tlv.cc:642
  2013 Aug 10 14:13:30 N7k last message repeated 8 times
  2013 Aug 10 14:13:30 N7k vsh[27390]: CLI-4-WARN_OUT_OF_MEMORY: Out of memory
  ../feature/vsh/cli/cli_common/cli_tlv.cc:675
  ```

  Conditions: This symptom might be seen when AAA authorization is enabled. The **show** commands sometimes cause a memory leak, which leads to CLI-4-WARN_OUT_OF_MEMORY errors and causes the VSH process to fail.

  Workaround: This issue is resolved.

- CSCuj06468

  Symptom: After an upgrade to Cisco NX-OS Release 6.2(2), DHCP relay sometimes fails on Cisco Nexus 7000 Series devices for DHCP requests when the source User Datagram Protocol (UDP) port is not bootpc (port 68).

  Conditions: This symptom might be seen for DHCP packets sent using a non-bootp UDP source port and occurs only in Release 6.2(2).

  Workaround: This issue is resolved.

- CSCuj07925

  Symptom: OSPF interface commands do no take effect after a reboot.

  Conditions: This symptom might be seen under the following conditions:

  - A Cisco Nexus 7000 Series switch with a Supervisor 2E module is running Cisco NX-OS Release 6.1(2) or Release 6.2(2).
  - The **passive-interface default** command is globally configured for a router OSPF process.
  - An interface is configured with the **no ip ospf passive-interface** command and the **ip ospf network point-to-point** command.
  - The affected interface might incorrectly appear as passive and in the default broadcast mode.

  Workaround: This issue is resolved.

- CSCuj10728

  Symptom: Error reporting on the Cisco Nexus 7000 32-port 10-Gigabit Ethernet SFP+ I/O module (N7K-M132XP-12) should be improved.

  Conditions: The output of the **show interface** command should show packet corruption within internal ASICs. Polling of the errors with SNMP should also be possible.

  Workaround: This enhancement request is resolved.

- CSCuj22930

  Symptom: On a Cisco Nexus7000 Series device, packets might be dropped in a FabricPath topology after a reload or bootup of an F1 Series module that is running Cisco NX-OS Release 6.2(2). This issue is not intermittent, and flows that are interrupted cannot pass any packets.

  Conditions: This symptom might be seen under the following conditions:

  - Cisco NX-OS Release 6.2(2) is running.
  - There is a FabricPath topology.
  - An F1 Series module is configured as a leaf or edge port module.
  - The F1 Series module reloads.

  Workaround: This issue is resolved.

- CSCuj24572

  Symptom: On a Cisco Nexus 7000 Series device running Cisco NX-OS Release 6.2(2), broadcast frames coming from the peer link might not be forwarded to host ports on a Cisco Nexus 2000 Series fabric extender (FEX), which causes incomplete ARP entries when the FEX is not connected.

  Conditions: This symptom might be seen on Cisco Nexus 7000 Series devices using module type N7K-F248XP-25 or N7K-F248XP-25E after the module or the chassis reloads. However, after a nondisruptive ISSU, this issue does not occur until the module reloads.

  Workaround: This issue is resolved.

- CSCuj25197

Symptom: A Cisco Nexus 7000 Series device sometimes forwards packets across virtual device contexts (VDCs) when it is configured as an emulated switch in a VDC with F2 modules, and with another VDC using M1 modules. The leaking packet is received over a virtual port channel. On the peer, it is received over the peer link on an F2 module and picked up by the M1 modules in another VDC.

Conditions: This symptom might be seen when the Cisco Nexus 7000 Series device is equipped with a supervisor 2 or 2E.

These are the triggers for this issue:

- After the emulated switch is configured, some ports are allocated or deallocated from the VDC.
- The problem can occur after a FabricPath topology change (clearing adjacencies, peer-link flap).

Workaround: This issue is resolved.

Resolved Caveats—Cisco NX-OS Release 6.2(2)

- CSCub54436

  Symptom: Prior to Cisco NX-OS Release 6.2(2), the spanning tree TCN is transmitted from Nexus 7000 Series device when a new **otv extend-vlan** command is configured on an overlay interface and the edge device gets selected as an AED. From Cisco NX-OS Release 6.2(2) onwards the TCN transmission does not happen.

  Conditions: When a new **otv extend-vlan** command is configured and the edge device is selected as an AED.

  Workaround: None

- CSCtf36357

  Symptom: A Cisco Nexus 7000 Series device does not support having ingress NetFlow sampling and DHCP relay configured on the same interface.

  Conditions: This symptom might be seen under normal operating conditions for a Cisco Nexus 7000 Series device.

  Workaround: This issue is resolved.

- CSCts51026

  Symptom: When tacac+ source-interface configuration is present, Small memory leak is seen in libipconf for each tacacs+ authentication and authorization request.

  Conditions: This can occur only if tacacs+ source-interface configuration is present.

  Workaround: Disabling and enabling tacacs+ service will recover the memory that is leaked.

- CSCue57018

  Symptom: An ISSU to Cisco NX-OS Release 6.2(2) from a release earlier than Release 6.1(3) becomes blocked.

  Conditions: This symptom might be seen under the following conditions:

  - There is an ISSU to Cisco NX-OS Release 6.2(2) from a release earlier than Release 6.1(3).

  - There is a VACL configuration in at least one VDC. The VACL can be active or inactive.

  Workaround: This issue is resolved.

- CSCug37851

  Symptom: A Cisco Nexus 7000 Series device might experience SNMP timeouts when using bulk Get requests.

  Conditions: This symptom might be seen with Bridge and Entity MIBs, especially when FEX modules are in use.

  Workaround: This issue is resolved.

- CSCug56477

  Symptom: Web Cache Control Protocol (WCCP) redirection does not work when F2e Series modules are used.

  Conditions: This symptom might be seen when using an M1 Series and an F2e Series module in a mixed VDC.

  Workaround: This issue is resolved.

- CSCuh19881

  Symptom: If a VLAN mapping of X to A exists, attempts to overwrite the mapping with the **otv vlan mapping X to B** command are rejected. In addition, attempts to copy a configuration to the running configuration are rejected if mapping conflicts exist.

  Conditions: This symptom might be seen when a mapping already exists and the same VLAN is used to map a replaced configuration command.

  Workaround: This issue is resolved.

- CCui22809

  Symptom: When you perform an ISSU to a Cisco NX-OS Release 6.2(2), a reload of the 32-port 10-Gigabit Ethernet SFP+ I/O module (N7K-M132XP-12) is required to enable the new DSCP based queuing functionality in Release 6.2(2).

  Conditions: This symptom might be seen following an ISSU to Cisco NX-OS Release 6.2(2).

  Workaround: This issue is resolved.

- CSCui43540

  Symptom: A random failure occurs with Layer 2 VPN.

  Conditions: This symptom might be seen when a remote provider edge (PE) device is going through ISSU and has VPWS and VPLS configured.

  Workaround: This issue is resolved.

Upgrade and Downgrade

To perform a software upgrade or downgrade, follow the instructions in the Cisco Nexus 7000 Series NX-OS Software Upgrade and Downgrade Guide, Release 6(x). For information about an In Service Software Upgrade (ISSU), see https://www.cisco.com/c/dam/en/us/td/docs/dcn/tools/nexus-7k-issu-matrix/index.html

Related Documentation

Cisco NX-OS documentation is available at the following URL:

http://www.cisco.com/en/US/products/ps9372/tsd_products_support_series_home.html

The Release Notes for upgrading the FPGA/EPLD is available at the following URL:

http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_1/epld/epld_rn.html

Cisco NX-OS includes the following documents:

Release Notes

Cisco Nexus 7000 Series NX-OS Release Notes, Release 6.x

NX-OS Configuration Guides

Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide

Cisco Nexus 7000 Series NX-OS Configuration Examples

Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide

Configuring Feature Set for FabricPath

Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide

Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide

Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide

Cisco Nexus 7000 Series NX-OS IP SLAs Configuration Guide

Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide

Cisco Nexus 7000 Series NX-OS LISP Configuration Guide

Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide

Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide

Cisco Nexus 7000 Series NX-OS OTV Configuration Guide

Cisco Nexus 7000 Series OTV Quick Start Guide

Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide

Cisco Nexus 7000 Series NX-OS SAN Switching Configuration Guide

Cisco Nexus 7000 Series NX-OS Security Configuration Guide

Cisco Nexus 7000 Series NX-OS System Management Configuration Guide

Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide

Cisco Nexus 7000 Series NX-OS Verified Scalability Guide

Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide

Cisco Nexus 7000 Series NX-OS Virtual Device Context Quick Start

Cisco NX-OS FCoE Configuration Guide for Cisco Nexus 7000 and Cisco MDS 9500

NX-OS Command References

Cisco Nexus 7000 Series NX-OS Command Reference Master Index

Cisco Nexus 7000 Series NX-OS FabricPath Command Reference

Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference

Cisco Nexus 7000 Series NX-OS High Availability Command Reference

Cisco Nexus 7000 Series NX-OS Interfaces Command Reference

Cisco Nexus 7000 Series NX-OS IP SLAs Command Reference

Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference

Cisco Nexus 7000 Series NX-OS LISP Command Reference

Cisco Nexus 7000 Series NX-OS MPLS Command Reference

Cisco Nexus 7000 Series NX-OS Multicast Routing Command Reference

Cisco Nexus 7000 Series NX-OS OTV Command Reference

Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference

Cisco Nexus 7000 Series NX-OS SAN Switching Command Reference

Cisco Nexus 7000 Series NX-OS Security Command Reference

Cisco Nexus 7000 Series NX-OS System Management Command Reference

Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference

Cisco Nexus 7000 Series NX-OS Virtual Device Context Command Reference

Cisco NX-OS FCoE Command Reference for Cisco Nexus 7000 and Cisco MDS 9500

Other Software Document

Cisco NX-OS Licensing Guide

Cisco Nexus 7000 Series NX-OS MIB Quick Reference

Cisco Nexus 7000 Series NX-OS Software Upgrade and Downgrade Guide

Cisco NX-OS System Messages Reference

Cisco Nexus 7000 Series NX-OS Troubleshooting Guide

Cisco NX-OS XML Interface User 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.