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

Date: June 3, 2014
Part Number: OL-29671-05 IO
Current Release: 6.2(8)

This document describes the features, caveats, and limitations for Cisco NX-OS software for use on the Cisco Nexus 7000 Series. Use this document in combination with documents listed in the "Related Documentation" section on page 142.

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

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

Table 1 shows the online change history for this document.

Table 1 Online History Change

Part Number	Revision	Date	Description
OL-29671-01	A0	August 22, 2013	Created release notes for Release 6.2(2).
	B0	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.
	C0	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.
	D0	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.
	E0	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.
	F0	September 16, 2013	Updated vPC and vPC+ in the "New Software Features" section.
	G0	September 18, 2013	Added "Upgrade With an M2 Series Module Installed" to the "Upgrade/Downgrade Paths and Caveats" section.
	H0	September 30, 2013	Updated FabricPath in the "New Software Features" section.

Table 1 Online History Change (continued)

Part Number	Revision	Date	Description
OL-29671-02	A0	October 16, 2013	Created release notes for Release 6.2(2a).
	B0	October 21, 2013	Added the Cisco Nexus B22 Fabric Extender for HP (N2K-B22HP) to Table 2 and Table 3.
	C0	October 22, 2013	Updated the "Security Features" section of the "Cisco NX-OS Release 6.2(2) Software Features" section.
	D0	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).
	E0	November 21, 2013	Updated the "Limitations" section to add "Behavior of Control Plane Packets on an F2e Series Module".
	F0	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.
	G0	December 10, 2013	Updated Table 9 to add support for Fe and F2e Series modules hardware as ERSPAN destinations.
OL-29671-03	A0	December 20, 2013	Created the release notes for Release 6.2(6).
	B0	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)
	C0	January 7, 2014	Added Physical Port vPC to the new features for Release 6.2(6).
	D0	January 7. 2014	Updated the "OTV Support"section and the "OTV" section under New Software Features for Release 6.2(6).
	E0	January 8, 2014	Added the Cisco Nexus B22HP Fabric Extender for HP (N2K-B22HP) to the "Cisco Nexus Fabric Extenders" section.
	F0	January 4, 2014	Updated the ISSU and ISSD information in the "Upgrade/Downgrade Paths and Caveats" section.
	G0	January 22, 2014	Updated the "Open Caveats—Cisco NX-OS Release 6.2" section to add CSCul91443.
	H0	January 23, 2014	Updated the "Interoperability Between Modules" information.
	I0	January 24, 2014	Updated the Updated the "Limitations" section to add "The no hardware ejector enable Command is Not Recommended for Long-Term Use"

Table 1 Online History Change (continued)

Part Number	Revision	Date	Description
	J0	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.
	K0	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.
	L0	February 20, 2014	• Updated "Non-ISSU Upgrade Steps" section.
OL-29671-04	A0	February 26, 2014	Created the release notes for Release 6.2(6a).
	BO	March 10, 2014	Revised Table 5.
	CO	April 2, 2014	Removed support for ERSPAN destination for F3 Series modules.
	DO	April 3, 2014	Revised Tables 2 and 3.
	EO	April 8, 2014	Correct Cisco NX-OS Release to 6.2(6) for CLI command for breakout capabilities.
OL-29671-05	AO	April 25, 2014	Created the release notes for Release 6.2(8).
	B0	April 26, 2014	Added two additional caveats to the "Open Caveats—Cisco NX-OS Release 6.2" section.
	C0	April 28, 2014	Added additional caveats to the "Caveats" section.
	D0	May 1, 2014	Added additional caveats to the "Caveats" section.
	E0	May 6, 2014	Added additional caveat to the "Open Caveats—Cisco NX-OS Release 6.2"section.
	F0	May 12, 2014	Updated information for LISP in "Cisco NX-OS Release 6.2(8) Software Features" section.
	G0	May 13, 2014	Added list of modules and required level of software to "CLI Command for Breakout Capabilities" section.
	H0	May 22, 2014	Added additional caveat to the "Caveats" section.
	I0	June 3, 2014	Updated Table 9.

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

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
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)
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)
N77-C7718-FAN	Fan, Cisco Nexus 7718 chassis	6.2(2)
N77-C7710-FAN	Fan, Cisco Nexus 7710 chassis	6.2(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)
N7K-C7018-FAN	Fan tray for the Cisco Nexus 7018 chassis	4.1(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)
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)
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
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)
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)
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)
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)
N7K-F312FQ-25	Cisco Nexus 7000 12-port 40-Gigabit Ethernet QSFP+ 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)
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)
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	6.2(6)
	N2K-C2248TP-1GE	
	N2K-C2232PP-10GE	
	N2K-C2232TM	
	N2K-C2248TP-E	
	N2K-C2232TM-E	
	N2K-C2248PQ	
	N2K-B22HP[1]	
24-port Cisco Nexus 7700 F3 Series 40-Gigabit Ethernet QSFP I/O module (N77-F324FQ-25)	N2K-C2224TP-1GE	6.2(8)
	N2K-C2248TP-1GE	
	N2K-C2232PP-10GE	
	N2K-C2232TM	
	N2K-C2248TP-E	
	N2K-C2232TM-E	
	N2K-C2248PQ	
	N2K-B22HP[1]	
48-port Cisco Nexus 7700 F3 Series 1/10-Gigabit Ethernet SFP+ I/O module (N77-F348XP-23)	N2K-C2224TP-1GE	6.2(6)
	N2K-C2248TP-1GE	
	N2K-C2232PP-10GE	
	N2K-C2232TM	
	N2K-C2248TP-E	
	N2K-C2232TM-E	
	N2K-C2248PQ	
	N2K-B22HP	
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	5.2(1)
	N2K-C2232PP-10GE	
	N2K-C2232TM	6.1(1)
	N2K-C2248TP-E	
	N2K-2232TM-E	6.2(2)
	N2K-C2248PQ	
	N2K-B22HP	

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	40GBASE-SR4 QSFP+	6.2(6)
	QSFP-40G-CSR4	40GBASE-CSR4 QSFP+	6.2(6)
	QSFP-40GE-LR4	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)
N77-F348XP-23	CWDM-SFP-xxxx[1]	1000BASE-CWDM	6.2(6)
	CWDM-SFP-1470	CWDM	6.2(8)
	CWDM-SFP-1530	CWDM	6.2(8)
	CWDM-SFP-1610	CWDM	6.2(8)
	DWDM-SFP-xxxx[1]	1000BASE-DWDM	6.2(6)
	DWDM-SFP-3033	DWDM	6.2(8)
	DWDM-SFP-4453	DWDM	6.2(8)

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

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	DWDM-SFP-6063	DWDM	6.2(8)
	FET-40G	Cisco Fabric Extender Transceiver (FET)	6.2(6)
	SFP-10G-SR	10GBASE-SR SFP+	6.2(6)
	SFP-10G-LR	10GBASE-LR SFP+	6.2(6)
	SFP-10G-ER	10GBASE-ER SFP+	6.2(6)
	SFP-10G-ZR	10GBASE-ZR SFP+	6.2(6)
	DWDM-SFP10G-xx.xx	10GBASE-DWDM SFP+	6.2(6)
	SFP-10G-LRM[4]	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)
	GLC-BX-U	1000BASE-BX10-U	6.2(8)
	GLC-EX-SMD	1000BASE-EX SFP	6.2(8)
N7K-F312FQ-25	QSFP-40G-SR-BD	Cisco 40G BiDi QSFP+	6.2(6)
	QSFP-40G-SR4	40GBASE-SR4 QSFP+	6.2(6)
	QSFP-40G-CSR4	40GBASE-CSR4 QSFP+	6.2(6)
	QSFP-40GE-LR4	40GBASE-LR4 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-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)
N77-F248XP-23E	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.2(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-10G-SR	10GBASE-SR SFP+	6.2(2)
	SFP-10G-LR	10GBASE-LR SFP+	6.2(2)
	SFP-10G-ER	10GBASE-ER SFP+	6.2(2)
	SFP-10G-LRM	10GBASE-LRM SFP+	6.2(2)
	SFP-10G-ZR[2]	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)

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

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	GLC-T	1000BASE-T SFP	6.2(2)
	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)
N7K-F248XP-25	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.0(1)
	SFP-10G-SR	10GBASE-SR SFP+	6.0(1)
	SFP-10G-LR	10GBASE-LR SFP+	6.0(1)
	SFP-10G-ER	10GBASE-ER SFP+	6.0(1)
	SFP-10G-LRM	10GBASE-LRM SFP+	6.0(1)
	SFP-10G-ZR[2]	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-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)

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

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	CWDM-SFP-xxxx[1]	1000BASE-CWDM	6.0(1)
	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	10GBASE-SR SFP+	6.1(2)
	SFP-10G-LR	10GBASE-LR SFP+	6.1(2)
	SFP-10G-ER	10GBASE-ER SFP+	6.1(2)
	SFP-10G-LRM	10GBASE-LRM SFP+	6.1(2)
	SFP-10G-ZR[2]	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-BX-D	1000BASE-BX10-D	6.1(2)
	GLC-BX-U	1000BASE-BX10-U	6.1(2)
	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	10GBASE-SR SFP+	5.2(1)

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

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	SFP-10G-LR	10GBASE-LR SFP+	5.1(1)
	SFP-10G-ER[2]	10GBASE-ER SFP+	5.1(1)
	SFP-10G-LRM	10GBASE-LRM SFP+	5.1(1)
	SFP-10G-ZR[2]	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)
	DWDM-SFP10G-xx.xx[1]	10-GBASE-DWDM SFP+	6.1(1)
N7K-M108X2-12L	SFP-10G-SR[3]	10GBASE-SR SFP+	5.2(3a)
	SFP-10G-LR[3]	10GBASE-LR SFP+	5.2(3a)
	SFP-10G-LRM[3]	10GBASE-LRM SFP+	5.2(1)
	SFP-H10GB-CUxM[3]	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)

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

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	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)
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)
	DWDM-SFP-xxxx[1]	1000BASE-DWDM	5.0(2a)

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

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	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	10GBASE-SR SFP+	4.2(6)
	SFP-10G-LR	10GBASE-LR SFP+	4.0(3)
	SFP-10G-ER	10GBASE-ER SFP+	4.0(1)
	SFP-H10GB-ACUxM[2]	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	10GBASE-SR SFP+	5.1(1)
	SFP-10G-LR	10GBASE-LR SFP+	5.1(1)
	SFP-10G-ER	10GBASE-ER SFP+	5.1(1)
	SFP-10G-LRM	10GBASE-LRM SFP+	5.1(1)
	SFP-10G-ZR[2]	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[2]	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	10GBASE-DWDM SFP+	6.1(1)
N7K-M224XP-25L	FET-10G	Cisco Fabric Extender Transceiver (FET)	6.1(1)
	SFP-10G-SR	10GBASE-SR SFP+	6.1(1)
	SFP-10G-LR	10GBASE-LR SFP+	6.1(1)
	SFP-10G-ER	10GBASE-ER SFP+	6.1(1)
	SFP-10G-ZR[2]	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[2]	SFP-H10GB-CUxM Twinax Cable Passive (1m, 3m, 5m)	6.1(1)

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

I/O Module	Product ID	Transceiver Type	Minimum Software Version
	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	QSPF-40G-SR-BD	Cisco 40G BiDi QSFP+	6.2(6)
	FET-40G	Cisco 40G Fabric Extender Transceiver (FET)	6.2(6)
	QSFP-40G-SR4	40GBASE-SR4 QSFP+	6.1(1)
	QSFP-40G-CSR4	40GBASE-CSR4 QSFP+	6.2(2)
	QSFP-40GE-LR4	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)
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)

1. For a complete list of supported optical transceivers of this type, go to the Cisco Transceiver Module Compatibility Information page.
2. Only version -02 is supported.
3. Requires CVR-X2-SFP10G, OneX Converter Module (X2 to SFP+ Adapter)

4. Multimode fiber supported on ports 41 to 48 only. Single-mode fiber support is applicable to all ports

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

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.

Table 6 Supported ISSU and ISSD Paths

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(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(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)	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
Previous Release	**Release Train**	**Releases That Support ISSU to the Previous Release**	**Releases That Support ISSD from the Previous Release**
NX-OS Release 6.2(2)	6.1	6.1(2), 6.1(3), 6.1(4)	No support

Table 6 Supported ISSU and ISSD Paths

Current Release	Release Train	Releases That Support ISSU to the Current Release	Releases That Support ISSD from the Current Release
	6.0	6.0(4)	No support
	5.2	5.2(4), 5.2(7), 5.2(9)	No support

Note Cisco NX-OS Release 6.2(2) does not support an in-service software downgrade (ISSD) to any earlier release because of the scale improvements in Release 6.2(2).

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

ISSU Upgrade Steps

To perform an ISSU upgrade to Release 6.2(2a) 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.

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.
4. Enter the **copy running-config startup-config vdc-all** command.
5. Enter the **reload** command to reload the switch.
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.*

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.

- VDC Ports Can Become Unallocated After a Downgrade

  If your Cisco Nexus 7000 Series device that is running Release 6.2(2) 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(2a). 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:

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

  1. Enter the **write erase** command on the switch.
  2. Reload the switch with the downgrade image.
  3. 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 downgrade to, follow these steps to change instances of **f2e** to **f2** in the **limit-resource module-type** command of the Release 6.2(2a) configuration, before reloading the switch:

  1. Copy your Release 6.2(2a) configuration to a file on a remote server by entering the **copy running-config** *saved-config* 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 **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, 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(2a). 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(2), 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(2), 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 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, 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, or after it.

  Before or after an upgrade from Cisco NX-OS Release 6.1(x) to Release 6.2, 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(2) or a later release 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(2) 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.

- 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. Enter the **no lisp** *feature* command to individually unconfigure the LISP commands. Then enter the **no feature lisp** command. After the ISSU completes, enter the f**eature lisp** command to reenable LISP and then reconfigure it.

- 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(8)

Diagnostic Optical Monitoring (DOM) is supported on the F3 Series modules in Cisco Nexus 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 29
- Cisco Nexus 7000 Series I/O Module, page 29
- Cisco Nexus 7700 Series I/O Modules, page 30

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, a nd VXLAN, MPLS, VPLS, DFA. This module supports Cisco Nexus 2000 Series Fabric Extenders.

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.

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 30
- Cisco Nexus 7718 Switch, page 30
- Cisco Nexus 7700 Series I/O Module, page 30
- Cisco Nexus Fabric Extenders, page 30
- Cisco Nexus 7000 Series Network Analysis Module, page 31

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

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 23 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). Queueing 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.

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 37
- BFD for IPv6 Static Routes, page 37
- Static Route to a VLAN, page 37
- USGv6 and IPv6 Phase 2 Ready Logo, page 37
- FIPS, page 37
- Dynamic Fabric Automation (DFA), page 37

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 Cisco's Evolutionary 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, page 38
- MAC Security, page 38

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 unicast 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.

MIBs

This section contains the following topics:

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

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:

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 nonmixed 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.

Caveats

This section includes the following topics:

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

- 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: Delete the vPC leg port channel and recreate the port channel using the **channel-group** command to clean up the stale translation table entries and restore translation table entries according to the current port mode of the VPC leg port-channel.

  Example:

Delete the port-channel:

```
switch-two(config)# no interface port-channel 10
```

Recreate the port-channel using the member port-config:

```
switch-two(config)# int e7/1-2
switch-two(config-if-range)# channel-group 10
```

Reconfigure the port-channel as a vPC:

```
switch-two(config-if-range)# int po10
switch-two(config-if)# vpc 1
```

- 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: Do a shut, no-shut on all of the ports that have neighbors that can see the RemIndex populated in the SNMPWALK.

- 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: Configure the **no otv isis hello padding always** command on all overlays. Change the value of the interface MTU on the overlay to less than the default of 1400 and change the value of the lsp-mtu to less than the default of 1392.

- CSCul81224

  Symptom: If you attempt to configure a host address of .0 or .255 for an SNMP host, the following message is displayed and the host address cannot be configured:

  ```
  Bad IPV4 host address
  ```

  Conditions: This symptom is seen when a new SNMP host is configured in Cisco NX-OS Release 6.2(2) and Cisco NX-OS Release 6.2(2a).

  Workaround: None.

- CSCty67801

  Symptom: Creation of a switch virtual interface (SVI) should fail if a virtual fabric interface (VFI) is configured on a VLAN. In addition, a VFI configuration on a VLAN should fail if a corresponding SVI is created.

Conditions: This symptom might be seen if both an SVI and a VFI are configured for a VLAN at the same time.

Workaround: Do not configure both an SVI and a VFI for a VLAN at the same time.

- CSCua92310

  Symptom: A pseudowire stays down after a router ID or loopback configuration change.

  Conditions: This symptom might be seen when there is a configuration change for a Layer 2 router ID or a loopback IP address of a Virtual Circuit (VC) that is already up.

  Workaround: Configure the Layer 2 router ID first and then provision the VFI.

- 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.

- CSCub27817

  Symptom: Some FEX vPC+s do not come up after a switch reload in a scale setup.

  Conditions: This symptom might be seen in a scale setup with approximately1000 VLANs when a switch reload is performed with a saved configuration and the configuration has port channels in access mode (but not in trunk mode). As a result, some of the VLANs might fail to come up on some interfaces.

  Workaround: Flap the port channel. In addition, configure port channels in trunk mode.

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

  - Move to a template that has no two queues that have the same IVL values.
  - Move the CoS between the different ingress queues so that no two queues have same IVL value.

- 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: Manually enter the **switchport trunk allowed vlan** command after the rollback.

- 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.

- CSCue96044

  Symptom: A port channel that has the Link Aggregation Control Protocol (LACP) enabled and is configured with an Ethernet Flow Point (EFP) rewrite push tag, goes into a suspended state.

  Conditions: This symptom might be seen when packets are dropped by a port ASIC check.

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

  Workaround 1:

  - Remove the EFP rewrite configuration of EFP.
  - Bring up the EFP.
  - Reapply the **rewrite push tag** command.

  Workaround 2:

  - Use a static port channel with EFPs.

- 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.

- CSCug21520

  Symptom: The Web Cache Control Protocol (WCCP) is not supported in a mixed VDC on F1 or F2 series modules. Policies would fail.

  Conditions: This symptom might be seen when the “WAE is not reachable” error message appears.

Workaround: None.

- CSCug69694

  Symptom: An EPLD upgrade of DPFPGA0 or DPFPGA1 might fail occasionally. The upgrade process starts and after some time might fail with a mismatch error.

  Conditions: This symptom might be seen under normal operating conditions for a Cisco Nexus 7000 Series device when a manual upgrade of the DPFPGA is needed.

  Workaround: Execute the upgrade sequence again. Typically, the upgrade is successful in the next one or two tries.

- CSCug83039

  Symptom: A Virtual Circuit (VC) does not go down when the label switch path goes down when the OSPF configuration is removed on the remote PE.

  Conditions: This symptom might be seen when a second fault happens such as an MPLS path or a Label Distribution Protocol (LDP) adjacency going down in the VC after a Layer 2 VPN process failure or a supervisor switchover. The URIB error notifications are not propagated to the VC after a stateful recovery.

  Workaround: None. Enter the **clear l2vpn service xconnect** command to reprovision the VC and bring it back to the correct state.

- CSCug83123

  Symptom: If a nondefault logging level is set for the Link Layer Discovery Protocol (LLDP) and a switchover occurs, the running configuration does not show the logging level LLDP 5.

```
switch(config)# show running-config | grep lldp
feature lldp
```

  The logging level should be shown as follows:

```
switch(config)# show running-config | grep lldp
feature lldp
logging level lldp 5
```

  Conditions: This symptom might be seen when LLDP sets a nondefault logging level explicitly with the **logging level lldp** command and a switchover occurs. The issue occurs in an ASCII configuration when the correct logging level is not displayed in the running configuration after a switchover.

  Workaround: None.

  You can check the logging level of LLDP after a switchover to verify that it synchronizes correctly.

```
switch(config)# show logging level lldp
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
lldp                    2                       5

0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)
```

- CSCug85241

Symptom: Interface configurations on a Cisco Nexus 7000 Series Network Analysis Module (Cisco NAM-NX1) are not available to the user for configuration and a meaningful error is not provided. If you copy the file to the running configuration, the errors appear on the switch console. For example:

```
switch# copy copy-r-s running-config
Error: Service card interface configuration is not editable.
Syntax error while parsing 'switchport'
Syntax error while parsing 'switchport mode trunk'
Syntax error while parsing 'switchport monitor ingress'
Syntax error while parsing 'mtu 9216'
Error: Service card interface configuration is not editable.
Syntax error while parsing 'switchport'
Syntax error while parsing 'switchport mode trunk'
Syntax error while parsing 'switchport monitor ingress learning'
Syntax error while parsing 'mtu 9216'
```

Conditions: This symptom might be seen when you try to change the configuration of ports or port channels on the Cisco Nexus 7000 Series Network Analysis Module.

Workaround: None.

- 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: None. Because the VDC is reloading, the error log is ignored.

- CSCug95664

  Symptom: A pseudowire (PW) comes up with Type 4 for bridge domains but there is no dummy tag inserted which results in traffic loss over the Ethernet Flow Point (EFP).

  Conditions: This symptom might be seen when a **vc type vlan** command is configured in long form PW, or when a port profile is inherited with that configuration. The issue occurs only in the VPLS service with the EFP member in the BD, when either VPLS PW peer has this configuration.

  Workaround: None.

- CSCuh09113

  Symptom: A FEX satellite interface remains up during an image download. Ideally, the interfaces should be down when-the FEX is downloading an image.

  Conditions: This symptom might be seen only when a FEX image version mismatches with a Cisco Nexus 7000 Series device and the FEX tries to download a new image.

  Workaround: None.

- 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.

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

  - Enter the **show run bfd** command to get the BFD configuration. Enter the **no feature bfd** command to disable the BFD feature and enter the **feature bfd** command to reenable it. Save all BFD configurations and reapply the configurations to the running configuration after BFD is reenabled.
  - Enter the **reload module** command to reboot the line cards.
  - Enter the **reload** command to reboot the switch.

- CSCuh21793

  Symptom: The Ethanalyzer tool does not capture certain packets when they originate from F1 or F2 Series modules.

  Conditions: This symptom might be seen when certain packets that originate from F1 or F2 Series modules have proprietary headers associated with them. When the Ethanalyzer tool is used with a capture filter, some packets that match the criteria are not captured because the filter gets applied incorrectly due to the proprietary headers.

  Workaround: Capture the packets without using any capture filters and analyze the captured packets using the “display-filter” option of the Ethanalyzer tool. Otherwise, copy the pcap file to an external workstation and use Wireshark software to analyze the packets.

- 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 sizeable 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.

- CSCuh39086

  Symptom: When an ACL QoS policy is pushed by the policy server and gets programmed correctly, a transient and brief failure for sending the status from the line card immediately following a switchover is logged. The following error appears:

  ```
  ACLQOS-SLOT1-2-ACLQOS_FAILED: ACLQOS failure: Error sending client status for verify
  session ret_val 0x801c006e
  ```

  Conditions: This symptom might be seen when a configuration change was in progress and a switchover occurred.

  Workaround: None required. The situation will automatically correct itself.

- CSCuh44476

  Symptom: Some neighbors are not discovered and the Virtual Circuits (VCs) 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: Configure all provider edges as separate devices.

- CSCuh46295

  Symptom: Port security works correctly when configured on two interfaces, and all MAC addresses are learned through the interfaces. If the **shut** command is entered on the interfaces, other interfaces have a security violation error.

  Conditions: This symptom might be seen on a Cisco Nexus 7000 Series device that is running Cisco NX-OS Release 6.2(2) or Release 6.1(2).

  Workaround: Move the two ports into a Link Aggregation Control Protocol (LACP) channel.

- CSCuh50203

  Symptom: Link Layer Discovery Protocol (LLDP) MIB support is incomplete for FEX interfaces on Cisco Nexus 7000 devices in Cisco NX-OS Release 6.2(2).

  Conditions: The following OIDs for snmpget/snmpwalk/snmpgetnext queries are not supported for FEX interfaces:

  - lldpStatistics
  - lldpRemUnknownTLVTable
  - lldpRemOrgDefInfoTable

  The following OIDs are not supported only for snmpset for FEX interfaces:

  - lldpPortConfigNotificationEnable
  - lldpPortConfigTlvsTxEnable

  Workaround: None.

- 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: To recover from this issue, enter the **shut** command followed by the **no shut** command on the affected ports.

  To prevent this error, enter the **monitor session** *x* **type-erspan destination** command followed by the **shut** command prior to a VDC reload or reload of the modules with ERSPAN destination ports. Repeat those commands for all operationally up ERSPAN destination sessions.

- 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: Disable the policies that were applied after the error and apply them again. If that does not correct the problem, disable the feature and enable it again.

- 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.

- CSCuh53070

  Symptom: Cisco Nexus 2000 Virtual Circuits (VCs) take a long time to come up. There are continuous Network Layer Reachability Information (NLRI) label advertisements/withdrawals that cause the route to be provisioned after an extended delay.

  Conditions: This symptom might be seen when a global Layer 2 router ID/loopback change causes the readvertisement with a new SAII.

  Workaround: If a router ID change needs to made, enter the **shut** command in the Layer 2 VPN submode, make the change, and enter the **no shut** command.

  Another workaround is to make the router ID configuration change and do a stateless restart of the Layer 2 VPN feature with the new configurations.

- CSCuh53852

  Symptom: Link Layer Distribution Protocol (LLDP) get, getnext, and set that use LldpPortNumber as the index into LLDP MIB table fail in some SNMP managers and NMS.

  Conditions: The following indices are used for LLDP MIB tables:

  - lldpPortConfigPortNum (index into the lldp port config table)
  - lldpStatsTxPortNum (index into the Tx stats MIB table)
  - lldpStatsRxPortNum (index into the Rx stats MIB table)
  - lldpLocPortNum (index into the Local System Data Table)
  - lldpRemLocalPortNum (index into the Remote System Data Table)

  These OIDs have the textual convention: LldpPortNumber.

  Because LLDP works for Layer 2 and Layer 3 ports, InterfaceIndex is used for all of them. The TC LldpPortNumber has a range of only 1 to 4096, and all the interface indexes are beyond this range. As a result, the SNMP queries work in some SNMP Managers or Network Management System (NMS), but do not work in others.

  Workaround: To work around these issues, follow these steps:

  1. Execute snmp queries from a MIB browser to return the correct output for all the LLDP MIB OIDs.
  2. Remove the MIB VARS setting so that the Textual Convention LldpPortNumber does not apply. The SNMP queries will return and set the correct value.

- CSCuh71741

  Symptom: Syslog error messages such as the following might appear occasionally:

  ```
  vsh: unknown enum:0, tid(hex):22a0142
  ```

  Conditions: This symptom might be seen when moving ports on the Cisco Nexus 7000 Series Network Analysis Module between VDCs, especially after an ISSU or a HA switchover.

  Workaround: None. These errors are transient and can be ignored. No functionality is affected.

- 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: Terminate the command at the CLI.

- CSCuh91865

  Symptom: The output of the **show interface tunnel** *number* command has ambiguous information if both the interface tunnel and the interface tunnel-te are configured with the same tunnel number.

  Conditions: This symptom might be seen when both GRE and TE tunnels are configured with the same tunnel number.

  Workaround: Avoid using the same tunnel number for both GRE and TE tunnels.

- CSCuh94708

  Symptom: The output of the **show hardware interface capacity** command shows only error counter drops. It does not include congestion drops.

  Conditions: This symptom might be seen under normal operating conditions for a Cisco Nexus 7000 Series device.

  Workaround: Use the **show interface counters** command to see the congestion and error drops for each port.

- 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.

- 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: Use a lower scale switch-wide, move to Multiple Spanning tree mode, or, if possible, move to a higher-end supervisor.

- CSCui02936

  Symptom: A port channel with the Link Aggregation Control Protocol (LACP) and a mismatched native VLAN configuration is suspended.

  Conditions: This symptom might be seen because Layer 2 PT requires trunk ports to have a native VLAN that is the same as the trunk VLAN for CDP tunneling. The CE native VLAN is the default for VLAN 1. There is a native VLAN mismatch on the two ends of the port-channel link that causes the Link Aggregation Control Protocol (LACP) to suspend the link.

  Workaround: Disable LACP on the port channel.

- 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: Disable the Gold StandbyFabric Loopback test and Rewrite engine test before power cycling the fabric modules or opening or closing the fabric ejectors. Enable the tests once the spine is replaced or removed permanently.

- 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.

- 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: To avoid the issue, clear the IPv6 URPF/Redirects information. Manually configure IPv6 URPF/IPv6 Redirects (in the same mode that you used before moving to switchport mode) and unconfigure IPv6 URPF/Redirects with the **no** version of the commands.

- 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.

- CSCui13170

  Symptom: When a switch boots up and an ASCII configuration is replayed from a file in a vPC+ scale configuration, the access port-channel interfaces go into an error disabled state. The interfaces in the vPC+ switch go into a loop from “initializing” to “link not connected.”

  Conditions: This symptom might be seen when a switch is freshly booted up and an ASCII configuration is replayed from a saved configuration file in vPC+ scale configuration.

  Workaround: Configure and bring up the peer link before applying another vPC+ scale configuration.

- CSCui13285

  Symptom: If the fabric module is in the POWER_UP state (during initialization) after a power on or insertion, an online insertion and removal (OIR) results in the module having a fabric synchronization failure and port ASIC PL congestion.

  Conditions: This symptom is rare and might be seen only if an OIR occurs during the time when the fabric module is initializing.

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

  - Perform the OIR after inserting a fabric module, or power it on from the CLI and wait until it comes online.
  - Enter the **poweroff xbar** command or **out-of-service xbar** command before the OIR to gracefully power down the fabric module.

- 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: None. A rollback works when a single overlay is defined with VLAN mappings. Other overlays should not have VLAN translations.

- 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.

- 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: If LDP is in this state, try again to unconfigure or reconfigure LDP to clear the problem.

- 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: None.

- CSCui26099

  Symptom: During an admin VDC migration, incomplete configurations on one interface can affect porting of configurations on some other interfaces.

  Conditions: This symptom might be seen if there are inline comments in the output of the **show running-config** command. Comments start with "!" or "#."

  Enter the following command to check if there are comments in the running configuration.

```
switch# show running-config | grep [!#]
!Command: show running-config
!Time: Fri Jul 26 23:32:14 2013
    ! Incomplete config, specify a neighbor
```

  The first two lines are normal comments that exist in the beginning of all running configurations, and they do not cause any issues. The third line is an inline comment that causes problems after the admin VDC migration. The configuration should be corrected before the admin VDC migration.

  Inline comments can be caused by incomplete or invalid configurations. Two cases have been identified that might cause inline comments. One is a Layer 2 VPN feature with an incomplete configuration for pseudowire. The other is the EFP/EVC feature with mismatched configurations. Following are examples of those cases:

```
interface pseudowire7
! Incomplete config, specify a neighbor
  encapsulation mpls

interface Ethernet2/2
    service instance 6 ethernet
    encapsulation dot1q 100
    !Invalid config. Encap vlan 100 does not match bridge-domain 6
```

  These inline comments cause problems after the admin VDC migration.

  Workaround: Before the admin VDC migration, correct the configuration by making sure there are no inline comments.

  If this issue does occur after an admin VDC migration, try to correct the problem through one of the following methods:

  - Reboot the switch with a clean database and reapply the configurations. Correct the inline comments before attempting another admin VDC migration.

  - Manually correct the errors in the configuration if the expected result is understood. Compare the running configuration with the expected running configuration and modify any mismatches.

- 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: Try the reload again. The VDC will reload again and become active.

- 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.

- 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: Enter the **shut** command followed by the **no shut** command on the port.

- 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.

- CSCui30756

  Symptom: After an Authoritative Edge Device (AED) failover in an OTV site, traffic going out of the site to receivers that are silent might experience a delay of about 15 seconds until the new AED issues a TCN flush. After this time, the traffic fully converges and there is no impact on functionality.

  Conditions: This symptom might be seen when source traffic originates from the site that is undergoing an AED failover. Silent receivers are present in the remote site(s).

  Workaround: None.

- 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.

- CSCui30906

  Symptom: For IPv4 and IPv6, when the owner_type is set to CONNECTED in the filter, the onep_routing_rib_get_route_list function does not return the connected (direct) route from a Cisco Nexus 7000 Series device.

  Conditions: This symptom might be seen in the RIB when the get route list function has the specific owner_type==CONNECTED.

  Workaround: None.

- CSCui31353

  Symptom: A port channel on a local customer edge (CE) device gets into a blocked state.

  Conditions: This symptom might be seen when a Spanning Tree Protocol (STP) dispute occurs due to a bridge protocol data unit (BPDU) not being exchanged correctly on a Layer 2 VPN setup.

  Workaround: Enter the **spanning-tree bpdufilter enable** command on the Ethernet interface on the far end CE device so that it does not send a BPDU at all.

- CSCui32036

  Symptom: When you enter the **no feature-set fabricpath** command and perform a switchover, VLANs disappear from the topology after the switchover. Enter the **show system internal m2rib topo** command to confirm that the VLANs are not displayed in the topology. FabricPath traffic loss can occur on these VLANs as a result.

  Conditions: This symptom might be seen after a switchover or ISSU. Before the switchover, the **no feature-set fabricpath** command was entered, which caused the VLANs and topologies to be disassociated.

Workaround: After the switchover, enter the **no mode fabricpath** command and the **mode fabricpath** command.

- 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.

- CSCui33320

  Symptom: After the **clear ospf neighbor** command is entered, OSPF takes 45 seconds to converge and bidirectional forwarding detection (BFD) takes approximately 2 minutes, 30 seconds to converge for a large number of sessions on physical interfaces. BFD takes approximately 2 minutes 10 seconds to converge after the **clear ospf** command is entered from a steady state.

  Conditions:. This symptom might be seen only when you clear a large number of OSPF or BFD neighbors together. There should be no major difference in convergence for a small number of BFD sessions.

  Workaround: None.

- CSCui33326

  Symptom: When a line card is reloaded or ejected and plugged in, the following syslog message appears:

  ```
  2013 Jun 19 08:18:32 N7k-114-102 %ETHPORT-2-IF_CRITICAL_FAILURE: (Debug
  syslog)Critical failure:
  ```

  The result is that the line card bringup is a delayed by a few minutes.

  Conditions: This symptom is rare, but might be seen under high scale scenarios.

  Workaround: None, but there is no functional impact because the line card does come up.

- 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: The application should call the function onep_routing_rib_get_route_list to get all the routes initially to register later for events.

- 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: After performing a non-ISSU upgrade using the startup configuration file, enter the following commands:

```
show running-config aclmgr inactive-if-config <<< This command will give all ACL
inactive config
clear inactive-config acl <<< This command will clear all the ACL config, and create a
file in bootflash.
```

Note The inactive interface configuration for the ACL manager is saved to /bootflash/aclmgr_inactive_if_config.cfg for the default VDC and for VDCs other than the default VDC to /bootflash/vdc_x/aclmgr_inactive_if_config.cfg (where x is vdc number).

Find this file and apply it as an ASCII replay. In the default VDC, enter the following command:

```
copy bootflash:aclmgr_inactive_if_config.cfg running-config
```

Ensure that there is no inactive configuration by entering the following command:

```
show running-config aclmgr inactive-if-config
```

- CSCui39120

  Symptom: A VPLS Virtual Circuit (VC) (Border Gateway Protocol [BGP] AD and Label Distribution Protocol [LDP] signaled) disappears after a supervisor switchover.

  Conditions: This issue might be seen after a switchover of a Supervisor 2 module. The VCs go down and come back up, but one VC disappears. This issue probably will not be seen following a supervisor switchover when the two peers are VDCs on the same device.

  Workaround: Ensure that the peers are on different devices and are not simultaneously switched over.

- CSCui40856

  Symptom: If physical port vPC paths are in mode access, and you remove and add back the VLAN in these vPC paths from the MCT, the vPC paths remain down.

  Conditions: This symptom might be seen under these conditions:

  - The vPC paths are physical ports and in mode access.
  - A VLAN that is configured on the vPC paths is removed and added back from the MCT.

Workaround: Manually flap the vPC paths.

- 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.

- CSCui41300

  Symptom: The **show mac address-table vlan 100** command and the **show mac address-table bridge-domain 100** command both return the same MAC address entries.

  Conditions: This symptom is seen in Cisco NX-OS Release 6.2(2). An error message should appear.

  Workaround: None.

- CSCui41381

  Symptom: The system has bridge domain (BD 200), but no VLAN 200. However, the **show hardware mac address-table vlan 200** command displays the hardware entries, but the **show hardware mac address-table 3 bridge-domain 200** command does not display an error:

```
switch (config)# sh hardware mac address-table 3 bridge-domain 200
ERROR: VLAN 0 does not exist!
```

  Conditions: This symptom might be seen under normal operating conditions for a Cisco Nexus 7000 Series switch.

  Workaround: None.

- 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: None.

- CSCui45267

  Symptom: If there is an error from the SPM before a switch reload, the policies are not reapplied after the error and reload. As a result, the TCAM becomes incorrect.

  Conditions: This error might be seen when an error from the SPM occurs before the switch reloads.

  Workaround: Disable the feature and enable it again.

- 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.

- CSCui49752

  Symptom: Traffic downtime occurs for VPLS traffic.

  Conditions: This symptom might be seen after a stateful switchover.

  Workaround: None.

- CSCui51103

  Symptom: Approximately 3 minutes after an ISSU completed, traffic loss occurred for a number of multicast routes. It takes approximately 69 seconds for the traffic to converge.

  Conditions: This symptom might be seen when an ISSU occurs in this sequence:

  - The standby supervisor was upgraded and reloaded.
  - A switchover occurred from the standby supervisor to the active supervisor.
  - The active supervisor was upgraded and reloaded.
  - All the line cards upgraded successfully.

  Approximately 3 minutes after the ISSU completed, the multicast packet loss started and took around 69 seconds to converge.

  Workaround: None. Due to some error, the routes were expired unexpectedly. The traffic recreates the multicast routes after that.

- CSCui51401

  Symptom: IPv6 Bidirectional Forwarding Detection (BFD) sessions go down when an IPv6 RACL is applied on the switch virtual interface (SVI) where BFD is applied. BFD sessions applied alone work correctly.

Conditions: This symptom might be seen when there is something wrong in the IPv6 merge compression logic. There are some entries in the IPv6 TCAM that are probably causing packets to get redirected to the wrong interface. The sessions go down as OSPF packets get dropped. BFD is still able to receive packets.

Workaround: None.

- CSCui52007

  Symptom: Multicast traffic drops occur because a route becomes unsynchronized between the control plane and the line card.

  Conditions: This symptom might be seen during a parallel ISSU. The switch is not in a steady state and the multicast route update occurs when the line cards are being upgraded.

  Workaround: Enter the **clear ip mroute grp_ip src_ip** command to clear the affected route.

- 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: Reapply the ASCII configuration.

- 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.

- CSCui54639

  Symptom: Custom queuing polices are incomplete after an admin VDC migration.

  Conditions: This symptom might be seen after you do the following:

  1. Create custom queuing polices.
  2. Migrate to a new admin VDC.

  Queuing policies do not have the class maps and actions that were configured before the migration.

  Workaround: Reconfigure queuing policies with the correct class maps and actions.

- CSCui54572

Symptom: An (S, G) state is created with a NULL OIF list on a transit router (after an ISSU), even when the last hop has the **ip pim spt-threshold-infinity** command configured.

Conditions: This symptom might be seen when the **ip pim spt-threshold-infinity** command is set for a last hop router. If the transit router undergoes an ISSU to Cisco NX-OS Release 6.2(2), it is possible that while unicast routing is still converging, data packets get redirected (due to the RPF check fail), which creates the (S, G) state. The result is traffic that is not forwarded for the (S, G) groups that were created.

Workaround: Enter the **clear ip mroute** command for the (S, G) groups that were created.

- 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.

- 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: Configure specific routes for the Map Server (MS), the Map Resolver (MR), and the Proxy Egress Tunnel Router (PETR). These routes can be static routes that are added to the router, or learned through a routing protocol

- 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: Enter the **clear l2vpn service vfi name** *VFI-context-name* command, or delete and reconfigure the PW.

- 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: Delete and readd the **short form member** command for the VFI.

- 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: Use no more than 126 HSRP secondary addresses per HSRP group to avoid problems with the **show running configuration** command and the loss of secondary addresses after a reload.

- 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 issuing 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.

- CSCui75788

Symptom: When you are working with vPC+ on the F2 series modules, you might see stale MAC addresses in the MAC address table after you reload the module. The hardware MAC address tables will show that the RD value is 1.

Conditions: This symptom might be seen when the VLAN is in FabricPath mode with traditional learning, and the aging time is low (typically lower than 10 minutes).

Workaround: Clear the MAC address table for the problem addresses.

- 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: There is no functional impact from this caveat. Reloading the switch should fix the issue.

- 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: None.

- CSCul42680

  Symptom: When you are working with vPC and the ports on a FEX module go up and down, some MAC addresses will be missing on the primary vPC. The MAC address is deleted from the Peer 1 as a part of the shutdown even though Peer 1 did not receive an age notification from Peer 2.

  Conditions: This symptom might be seen when the MAC address is on Peer 2 synchronized to Peer 1. As there is traffic on Peer 2, the MAC address does not age out on Peer 2. However, the MAC address ages out on Peer 1, which sends an age notification to Peer 2. Then, the fabric port channel on Peer 1 goes up and down.

  Workaround: Configuring **clear mac address-table dynamic** on both peers causes the MAC addresses to be relearned on Peer 2 and synced to Peer 1.

- 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: Make the shared tree and the source tree the same path for the sparse mode host or set the SPT thresholds to infinity for hosts only.

- CSCul50626

Symptom: You might see a ~20K memory leak when you bring the port up and down when you are working with more than 2,000 VLANs configured with a vPC peer and FEX module connected through breakout ports.

Conditions: This symptom might be seen when the FEX module is connected only on one side of vPC leg.

Workaround: Remove the FEX module.

- 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: Enter the **show hardware internal statistics** command to see the counters.

- 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: Enable dot1x after MVRP is enabled.

- CSCul70427

  Symptom: After removing and reinserting the supervisor, you might see ACL QoS fail.

  Conditions: This symptom might be seen when you remove and reinsert the supervisor module.

  Workaround: Use Autorecovery.

- CSCul75153

  Symptom: Leaves not sent after non MVRP trunk port VLAN membership is updated.

  Conditions: This symptom might be seen when a FEX trunk is configured. The register is sent upstream for all VLANs because FEX ports do not run MVRP. However, if changes are made to the configuration for ex set trunk allowed vlan to one VLAN only, then the leave is not sent for the others. This applies to all ports that do not run MVRP.

  Workaround: None.

- 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: If the symptom appears, enter the **terminal reset vlan-config-mutex** command.

- CSCug04719

  Symptom: When you perform an ISSU from Release 6.1(4) to Release 6.2(1.59)S5, the Label Distribution Protocol (LDP) might stop on the newly active supervisor.

  Conditions: This symptom might be seen when you perform an ISSU from Release 6.1(4) to Release 6.2(1.59)S5.

  Workaround: None.

- CSCum01261

  Symptom: The SNMP counters do not clear the OutUcastPkts counters.

  Conditions: This symptom might be seen when you enter the **clear counters int eth** *x*/*y* **snmp** command.

  Workaround: None.

- CSCul84463

  Symptom: When you add a vPC peer link as a static router port with an 'SVF' flag in the IGMP Snooping mrouter cache, it could lead to unnecessary drawing of traffic to the vPC+ peers.

  Conditions: This symptom might be seen in IGMP snooping in vPC+ environment where the vPC peer-link gets added as static router port.

  Workaround: Either bring the vPC peer link down and up, or unconfigure/configure IGMP snooping on the problem VLAN.

- 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: If the phy-port vPC leg must be shut on both switches, shut the one on the primary first, followed by shutting the secondary one.

- CSCum06933

  Symptom: When a module receives CRC errors marked FATAL, then the module is reloaded.

  Conditions: This symptom might be seen when a fabric is removed. It is not seen when a fabric is reloaded, powered-down, or powered-up.

  Workaround: Put the fabric out of service by issuing the **out-of-service xbar** *x* command before removing the fabric.

- 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: None.

- CSCuj21989

  Symptom: The VLAN mapping configuration is not failing on a vPC peer link.

  Conditions: This symptom might be seen when you have configured a vPC.

  Workaround: You should not have VLAN translation configured before the link is configured as a peer link. After the link is configured as a peer link, do not configure VLAN translation on this link.

- 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: Reduce the logging level 2 for ARP by entering the **logging level arp 2** command.

- 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.

- CSCuj69529

  Symptom: The multicast traffic drops after an AED failover in the same site.

  Conditions: This symptom might be seen in VLANs that do not have directly connected receivers or senders. Those VLANs contain only PIM routers.

  Workaround: Bring down the OTV site interfaces, wait until the OTV mroutes time out, and then bring up the OTV site interfaces.

- 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.

- CSCul96740

  Symptom: Cannot configure all the VLANs; limited by resource template.

  Conditions: This symptom might be seen when the VLAN mgr creates 3 VLANs—1, 4040, and 4042—on system boot up. There are now from 622 onwards. If the resource manager counts these 3 Boot Up reserved VLAN as resources, then you can create only 1021 VLANs and will receive an error.

  Workaround: None.

- CSCum09180

  Symptom: The ACL log is not getting rate limited when RACL merges with the NetFlow sampler.

  Conditions: This symptom might be seen when a NetFlow sampler is configured with a RACL that is log enabled.

  Workaround: None.

- CSCum15561

  Symptom: On scale/batched triggers such as FEX reloads and peer reloads you might see STP set port state failures.

  Conditions: This symptom might be seen with a scale configuration of approximately 20 host interfaces or more with 75 VLANs per host interface on a 2248PQ FEX type. With this configuration, when batched triggers are executed, STP set port failure are seen.

  Workaround: Bring the interfaces where the error is seen down and up.

- 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.

- CSCum01596 and CSCum01599

  Symptom: The F1 and F2 Series modules are slow to execute the cshcMacUsageTable walk.

  Conditions: This symptom might be seen when you are working with the F1 or F2 Series modules.

  Workaround: Reduce the frequency of polling.

- CSCuj75717

  Symptom: ORIB and ISIS redistribute databases are out of sync for unicast-only overlay VLANs on clearing OTV-ISIS neighbors.

  Conditions: This symptom might be seen after you enter the **clear otv isis adjacency** command multiple times.

  Workaround: None.

- 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: Restart PIM, or you can stop traffic and joins and wait for the routes to time out. Then start traffic and joins at around the same time.

- 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: None.

- CSCul08181

  Symptom: When you bring down the trunk in a promiscuous vPC port channel, the BPDUs time out on the vPC peer link in all VLANs.

  Conditions: This symptom might be seen when you bring down the trunk in a promiscuous vPC port channel.

  Workaround: None.

- 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: Do not change the join interface for an overlay in order to avoid the problem entirely. If you already see the problem, clear the IP mroute.

- CSCum11808

  Symptom: The Egress WCCP policy is not applied and the packet is dropped.

  Conditions: This symptom might be seen when the ingress port is a FabricPath VLAN, the egress port is a port channel subinterface, subinterface, or SVI, and the policy type is an egress WCCP policy on the routed packet.

  Workaround: None.

- CSCul67147

Symptom: ISSD might not work if you have entered the **fabricpath domain default** command under the interfaces.

Conditions: This symptom might be seen when you have entered the **fabricpath domain default** command under the interfaces.

Workaround: Enter the **no fabricpath domain default** command.

- CSCul81991

  Symptom: When you apply a breakout configuration, you might see large invalid counters on an interface.

  Conditions: This symptom might be seen when you are doing breakout/ no breakout on high bandwidth ports.

  Workaround: Clear the counters manually after you do a breakout/no breakout by entering the **clear counters interface** command on the created broken out/parent port(s).

- 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: Clear the MAC address table on the device that is the current owner of the local MAC addresses, or bring the overlay down and up.

- 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: None.

- 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: None.

- CSCuj17442

  Symptom: When TACACS authorization is enabled, you cannot add the **inherit** command.

  Conditions: This symptom might be seen when TACACS authorization is enabled.

  Workaround: Remove TACACS authorization commands.

- CSCuj33260

  Symptom: A core might be unexpectedly written out for the Netstack application.

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

  Workaround: None.

- CSCuj51652

  Symptom: You might see the Aclmgr go down when the you enter the **channel-group** *channel-group* on a port or a group of ports.

  Conditions: This symptom might be seen if the channel-group command fails with this error:

  ```
  command failed: no such pss key [membership update failed for port-channel999
  [service: unknown err: 0x40480003 - no such pss key]]
  ```

  Workaround: Determine why the channel-group command is failing and fix it.

- CSCuj81349

  Symptom: You might see a vsh core.

  Conditions: This symptom might be seen after applying a configuration.

  Workaround: None.

- CSCul88079

  Symptom: DHCPv6 cannot be used with POAP.

  Conditions: This symptom might be seen if you are working with POAP.

  Workaround: None.

- 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.

- CSCul90173

  Symptom: When using the vPC track feature, the vPC is brought down at the moment when the tracking object goes down. After some time (30 sec—configured UP interval) the VPC is restored although the tracking object is down.

  Conditions: This symptom might be seen when you configure the vPC track feature and bring down the remote side of the link.

  Workaround: None.

- 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: None.

- 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.

- CSCum07092

  Symptom: When you enter the **show interface ex/y capabilities breakout** command, the switch displays the available breakout maps for CB40 interfaces, even though the breakout is not supported on it.

  Conditions: This symptom might be seen if you are working on an interface of CB40 on an F3 Series module.

  Workaround: None.

- CSCul69597

  Symptom: The FabricPath feature might be lost after reloading the VDC when the running config file is copied to bootflash as configuration file and then copied to the start-up config.

  Conditions: This symptom might be seen after you reload the VDC and copy the config file from the bootflash.

  Workaround: Reconfigure FabricPath after you reload the VDC.

- 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.)

- 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: Enter the **clear mac address-table dynamic** command

- CSCum91985

  Symptom: VDC goes to a Failed state when you enter the **vdc reload** command.

  Conditions: This symptom might be seen when you reload the VDC.

  Workaround: None.

- CSCum11100

  Symptom: After you change the maximum path in a FabricPath domain, equal cost multipathing (ECMP) cannot recover.

  Conditions: Because of an inconsistency in ISIS and RIB/FIB o/p in the number of ECMPs, ISIS shows more ECMPs than the forwarding layer can support when configured with a higher than supported value.

  Workaround: Ensure that you do not configure a bigger value.

- 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: If you want to disable PIM protocol from the switch, stagger this process at one interface at a time. Also, when we have a very large number of mroutes, try to shut the interface first, before disabling PIM from it and wait for the routes to age out in the scenario.

- CSCun03952

  Symptom: Even after you stop traffic, you may see that the routes at first hop router (FHR) are still held by IP.

  Conditions: You might see that the routes in MRIB at the FHR do not expire and are held owned by IP, even after all traffic is stopped.

  Workaround: Enter the **clear ip mroute data-created *** command and then enter the **clear ip mroute *** command.

- CSCun16347

  Symptom: After executing the CLI to migrate to the admin VDC, the CLI might not return. This happens because the **banner motd** command is not processed properly by the admin vdc migration script.

  Conditions: You might see this symptom after you execute the CLI to migrate to the admin VDC.

  Workaround: Either reload the router and go back to the default VDC, or remove the CLI from the configuration and then migrate to admin VDC.

- CSCun03944

  Symptom: When the router boots up and it is the first HOP router and you have about 32K multicast streams already started, there are error messages that slab allocation for a PIM route entry failed.

  Conditions: You might see this symptom when the PIM process starts up on the FHR and more than 32K streams are ingressing the router.

  Workaround: Start the multicast streams after all the interfaces are properly enabled with PIM and the unicast routing protocols go to the steady state.

- CSCun16070

  Symptom: The mroutes steady state might fluctuate after you perform an ISSU from Cisco NX-OS Release 6.2(2a) to Release 6.2(6a) with the ASM 32K Scale setup.

  Conditions: You might see this symptom after you perform an ISSU from Cisco NX-OS Release 6.2(2a) to Release 6.2(6a) with the ASM 32K Scale setup. After an ISSU, the mroutes expire and are added back on R4. The problematic route and the RPF on R4 was PC2. However, PC2 is not seen as OIF on the first hop router (FHR) R5.

  Workaround: Enter the **clear ip mroute *** command.

- 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: None.

- 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: Reload the module after you perform the ISSU if OTV-enabled VLANs are present in the Cisco NX-OS Release 6.2(6) from which you are upgrading.

- CSCum76673

  Symptom: When you downgrade from Cisco NX-OS Release 6.2(6) to Release 6.2(2a) and have a large-scale configuration, the older version may not be able to handle the same scale which might cause the VDC and/or ports to move to failure state.

Conditions: You might see this symptom when you are working with a large-scale configuration and you downgrade from Cisco NX-OS Release 6.2(6) to Release 6.2(2a).

Workaround: Reduce the scale of the configuration and reload the affected VDC or port.

- 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.

- CSCun13740

  Symptom: You might see a sequence timeout with Private VLANs during port bringup.

  Conditions: You might see this symptom after you perform an ASCII replay after an ISSU from Cisco NX-OS Release 6.2(2) to Release 6.2(6a).

  Workaround: None.

- CSCun27058

  Symptom: You might see the MAC address learned on the VPC leg, but not sent onto the supervisor on M-Series modules; although the MAC address is seen in the hardware.

  Conditions: You might see this symptom after you perform an ASCII replay with VPCs configured.

  Workaround: Enter the **clear mac address-table dynamic** command.

- CSCun05808

  Symptom: You might see the following error:

  ```
  Unable to perform the action due to incompatibility:  Module 1, 2, 3, 4, 7, 15
  returned status "Aclqos error: no resource found"
  ```

  Conditions: You might see this symptom after you apply RACL an/or VACL with statistics and QoS configured on VLANs. The system tries to remove bank-mapping, which is failing, and so it fails to remove egress the QoS configured on the VLAN.

  Workaround: None.

- 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: 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.

- CSCum11655

  Symptom: When you are working with WCCP and/or PBR policies on an F3 Series module, traffic might be affected.

  Conditions: You might see this symptom if a WCCP and/or PBR policy is applied on an F3 Series module with the action of forwarding the packet to a different module and/or switch; then, you bring down and bring up the interfaces on the local ports on the F3 Series module.

  Workaround: Enter the clear **ip arp * force-delete** command.

- CSCum11655

  Symptom: When you are working with WCCP and/or PBR policies on an F3 Series module, traffic might be affected.

  Conditions: You might see this symptom if a WCCP and/or PBR policy is applied on an F3 Series module with the action of forwarding the packet to a different module and/or switch; then, you bring down and bring up the interfaces on the local ports on the F3 Series module.

  Workaround: Enter the clear **ip arp * force-delete** command.

- CSCum70660

  Symptom: You might see the MAC address on one vPC peer pointing to the local vPC leg, while the MAC address on the vPC peer points to a LID of 0xFFFF. This issue is seen with ARP packets.

  Conditions: You might see this symptom when the ARP request lands on the peer on which the MAC address is correctly learned, and the ARP response lands on the peer where the MAC address is not learned properly.

  Workaround: Sending a data packet corrects the issue.

- 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.

- CSCuo15686

  Symptom: The message 'DOM not supported' is displayed only for the first broken out port.

  Conditions: You might see this issue when you are working with DOM on an F3 Series module.

  Workaround: Check the first broken out port transceiver details for DOM values.

- CSCuj40328

  Symptom: The **show transceiver detail** command displays the SFP type as unknown for the AVAGO Twinax (2m) SFP. This is a display issue only, and there is no impact in functionality. The following is an example of the display:

```
Ethernet6/15
    transceiver is present
    type is unknown
    name is CISCO-AVAGO
    part number is AFBR-7CER02Z-CS1
    revision is B1
    serial number is AVE1650K061-A
    nominal bitrate is 10300 MBit/sec
    Link length supported for copper is 2 m
    cisco id is --
    cisco extended id number is 4
    number of lanes 1
```

  Conditions: You see this issue when you are working with an AVAGO Twinax (2m) SFP.

  Workaround: None.

- CSCuj40171

  Symptom: The **show transceiver detail** command displays the SFP type as unknown for the MOLEX copper SFP (1.5m,2m,2.5m). This is a display issue only, and there is no impact in functionality. The following is an example of the display:

```
Ethernet6/7
    transceiver is present
    type is unknown
    name is CISCO-MOLEX
    part number is 74752-9641
    revision is 09
    serial number is MOC17080232
    nominal bitrate is 10300 MBit/sec
    Link length supported for copper is 1 m
    cisco id is --
    cisco extended id number is 4
    number of lanes 1
```

  Conditions: You see this issue when you are working with a MOLEX copper SFP (1.5m,2m,2.5m).

  Workaround: None.

- CSCuj40109

  Symptom: The **show transceiver detail** command displays the SFP type as unknown for the TYCO Twinax copper SFP (1.5m,2m,2.5m). This is a display issue only, and there is no impact in functionality. The following is an example of the display:

```
Ethernet6/1
    transceiver is present
```

```
type is unknown
name is CISCO-TYCO
part number is 1-2053783-5
revision is R
serial number is TED1646GC18
nominal bitrate is 10300 MBit/sec
Link length supported for copper is 1 m
cisco id is --
cisco extended id number is 4
number of lanes 1
```

Conditions: You see this issue when you are working with a MOLEX copper SFP (1.5m,2m,2.5m).

Workaround: None.

- 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.

- CSCun81705

  Symptom: You might see the snmpwalk/snmpget for an IPv6-specific instance return a no such instance error, while the getnext works OK.

  Conditions: This symptom might be seen when you are performing an snmp query for IP-MIB::ipNetToPhysicalPhysAddress.

  Workaround: Use snmpgetnext for that IPv6-specific instance.

- CSCuo22348

  Symptom: You might see one of the Layer 3 protocols flap when you issue the **show interface trunk** command.

  Conditions: This symptom might be seen when the system has multiple VDCs with M1 and F1 Series modules in same VDC. With these conditions, issuing this command can generate excessive traffic in Tx direction from CPU and may drop certain packets causing Layer 3 instability.

  Workaround: Do not issue this command under these circumstances.

- 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: None.

- 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: Use unicast mode.

- 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 switchID and does drive ES switchID. 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: Clear the MAC address dynamic for the destination MAC.

- CSCun76395

  Symptom: TACACS authorization request packets are not created by the switch, so all TACACS based authentication 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: Enable local login accounts, back up the VDC running configuration, delete and re-create VDC. The same configuration may be applied again once the VDC is back up

- 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.

- 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: None.

- 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: None.

- CSCun66506

  Symptom: NetScaler shows channel LA/1 bundle ports in down state, after their link is down.

  Conditions: You might see this symptom when you remove ports in quick succession among Nexus and NetScaler.

  Workaround: Bring up the affected links. Remove the links, waiting for 30 seconds before removing the next one each time.

- 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: Remove unsupported tunnel configuration.

- 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.

- 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: Set these configurations to the default values on both sides of the link.

- CSCum05788

  Symptom: You might see ping loss in the request path to the ping destination.

  Conditions: This symptom might be seen with a simple diamond topology. The source and destination client have single gateway devices that have ECMP to the destination. The gateway devices are connected via ECMP to two middle devices and then to the destination gateway that has the same mirrored topology. The failure is observed when on the source gateway the in use path is

broken by shutting the interface to the intermediate routers. The path in use is checked with trace route on the end clients. The end clients are windows PC running the ping utility that allows for modifying the ping frequency/response time.

Workaround: None

- CSCuo07507

  Symptom: The display after you enter the **show interface** *ethernet X/Y* **transceiver sprom** command, shows the QSFP transceiver revision as blank.

  Conditions: This symptom might be seen when you are working with a QSFP transceiver.

  Workaround: None

- CSCun26418

  Symptom: SPAN, OTV encapsulation, and multicast replication may fail because the replication ASIC responsible these processes stops replicating all traffic. This is particular to the SPAN/multicast replication pipeline, so working in an OTV environment you might 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 symptom might be seen when you are working in an OTV environment with a high rate of traffic egressing the OTV join-interface and with ERSPAN configured to span the join-interface.

  Workaround: Reload the module to clear the problem, and remove the ERSPAN session to prevent the issue from reoccurring.

- 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.

- CSCul66816

  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.

- CSCun24113

  Symptom: You might see the following message when there is a BGP flap in scale setup with 200K IPv4 routes:

  ```
  2014 Feb 19 11:57?28 CST3D5-D%Core %URIB-3-PTDEL_ERROR: urib[5879] (default-base)
  Could not delete 45.0.110.192/27 from URIB pt, error code r
  ```

  Conditions: This symptom might be seen when you are running high scale (more than200K IPv4 routes).

Workaround: Reduce the number of routes.

- 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.

- CSCun99702

  Symptom: If you are working with a running configuration with has one or more IPv6 Object-groups with wildcard masking, ISSD is being allowed from Cisco NX-OS Release 6.2(8) to any release that does not support IPv6 wildcard masking; that is, any previous release.

  Conditions: Object-group configurations with IPv6 wildcard masking must be present. With this configuration, an ISSD from the Cisco NX-OS Release 6.2(8) with support for IPv6 wildcard to a lower Release without support for IPv6 wildcard matching did not throw an incompatibility error.

  Workaround: Delete all the Object-group rules with wildcard masks and then perform the ISSD.

- 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: None; FabricPath ISIS restarts and fixes the issue.

- CSCun80157

  Symptom: You might see the OTV tunnel adjacency fail to program for one of the peers after VLAN goes up and down.

  Conditions: This symptom might be seen after back-to-back VLAN flapping.

  Workaround: For the VLANs, enter the shutdown command followed by the no shutdown command. If you still see problems, enter the same commands for the corresponding overlay on which the VLAN flapped.

- 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.

- CSCuo07626

  Symptom: You might see a map-request for an IPv6 EID not be resolved correctly.

  Conditions: This symptom might be seen when you get a link local IPv6 address for the locator when building a map request.

  Workaround: Issue the following command to configure the source locator to be used for LISP-encapsulated packets. To remove the configured source locator, issue the **no** form of this command:

  - **ip lisp source-locator** *interface*
  - **no ip lisp source-locator** *interface*

- CSCuo35291

  Symptom: You might see routing to a local subnet fails after you remove the LISP confguration.

  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: After you remove the LISP configuration, add the following commands to each interface that had LISP mobility enabled:

  interface *slot/number*

  ip verify unicast

  no ip verify unicast

- CSCun90161

  Symptom: When you issue the show rise command, the display might show an APBR entry as. stuck in ADD IN PRO state.

  Conditions: This symptom might be seen when you use a default route with the NS IP on netScaler. When a specific route for realer server IP address is removed on netScaler, netScaler attempts to install APBR entry with NS IP as next hop, During this process you might see some entries in ADD IN PRO state.

  Workaround: Issue the **shutdown** command followed by the **no shutdown** command, or issue the **usip -NO** command followed by the **usip -YES** command for the service that has the ADD IN PRO entries.

- 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 issuing 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.

- CSCuo02632

  Symptom: When you are working with IPv6 and you add a large number of interfaces, you might see both supervisors reload.

  Conditions: This symptom might be seen when you add a large number of ingress interfaces using the interface range. This symptom is rarely seen when >400 interfaces are added with range command.

  Workaround: Add the ingress interfaces one by one.

- 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: Delete the vPC leg port channel and recreate it by issuing the **channel-group** *number* command for member ports. This adds the correct mapping and displays it in the running configuration of the vPC leg port channel. Then make the port channel a vPC again.

- CSCuo25489

  Symptom: You might not see the private VLAN list in the display after you reload the switch by issuing the **copy r s** command, and then you issue 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: Delete and recreate the vPC leg port-channel.

- 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: Delete and recreate the vPC leg port-channel. This adds the correct mapping.

- 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 pvlan 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 recreate the PVLAN association between the primary and secondary VLANs.
  - You change the mode back by entering the **switchport mode private-vlan promiscuous** command.

  Workaround: Remove and recreate the promiscuous mapping configuration on the port for port channel.

- 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: Delete and recreate the vPC leg port-channel.

- 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: Take the following steps:

  - First disassociate the secondary VLAN from the primary VLAN.
  - Delete the primary and secondary VLANs.
  - Recreate the same primary and secondary VLANs and apply the configuration.
  - Wait for some time for the protocols to converge on both VLANs.
  - Finally, associate the secondary to the primary VLAN.

- CSCuo51846

  Symptom: If you have upgraded to the 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: Because the Cisco TAC is not responsible for the use of third-party modules and their possible malfunction, an immediate workaround is to use Cisco-branded transceivers. Alternately our future Cisco NX-OS Release 6.2(10) will provide appropriate fixes to resolve this issue.

Resolved Caveats—Cisco NX-OS Release 6.2(8)

- CSCum40651

  Symptom: You cannot perform a CLI [config cli or regular show cli] with more than 64 characters [including white spaces and full CLIs]. Cisco NX-OS will print 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 securityd 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 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 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.

- CSCum46499

  Symptom: The SNMP application might fail to send heartbeat messages to the System Manager when it is in a state with heavy polling, and the SNMP application could restart.

  Conditions: This symptom might be seen when you configure the device for SNMP management.

  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.

- CSCul52399

  Symptom: When you enter the **switchport monitor** command, the system might cause the port's CBL to be programmed to the wrong value, and you might see unexpected traffic sent out the SPAN destination port.

  Conditions: This symptom might be seen when you are working on a M1 Series module with the ports in dedicated mode.

  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.

- CSCui10084

Symptom: In Python mode, you cannot change Python's socket operations to the default VRF from the management VRF.

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 have 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.

- CSCuo42047

  Symptom: When you are running Cisco NX-OS Release 6.2(6) or Release 6.2(6a), because of a specific sequence of PVLAN events the PIXM process can experience memory corruption, which can lead to one of the following issues

  - PC and its members are error disabled
  - Incorrect CBL programming
  - PIXM crash.

  Conditions: This symptom might be seen when you are running Cisco NX-OS Release 6.2(6) or Release 6.2(6a). You might encounter this symptom if you are using port channels and private VLANs, and you have the following conditions:

  - The primary VLAN is associated to the secondary VLAN.
  - The port channel is a trunk port channel carrying both primary and secondary VLANs
  - STP operates on either the primary or the secondary VLAN on the trunk port channel.
  - Any modify operation (ADD/REMOVE) on the port channel will result in above symptoms.

  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.

Resolved Caveats—Cisco NX-OS Release 6.2(6a)

- CSCui53828

  Symptom: You cannot filter or modify OSPF routes before sending them to the RIB.

  Conditions: This symptom might be seen when you attempt to filter or modify OSPF routes before sending them to the RIB.

  Workaround: This issue is resolved.

- 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.

- CSCtz59444

  Symptom: You might be unable to add ports to a port channel; when you attempt to add a port, the member links move to the error-disabled state.

  Conditions: This symptom might be seen when the port channel or member is deleted, and the cleanup does not happen properly in ELTM. It can also appear in an ISSU if there are new updates and the port-channel configuration is manually modified.

  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.

- CSCul93937

  Symptom: The switch unexpectedly writes out a L2FM core during an upgrade to the Release 6.1.(3).

  Conditions: This symptom might be seen when you are performing an ISSU when you are running Release 6.1.(3).

  Workaround: This issue is resolved.

- CSCug40835

  Symptom: A switch that is running Cisco Nexus Release 6.1(2) faces an SPM service crash after you apply a policy on many interfaces in a single command.

  Conditions: This symptom might be seen when you are working with a supervisor 1.

  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 ACEs 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.

- CSCul69817

  Symptom: The LDP core might come down.

  Conditions: This symptom might be seen after the Netstack process is restarted.

  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.

- CSCul70142

  Symptom: The switch reloads.

  Conditions: This symptom might be seen when the sac_usd process crash is followed by snmpd and m2rib process cores.

  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.

- CSCui99899

  Symptom: The ARP entry for the server in the HSRP standby switch is not correct and it is pointing the server IP to the SVI VMAC instead of the host's own MAC. You might lose connectivity.

  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.

- 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.

- CSCuh83775

  Symptom: Trace route prints ?*? for the switch nodes when packets come in through an SVI interface in the MPLS environment.

  Conditions: This symptom might be seen you are working with MPLS.

  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.

- CSCuj51290

  Symptom: The WCCP service restarts on a switch.

  Conditions: This symptom might be seen if you have enabled WCCP.

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.

- CSCul52268

  Symptom: Failover does not occur until the ARP/ND cache expires in IP failover scenarios where a shared MAC address is used with OTV arp-nd-cache.

  Conditions: This symptom might be seen when the members of the virtual IP are on different sides of an OTV adjacency or when the virtual IP address of the device uses BIAs of the NICs.

  Workaround: This issue 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.

- CSCui06433

  Symptom: The available extensions for the **show tech** command need to be enhanced to improve MTTR.

  Conditions: This symptom might be seen when you are working in the **show tech** command.

  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.

- CSCud89826

  Symptom: The switch has no configuration option to present, only the hostname in the CDP Neighbor Device ID field.

  Conditions: This symptom might be seen when you are working with CDP.

  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)

- 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.

- 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.

- CSCue93131

  Symptom: Because a physical port vPC uses a Layer 2 FM MAC address move, there is a performance issue when the peer link is an M Series.

  Conditions: This symptom might be seen when the peer link is an M Series in a physical port vPC setup.

  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.

- CSCuh23105

  Symptom: Forwarding plane VLAN translation can be out of synchronization with a configured translation when multiple overlays are used.

Conditions: This symptom might be seen when the VLAN mapping X -> A is defined under overlay1, but X is then extended under overlay2. You can enter this configuration from the CLI, but it is a misconfiguration and can lead to incorrect translations.

Workaround: This issue is resolved.

- CSCuh79244

  Symptom: The **show l2vpn signaling rib** command and the **show l2vpn rib** command do not work correctly.

  Conditions: This symptom might be seen under normal operating conditions for a Cisco Nexus 7000 Series device.

  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.

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.

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