Cisco Nexus 7000 Series NX-OS Release Notes, Release 4.0

Date: February 2, 2009
Part Number: OL-16034-05 I0

This document describes the features, caveats, and limitations for Cisco NX-OS software. Use this document in combination with documents listed in the "Related Documentation" section on page 62.

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 NX-OS Release Notes: *http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_0/nx-os/release/notes/401_nx-os_release_note.html*

Note Table 1 shows the online change history for this document.

Table 1 Online History Change

Part Number	Revision	Date	Description
OL-16034-01	A0	April 1, 2008	Created release notes.
OL-16034-02	A0	April 21, 2008	Created release notes for Release 4.0(1a).
	B0	April 24, 2008	Added CSCso92283 to the open caveats.
	C0	April 25, 2008	Removed references to Data Center Network Manager (DCNM)
	D0	May 7, 2008	Added CSCso84540 and CSCsq03175 to the open caveats.
	E0	May 8, 2008	Removed CSCsq03175 from the open caveats,

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

Table 1 Online History Change

Part Number	Revision	Date	Description
OL-16034-03	A0	June 13, 2008	Created release notes for Release 4.0(2).
	B0	June 16, 2008	Removed CoPP default policy assignment from the "Cisco NX-OS Release 4.0(2)" section on page 5.
	C0	June 17, 2008	Added CoPP default policy assignment to the "Cisco NX-OS Release 4.0(2)" section on page 5.
	D0	June 24, 2008	Added resolved caveat CSCsq60582.
OL-16034-04	A0	August 21, 2008	Created release notes for Release 4.0(3).
	B0	August 22, 2008	Added CSCsq47196 to the open caveats, moved CSCsr56858 from the open caveats to the resolved caveats, and added CSCsr39659 to the resolved caveats.
	C0	August 25, 2008	Added CSCsr30773 to the resolved caveats.
	D0	August 29, 2008	Added CSCsr96589 to the open caveats.
	E0	September 11, 2008	Added CSCsu41395 and CSCsu45752 to the open caveats.
OL-16034-05	A0	November 3, 2008	Created release notes for Release 4.0(4)
	B0	November 7, 2008	Added CSCsv47908 to the open caveats.
	C0	November 10, 2008	Added CSCsv49677 to the open caveats.
	D0	November 22, 2008	Added CSCsv84522 to the open caveats.
	E0	November 26, 2008	Added specific information about no new features in Release 4.0(4).
	F0	December 11, 2008	Added CSCsw35996 to the open caveats.
	G0	December 18, 2008	Added CSCsw64054 to the open caveats.
	H0	February 2, 2009	Removed CSCsw35996 from the open caveats.
	I0	September 14, 2009	Added a Limitation about multicast over tunnel interfaces.

Contents

This document includes the following sections:

Introduction

The Cisco NX-OS software is a data center-class operating system that is based on the Cisco SAN-OS software.

The Cisco NX-OS software 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:

Hardware Supported

Cisco NX-OS supports the Nexus 7000 Series 10-slot chassis. You can find detailed information about supported hardware in the *Cisco Nexus 7000 Series Hardware Installation and Reference Guide*.

Memory Requirements

Cisco NX-OS requires 4 GB of memory.

New Software Features

This section briefly describes the new features introduced in the releases of the Cisco NX-OS software. For detailed information about the features listed, see the documents listed in the "Related Documentation" section on page 62. 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:

Cisco NX-OS Release 4.0(4)

There are no new hardware or software features or enchancements in Release 4.0(4).

Cisco NX-OS Release 4.0(3)

This section briefly describes the new features introduced in this release and includes the following topics:

IPv6 Routing Protocols—OSPFv3 and PIM ASM

Open Shortest Path First version 3 (OSPFv3) is a link-state protocol that uses Dijkstra's algorithm to find the shortest path to a destination. OSPFv3 is defined in IETF RFC 2740. OSPFv3 expands on OSPFv2 to provide support for IPv6 routing prefixes and the larger size of IPv6 addresses. OSPFv3 uses link-local IPv6 addresses for neighbor discovery and IPv6 for authentication.

Protocol Independent Multicast Any Source Multicast (PIM ASM) for IPv6 provides support for IPv6 addresses and is called PIM6.

Tunnels (GRE)

Cisco NX-OS supports Generic Route Encapsulation (GRE) tunnels.

Tunneling allows you to encapsulate arbitrary packets inside a transport protocol. This feature is implemented as a virtual interface to provide a simple interface for configuration. The tunnel interface provides the services necessary to implement any standard point-to-point encapsulation scheme.

VRRP

Virtual Routing Redundancy Protocol (VRRP) allows for a transparent failover at the first-hop IP router, by configuring a group of routers to share a virtual IP address. VRRP selects a master router in that group to handle all packets for the virtual IP address. The remaining routers are in standby and take over in the event that the master router fails.

SNMP Multiple Instances

Cisco NX-OS supports the CISCO-CONTEXT-MAPPING-MIB to map between Simple Network Management Protocol (SNMP) contexts and logical network entities. You can associate an SMNP context to a virtual routing and forwarding instance (VRF), protocol instance, or topology.

CMP Enhancements

The following enhancements to the Connectivity Management Processor (CMP) were added in Release 4.0(3):

- Updates to the user authentication process—The CMP accepts users with network-admin privileges. When the control processor (CP) and the CMP are both operational, you can log into the CMP using your NX-OS-configured username and password. The CP synchronizes the password for admin user with the CMP. When the CP is not operational, you can log into the CMP using the admin userid and password.
- Logging levels—You can save up to 256 CMP messages in a log file, and you can specify a severity threshold for the saved messages. Currently, the CMP records alert level and critical level messages. When the file has 256 messages, the CMP automatically removes the oldest message when it saves a new message. You can also clear the log file of all saved messages.
- Configuring serial communication settings—Use the following commands to configure serial communication characteristics for the CMP:
  - Speed (baud rate) (300 to 115,200 baud) by using the **speed** command
  - Number of bits (5 to 8) in a transmitted character by using the **databits** command
  - Parity checking (even, odd, or none) by using the **parity** command
  - Asynchronous stop bits (1 or 2) by using the **stopbits** command
  - Hardware version of the flow control (enable or disable) by using the **flowcontrol** or **no flowcontrol** command

Cisco NX-OS Release 4.0(2)

This section briefly describes the new features introduced in this release and includes the following topics:

Telnet IPv6 Support

You can use the **telnet6** command to create Telnet sessions with IPv6 addressing.

CoPP Configuration Status

You can use the **show copp status** command to display the control plane policing (CoPP) configuration status information.

Prestandard MST Interoperability

Although the Cisco NX-OS software does not run prestandard Multiple Spanning Tree (MST), the NX-OS software allows an interface running MST to respond with a prestandard MST message if it receives a prestandard message from the device at the other end of a link. In Cisco NX-OS Release 4.0(2) and later releases, you can force the interface running MST to send prestandard, rather than standard, MST messages using the **spanning-tree mst pre-standard** command in interface configuration mode. This example shows how to enable prestandard MST interoperability:

```
switch# configure terminal
switch(config)# interface ethernet 2/1
switch(config-if)# spanning-tree mst pre-standard
```

EIGRP Maximum Paths Default Change

The default number of EIGRP maximum path changed to 8.

CoPP Default Policies

You can assign a different default CoPP policy using the **setup** command at the CLI prompt. Also, the CoPP default policies have the following changes:

- Added Secure Shell FTP (SFTP) to the copp-system-class-management class map.
- Added access-control lists (ACLs) to match the source ports for TACACS+, RADIUS, Network Time Protocol (NTP), FTP, TFTP, SFTP, Secure Shell (SSH), and Telnet.
- Increased the policing bandwidth to 10 Mbps for the copp-system-class-management class.

Cisco NX-OS Release 4.0(1a)

This section briefly describes the new features introduced in this release and includes the following topics:

- QoS Maximum Policing Rate Increased, page 6

QoS Maximum Policing Rate Increased

The QoS maximum policing rate is increased to 80 Gbps.

Cisco NX-OS Release 4.0(1)

This release is the initial release of the Cisco NX-OS software and includes features in the following categories:

- Software Compatibility, page 7
- Serviceability, page 7
- Manageability, page 9
- Layer 2 Switching, Layer 3 Routing, and IP Services, page 10
- Quality of Service, page 12

Software Compatibility

Cisco NX-OS Release 4.0(1) interoperates with Cisco products that run any variant of the Cisco IOS software operating system. Cisco NX-OS Release 4.0(1) also interoperates with any networking operating system that conforms to the networking standards listed as supported in the *Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide, Release 4.0*.

This section includes the following topics:

Common Software Throughout the Data Center

The Cisco NX-OS software provides a unified operating system (OS) that is designed to run the data center network LAN and Layer 4 through Layer 7 network services. The NX-OS software integrates technologies such as Ethernet, Layer 4 through Layer 7 services (such as firewall services), and virtualization.

Modular Software Design

The Cisco NX-OS software supports distributed multithreaded processing on symmetric multiprocessors (SMPs), multi-core CPUs, and distributed module processors. Computationally intensive tasks, such as hardware table programming, can be offloaded to dedicated processors distributed across the modules. The Cisco NX-OS software creates modular processes on demand, each in a separate protected memory space. These processes are started and system resources allocated only when a feature is enabled. A real-time preemptive scheduler helps to ensure the timely processing of critical functions.

Virtual Device Contexts

The Cisco NX-OS software can segment OS and hardware resources into virtual contexts that emulate virtual devices. Each virtual device context (VDC) has its own software processes, dedicated hardware resources (interfaces), and an independent management environment. With VDCs, you can consolidate separate networks onto a common infrastructure, maintaining the administrative boundary separation and fault isolation characteristics of physically separate networks while providing many of the operational cost benefits of a single infrastructure. For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0*.

Serviceability

The Cisco NX-OS software has serviceability functions that allow you take early action based on network trends and events. These features help with network planning and improving response times.

This section includes the following topics:

Switched Port Analyzer

The switched port analyzer (SPAN) feature allows you to analyze all traffic between ports (called the SPAN source ports) by nonintrusively directing the SPAN session traffic to a SPAN destination port that has an external analyzer attached to it. For more information about SPAN, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

Ethanalyzer

Ethanalyzer is a Cisco NX-OS protocol analyzer tool based on the Wireshark (formerly Ethereal) open source code. Ethanalyzer is a command-line version of Wireshark for capturing and decoding packets. You can use Ethanalyzer to troubleshoot your network and analyze the control-plane traffic. For more information about Ethanalyzer, see the *Cisco Nexus 7000 Series NX-OS Troubleshooting Guide, Release 4.0*.

Call Home

The Call Home feature continuously monitors hardware and software components to provide e-mail-based notification of critical system events. A versatile range of message formats is available for optimal compatibility with pager services, standard e-mail, and XML-based automated parsing applications. It offers alert grouping capabilities and customizable destination profiles. This feature can be used, for example, to directly page a network support engineer, send an e-mail message to a networks operation center (NOC), and employ Cisco AutoNotify services to directly generate a case with the Cisco Technical Assistance Center (TAC). This feature enables networking devices to inform IT when a problem occurs and helps to ensure that the problem is acted on quickly, reducing the time for a resolution and maximizing the system uptime. For more information about Call Home, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

Online Diagnostics

The Cisco generic online diagnostics (GOLD) are a suite of diagnostic facilities to verify that hardware and internal data paths are operating as designed. Boot-time diagnostics, continuous monitoring, and on-demand and scheduled tests are part of the Cisco GOLD feature set. GOLD allows rapid fault isolation and continuous system monitoring. For information about configuring GOLD, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

Embedded Event Manager

The Embedded Event Manager (EEM) is a device management technology built into the Cisco NX-OS software. EEM allows you to customize the behavior of the device based on network events as they occur. For information about configuring EEM, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

NetFlow

The Cisco NX-OS NetFlow implementation supports version 5 and version 9 exports. It also supports the Flexible NetFlow configuration model and hardware-based Sampled NetFlow for enhanced scalability. For more information about NetFlow, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

Manageability

This section includes the following topics:

Simple Network Management Protocol

The Cisco NX-OS software is compliant with Simple Network Management Protocol (SNMP) version 1, version 2, and version 3. A rich collection of Management Information Bases (MIBs) is supported. For more information about SNMP, see the
Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0.

Configuration Verification and Rollback

With the Cisco NX-OS software, you can verify the consistency of a configuration and the availability of necessary hardware resources prior to committing the configuration. You can preconfigure a device and apply the verified configuration at a later time. Configurations also include checkpoints that allow you to roll back to a known good configuration as needed. For more information about rollbacks, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

Role-Based Access Control

With role-based access control (RBAC), the Cisco NX-OS software enables you to limit access to device operations by assigning roles to users. You can customize access and restrict it to the users who require it. For more information about RBAC, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 4.0*.

Connectivity Management Processor

The Cisco NX-OS software supports the use of a Connectivity Management Processor (CMP) for lights-out remote platform management. The CMP provides an out-of-band access channel to the device console. For more information about CMP, see the *Cisco Nexus 7000 Series Connectivity Management Processor Configuration Guide*.

Cisco NX-OS Device Configuration Methods

You can configure NX-OS features on your device using the following methods:

- CLI—You can configure devices using the CLI from an SSH session or a Telnet session. SSH provides a secure connection to the device. The CLI command references are organized by feature. For more information, see the *Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide, Release 4.0* or the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 4.0*.
- XML management interface over SSH—You can configure devices using the XML management interface, which is a programmatic method based on the NETCONF protocol that complements the CLI functionality. For more information, see the *Cisco Nexus 7000 Series NX-OS XML Management Interface User Guide, Release 4.0*.

Layer 2 Switching, Layer 3 Routing, and IP Services

This section includes the following topics:

Ethernet Switching

The Cisco NX-OS software supports high-density, high-performance Ethernet systems and provides the following data center-class Ethernet switching features:

- IEEE 802.1D-2004 Rapid and Multiple Spanning Tree Protocols (802.1w and 802.1s)
- IEEE 802.1Q VLANs and trunks
- 16,000-subscriber VLANs
- IEEE 802.3ad link aggregation
- Private VLANs
- Cross-chassis private VLANs
- UniDirectional Link Detection (UDLD) in aggressive and standard modes
- Traffic suppression (unicast, multicast, and broadcast)

Spanning Tree Protocol enables transparent upgrades using in-service software upgrades (ISSUs) in Spanning Tree Protocol environments, Bridge Protocol Data Unit (BPDU) guard, loop guard, root guard, BPDU filters, bridge assurance, and jumbo frame support.

For more information, see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 4.0* and the
Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide, Release 4.0.

IP Unicast Routing

The Cisco NX-OS software supports IP versions 4 and 6 (IPv4 and IPv6) and the following routing protocols:

- Open Shortest Path First (OSPF) Protocol Versions 2 (IPv4) and 3 (IPv6)
- Intermediate System-to-Intermediate System (IS-IS) Protocol
- Border Gateway Protocol (BGP)

- Enhanced Interior Gateway Routing Protocol (EIGRP)
- Routing Information Protocol Version 2 (RIPv2)

The implementations of these protocols are fully compliant with the latest standards and include 4-byte autonomous system numbers (ASNs) and incremental Shortest Path First (SPF). All unicast protocols support Non-Stop Forwarding Graceful Restart (NSF-GR). All protocols support all interface types, including Ethernet interfaces, switched virtual interfaces (VLAN interfaces) and subinterfaces, port channels, tunnel interfaces, and loopback interfaces. For more information, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 4.0.*

Layer 3 and Layer 2 Multicast

The Cisco NX-OS software includes the following multicast protocols and functions:

- Protocol Independent Multicast Version 2 (PIMv2)
  - Source Specific Multicast (SSM)
  - PIM sparse mode (Any-Source Multicast [ASM] for IPv4 and IPv6)

Note Cisco NX-OS does not support PIM dense mode.

  - Bidirectional Protocol Independent Multicast (Bidir PIM)
- Anycast rendezvous point (Anycast-RP)
- Multicast NSF for IPv4 and IPv6
- RP-Discovery using bootstrap router (BSR): Auto-RP and static
- Internet Group Management Protocol (IGMP) Versions 1, 2, and 3 router role
- IGMPv2 host mode
- IGMP snooping
- Multicast Listener Discovery (MLD) Protocol Version 2 (for IPv6)
- Multicast Source Discovery Protocol (MSDP) (for IPv4 only)

All protocols support all interface types, including Ethernet interfaces, switched virtual interfaces (VLAN interfaces) and subinterfaces, port channels, tunnel interfaces, and loopback interfaces.

For more information, see the *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide, Release 4.0.*

IP Services

The following IP services are available in the Cisco NX-OS software:

- Virtual Routing and Forwarding (VRF)

  All routing protocols and IP services are VRF aware.
- Dynamic Host Configuration Protocol (DHCP) helper
- Hot-Standby Routing Protocol (HSRP)
- Gateway Load Balancing Protocol (GLBP)
- Enhanced object tracking
- Policy-Based Routing (PBR)

- Graceful restart routing protocol extensions

For more information about the IP services, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 4.0*.

The Cisco NX-OS software also supports Unicast Reverse Path Forwarding (Unicast RPF).

For more information about Unicast RPF, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 4.0*.

Quality of Service

The Cisco NX-OS software provides Quality of Service (QoS) functions for classification, marking, queuing, policing, and scheduling. The Modular QoS CLI (MQC) supports all QoS features. You can use MQC to uniformly configure QoS across various Cisco platforms. For more information, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide, Release 4.0*.

Network Security

Note We recommend that you perform ACL configurations using the Session Manager. This feature allows you to verify ACL configuration and confirm that the required resources are available before committing the ACL configuration to the running configuration. For more information about Session Manager, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

This section includes the following topics:

- Cisco TrustSec, page 12
- Additional Network Security Features, page 12

Cisco TrustSec

The Cisco TrustSec security feature provides data confidentiality and integrity and supports standard IEEE 802.1AE link-layer cryptography with 128-bit Advanced Encryption Standard (AES) cryptography. Link-layer cryptography guarantees end-to-end data privacy while allowing the insertion of security service devices along the encrypted path. Security-group access control lists (SGACLs) are based on security group tags instead of IP addresses, which enables policies that are more concise and easier to manage due to their topology independence. For more information, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 4.0*.

Additional Network Security Features

In addition to Cisco TrustSec, the Cisco NX-OS software includes the following security features:

- Data path intrusion detection system (IDS) for protocol conformance checks
- Control Plane Policing (CoPP)
- Message-digest algorithm 5 (MD5) routing protocol authentication
- Cisco integrated security features, including Dynamic Address Resolution Protocol (ARP) inspection (DAI), DHCP snooping, and IP Source Guard
- Authentication, authorization, and accounting (AAA)
- RADIUS and TACACS+

- Secure Shell (SSH) Protocol Version 2
- Port security
- IEEE 802.1X authentication
- Layer 2 Cisco Network Admission Control (NAC)
- Policies based on MAC and IPv4 addresses supported by named ACLs (port-based ACLs [PACLs], VLAN-based ACLs [VACLs], and router-based ACLs [RACLs])

For more information about the above features, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 4.0*.

The Cisco NX-OS software also supports SNMP Version 3 (SNMPv3)

For more information about SNMPv3, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0*.

Licensing

The Cisco NX-OS licensing feature allows you to access premium features on the device after you install the appropriate license for that feature. Any feature not included in a license package is bundled with the Cisco NX-OS software and is provided to you at no extra charge.

You must purchase and install a license for each device.

Note With the exception of the Cisco TrustSec feature, you can enable a feature without installing its license. The Cisco NX-OS software gives you a grace period during which time you can try out a feature before purchasing its license. You must install the Advanced Services license package to enable the Cisco TrustSec feature.

For detailed information about the features that require licensing and NX-OS license installation, see the *Cisco Nexus 7000 Series NX-OS Licensing Guide, Release 4.0*.

For information about troubleshooting licensing issues, see the *Cisco Nexus 7000 Series NX-OS Troubleshooting Guide, Release 4.0*.

Limitations

This section describes the limitations in Cisco NX-OS Release 4.0.

This section includes the following topics:

Cisco TrustSec

The Cisco NX-OS Release 4.0(2) and earlier releases do not fully support the following commands:

- **clear cts cache**
- **clear cts policy**
- **cts cache**
- **cts l3 spi** (global configuration)
- **cts l3 spi** (interface configuration)
- **show cts l3 interface**
- **show cts l3 mapping**

SNMP MIB Traps

Cisco NX-OS Release 4.0(2) and earlier releases support only SNMP MIBs and traps for the default VRF of the first instance of the Layer 3 protocol.

QoS

Cisco NX-OS Release 4.0(2) and earlier releases do not support egress policing on VLAN interfaces.

Tunnel Interfaces and VRFs

Cisco NX-OS Release 4.0(3) and earlier releases support assigning tunnel interfaces only to the default VDC and the default Virtual Routing and Forwarding instance (VRF).

VLANs

The Cisco Nexus series 7000 device can scale a maximum of 4000 VLANs across the entire system. These VLANs can be configured in single VDC or across multiple VDCs. If the total number of VLANs configured on the device across all VDCs exceeds 4000, there are known issues with multiple modules.

Multicast over Tunnel Interfaces

In Cisco NX-OS Release 4.0(3) and earlier releases, tunnel interfaces do not support Protocol-Independent Multicast (PIM).

Caveats

This section includes the following topics:

Open Caveats—Cisco NX-OS Release 4.0(4)

- CSCsll44778

  Symptom: A maximum of 60 concurrent SSH and Telnet sessions are supported.

  Conditions: If more than 60 concurrent SSH and Telnet sessions are attempted, the results are unpredictable.

  Workaround: No workaround.

- CSCsl71366

  Symptom: A maximum of 200 VRFs are supported.

  Conditions: If more than 200 VRFs are configured, the results are unpredictable.

  Workaround: No workaround.

- CSCsl97793

  Symptom: ACL logging does not occur for packets matched by software ACL processing.

  Conditions: Packets processed in the software are not logged when they match an ACL with logging enabled.

  Workaround: No workaround.

- CSCsm09007

  Symptom: QoS match-all criteria is not supported.

  Conditions: When you configure match all for a QoS class map using the **class-map type qos match-all** command, the **match-all** option does not work. Instead, the match criteria is always treated as match any.

  Workaround: No workaround.

- CSCsm13589

  Symptom: Record-route does not work correctly when Policy Based Routing (PBR) is configured.

Conditions: Any IP traffic redirected due to PBR is not sent to the supervisor module. As a result, record-route does not work for packets redirected due to PBR.

Workaround: No workaround.

- CSCsm15545

  Symptom: Adjacency statistics reset after a supervisor module switchover.

  Conditions: The counter values in the output of **show ip adjacency** {**statistics** | **detail**} command are cleared after a supervisor module switchover.

  Workaround: No workaround.

- CSCsm22329

  Symptom: QoS statistics require a policing action in order for marking actions to produce statistics.

  Conditions: When you define a QoS service policy with only marking actions, the statistics do not work. The statistics features works only when the service policy has a policing action defined also.

  Workaround: You can get statistics for marking only policy by applying a dummy policing action to the policies. For example, in addition to the marking actions, you should define a policing action that permits 100 percent traffic. Configure the violate and conform action as transmit.

- CSCsm63331

  Symptom: The on-demand diagnostics for the port loopback test are not supported on the 32-port 10-Gbps Ethernet modules.

  Conditions: The **show diagnostic result module command** output indicates untested (U) for the 32-port 10-Gbps Ethernet modules after on-demand diagnostic testing of the port loopback feature with the **diagnostic start module** command.

  Workaround: No workaround.

- CSCsm70593

  Symptom: An interface is disabled when more than 50,000 port-VLAN instances go down at the same time.

  Conditions: When more than 50,000 port-VLAN instances go down at the same time, the interface times out and becomes disabled. The following system message displays:

  ```
  %$ VDC-1 %$ %ETHPORT-2-SEQ_TIMEOUT: Component MTS_SAP_L2FM timed out on response to
  opcode:MTS_OPC_ETHPM_PORT_LOGICAL_CLEANUP (for:RID_PORT: Ethernet9/46)
  ```

  Workaround: No workaround. This message is not seen when less than 500,000 Port-VLAN instances go down.

- CSCsm75863

  Symptom: Logging to an external syslog server using an IPv6 address does not work.

  Conditions: If you configure IPv6 addresses for an external syslog server, then logging does not work for the server.

  Workaround: No workaround.

- CSCsm79619

  Symptom: Removing the management IP and VRF configuration with the **write erase boot** command does not work.

  Conditions: The **write erase boot** command does not remove the management IP and VRF configuration.

  Workaround: To erase the management IP or VRF configuration, use the following command sequence:

  1. **write erase**
  2. **write erase boot**

- CSCsm98229

  Symptom: A checkpoint creation or rollback operation can fail when an in-service software upgrade (ISSU) is in progress.

  Conditions: If you roll back the configuration or create a checkpoint while an ISSU is in progress, then the rollback or checkpoint creation operation can fail.

  Workaround: Avoid performing a checkpoint creation or rollback operation at the same time while an ISSU is in progress. Instead, perform the checkpoint creation or rollback operation after the ISSU is complete.

- CSCsm98733

  Symptom: One checkpoint is missing after a supervisor module switchover.

  Conditions: If the ascii-cfg-server process restarts or if the active supervisor module switches over to the standby supervisor module while a checkpoint operation is in progress, then the checkpoint operation may not complete.

  Workaround: Recreate the checkpoint after a supervisor module switchover if the checkpoint is missing.

- CSCso02550

  Symptom: CoPP crashes with large policy maps.

  Conditions: CoPP crashes if you attach more than 300 classes to the policy map.

  Workaround: Reduce the number of classes attached to the CoPP policy map.

- CSCso03128

  Symptom: There is no warning that configuration changes are not saved.

  Conditions: Under the following conditions, the device does not warn you about unsaved changes:

  - If you exit after making an additional change while the running configuration is being copied to the startup configuration.
  - When you reload the device and you have not made any configuration changes since the last time the running configuration was copied to the startup configuration.

  Workaround: No workaround.

- CSCso03889

  Symptom: Address Resolution Protocol (ARP) ACLs are not supported on private VLANs.

  Conditions: If you configure an ARP ACL on a primary VLAN using the **ip arp inspection filter** *vlan-id* command, it is not propagated to the secondary VLAN.

  Workaround: No workaround.

- CSCso09082

  Symptom: The “use burn-in address (BIA)” feature for HSRP is not automatically applied to the main interface and all subinterfaces.

  Conditions: If you configure HSRP to use the BIA for an interface or subinterface using the **hsrp use-bia** command, the configuration is only applied to that interface or subinterface. The configuration is not, then, also applied to the main interface and all subinterfaces.

  Workaround: Manually enter the **hsrp use-bia** command for all the interfaces and subinterfaces on which it is required.

- CSCso27690

  Symptom: The device name does not display with the login prompt.

  Conditions: If you configure a device name using the **switchname** command, the name does not display at the login prompt on the standby.

  Workaround: If a supervisor module switchover occurs, the device name can be restored on the new active supervisor module by reentering the **switchname** command.

- CSCso31974

  Symptom: If you open the ejector levers on the supervisor and reload the chassis, the supervisor module attempts to come up and as the ejector levers are detected as open, the system reloads the supervisor module again. This situation results in the standby supervisor module going through repeated reboot cycles.

  Conditions: This symptom occurs when you attempt to reload the chassis with the supervisor module still seated but with the ejector levers open.

  Workaround: Ensure that you either completely remove the supervisor module from the chassis or insert the supervisor module completely into the chassis and close the ejector levers before you reload the chassis.

- CSCso43538

  Symptom: IGMP reports received on a VLAN interface cannot be policed with CoPP.

  Conditions: IGMP reports and queries received on a VLAN interface are not subjected to control plane policing. The packets can only be rate limited using the receive rate limiter.

  Workaround: Configure the **platform rate-limit receive** to rate limit IGMP packets received on VLAN interfaces.

Note The receive rate limiter matches and also rate limits all packets sent to the supervisor module. It does not differentiate IGMP traffic from other data traffic.

- CSCso43922

  Symptom: If significant traffic triggers ICMP redirects, it can cause the loss of OSPF adjacency.

  Conditions: ICMP redirect is enabled by default on all Layer 3 interfaces. If enough traffic is present to trigger ICMP redirects, it can affect OSPF control traffic. If OSPF packets are dropped because data packets are being copied to the supervisor module for ICMP redirect, it can lead to OSPF adjacency loss.

  Workaround: Disable ICMP redirect on Layer 3 interfaces by using the **no ip redirects** command in interface configuration mode.

- CSCso46631

  Symptom: No binding entries are created for VLAN 1 when you enable DHCP snooping on a trunk interface with multiple VLANs.

  Conditions: After you enable DHCP snooping on a trunk interface that has multiple VLANs, the NX-OS software creates binding entries for all VLANs except VLAN 1.

  Workaround: Do not use VLAN 1 as a trunking VLAN.

- CSCso74111

  Symptom: The device does not apply the shutdown process for following OSPF and OSPFv3 commands: the **shutdown** command in the router configuration mode and the **ip ospf/ospf3 shutdown** command in the interface configuration mode.

  Conditions: This situation occurs under all conditions.

  Workaround: Enter the **no** form of the command and then reenter the shutdown command.

- CSCsq04350

  Symptom: A VRF remains in the Admin Down pending state after a VRF shutdown and supervisor module switchover.

  Conditions: If you perform a supervisor module switchover immediately after shutting down a VRF, the VRF remains in the Admin down pending state.

  Workaround: Wait from 5 to 10 seconds after shutting down the VRF before you perform a supervisor module switchover.

- CSCsq06161

  Symptom: Configuring Equal Cost Multipath Protocol (ECMP) load sharing may cause some packets to be duplicated in some exceptions.

  Conditions: Packets that are sent to the software because of the same interface exception may be forwarded in both the hardware and software.

  Workaround: No workaround.

- CSCsq25183

  Symptom: With more than 1,000 interfaces or subinterfaces in the startup configuration, the device may fail.

  Conditions: If you are running an extremely large startup configuration, such as more than 1,000 interfaces or subinterfaces, the configuration server may exhaust its memory and fail.

  Workaround: No workaround.

- CSCsq28404

  Symptom: The IP EIGRP topology table does not show the next hop after changing the delay.

  Conditions: After you change the delay and enter the **show ip eigrp topology** command, the next hop information displayed is incorrect.

  Workaround: No workaround.

- CSCsq43292

  Symptom: Changing the LACP hello timers from normal to fast or from fast to normal may not work.

  Conditions: This symptom can occur in all conditions.

  Workaround: Configure the port channels in on mode, rather than using LACP.

- CSCsq44385

  Symptom: A rollback does not work correctly if the NetFlow record is modified.

  Conditions: If a NetFlow record is modified during a rollback, the rollback does not work properly.

  Workaround: If you are using a rollback, create a different NetFlow record.

- CSCsq66001

  Symptom: The tunnel interface is not detected when you are processing an SNMP MIB walk.

  Conditions: This situation occurs under all conditions and does not affect functionality.

  Workaround: No workaround.

- CSCsq73090

  Symptom: When you enter the **show interface tunnel** *number* command, the device displays the operational state of the tunnel as up when that tunnel source interface is down.

  Conditions: This situation occurs under all conditions.

  Workaround: Enter the **no shutdown** command on the tunnel source interface to bring it up or configure another interface as the tunnel source.

- CSCsq74911

  Symptom: The show blink function that displays the blink/beacon status for all devices is not available.

Conditions: This symptom exists under all conditions.

Workaround: No workaround.

- CSCsq79703

  Symptom: NX-OS supports only prefix length; it does not support wildcard masks that have a 0 bit anywhere after the first 1 bit. You cannot have an ACL that offers the same granularity as Cisco IOS ACL provides.

  Conditions: This symptom occurs under all conditions.

  Workaround: No workaround.

- CSCsq95595

  Symptom: The **clear counters** command does not clear the counters for tunnel interfaces.

  Conditions: This situation occurs under all conditions.

  Workaround: No workaround.

- CSCsr07444

  Symptom: When tracking a Layer 2 interface using Virtual Router Redundancy Protocol (VRRP), the VR priority is not correctly updated.

  Conditions: When you configure VRRP to track a Layer 2 interface, the VR priority is not updated correctly to reflect the state of the interface.

  Workaround: Enter the **shutdown** and **no shutdown** commands for the specified interfaces.

- CSCsr43915

  Symptom: You cannot work with EIGRP multi-instance MIBs without defining the SNMP context.

  Conditions: This symptom occurs when you are running more than one instance of EIGRP on a single device or operating an EIGRP process in a nondefault VRF.

  Workaround: Create an SNMP context on the switch by entering the **snmp-server context** *context-name* **instance** *instance-name* **vrf** *vrf-name* **topology** *topology-name* command. When you are using SNMPv3, supply the context name in the walk command; when you are using SNMPv2, supply the community string in the walk command, map the community string on the device, and enter the **snmp-server mib community-map** *community-string* **context** *context-name* command.

- CSCsr46956

  Symptom: The Nexus 7000 Series device scales up to a maximum of 4000 VLANs across the entire system. These VLANs can be configured in single VDCs or across multiple VDCs. Problems can occur with multiple modules if the total number of VLANs configured on the device across all VDCs exceeds 4000.

  Conditions: This symptom can occur in all conditions.

  Workaround: Restrict the total number of VLANs configured on the device to be fewer than 4000.

- CSCsr52252

  Symptom: After you upgrade to Release 4.0(4) from a previous release and you enter the **show eltm table** command from a module, the display may not show output for the module.

  Conditions: This symptom can occur in all conditions.

  Workaround: Run the command from the supervisor module.

- CSCsr61947

  Symptom: When you move a tunnel source interface to another VDC, the device should bring that tunnel interface down, but it is still up.

  Conditions: This situation occurs whenever you move a tunnel interface to another VDC.

  Workaround: Enter the **shutdown** command to bring the tunnel interface down manually or configure another interface as the tunnel source interface.

- CSCsr68326

  Symptom: When the device restarts the Netstack process, some IPV6 multicast protocols, such as OSPFv3, do not receive protocol packets.

  Conditions: After the device restarts the Netstack process, those IPV6 multicast protocols that do not receive protocol packets do not establish neighbors.

  Workaround: Restart the affected IPv6 multicast protocol.

- CSCsr75691

  Symptom: The device displays the CMP as operationally up, even when there is no cable connection to the CMP.

  Conditions: The output for the **show interface cmp-management** command shows the interface as up, even when there is no cable connection to the CMP.

  Workaround: No workaround.

- CSCsr82153

  Symptom: When you are saving the configuration in a nondefault VDC using the **show running-config startup-config** command and you enter the **show startup-config** command in the default VDC, the device does not display the startup-config and returns the following error:

  ```
  configuration change in progress
  ```

  Conditions: If you enter the **show startup-config** command in the default VDC when there is an ongoing **show running-config startup-config** command in a nondefault VDC.

  Workaround: Reenter the **show start-up config** command after the copy command mentioned above completes.

- CSCsr86071

  Symptom: When two devices are connected using CTS, do not perform ISSU simultaneously on both switches.

  Conditions: This symptom may occur in all conditions.

Workaround: Perform an ISSU on one device and wait for the process to complete. Then, perform an ISSU on the second device connected with CTS.

- CSCsr87423

  Symptom: No syslog message is sent when you insert either the standby supervisor or the fabric module.

  Conditions: The device does not send a syslog message when you insert either the standby supervisor or the fabric module.

  Workaround: Enter the **show module** command to check that the standby supervisor or the fabric module has been inserted.

- CSCsr90977

  Symptom: Ports may go into the error-disabled state when you apply a large ACL to a port channel with many interfaces and you reload the module with the interfaces.

  Conditions: This situation may occur when you restart a module with a large ACL applied to a port channel with many interfaces on that module. When the module restarts, the ACL policies may not reach that module and cause the related ports to remain down and move into the error-disabled state.

  Workaround: Manually bring up each port that is in the error-disabled state.

- CSCsr91565

  Symptom: After you reinitialize a module with port-channel subinterfaces that run Relay ACLs, the Relay ACL is removed from the port-channel interfaces.

  Conditions: You enabled the Relay function on the device by entering the **service dhcp** command. The module is configured with subinterfaces on a port channel with active members. After you reinitialize the module, some of the port-channel member interfaces are moved to a different VDC. The Relay ACL is removed from the port channel and port channel-subinterface on that module.

  Workaround: After you reinitialize a module with port-channel subinterfaces that run Relay ACLs, disable the Relay ACLs by entering the **no service dhcp** command and then reenable the Relay ACLs by entering the **service dhcp** command.

- CSCsr93674

  Symptom: When you enter the **show ip arp vrf** *nondefault-vrf* | **last** *num* command for a nondefault, VRF, the device does not return the shell prompt.

  Conditions: When this situation occurs, you can press Ctrl- C to return the device to its normal state.

  Workaround: Enter the **show ip arp vrf** *nondefault-vrf* | **tail lines** *num* command.

- CSCsr96589

  Symptom: When you are replaying ASCII configuration scripts in nondefault VDCs, various private-vlan configuration commands fail.

  Conditions: When you replay ASCII configuration scripts in a nondefault VDC, the generated **feature private-vlan** command does not fall in the correct place. As a result, all other private-vlan commands fail.

  Workaround: Manually move the **feature private-vlan** command to come after the other commands that enable features in your ASCII configuration script for nondefault VDCs.

- CSCsr99927

  Symptom: If you configure a minimum MTU value for path-mtu-discovery that is greater than the actual value discovered, the device does not fall back to the default value until the default ager times out in 10 minutes.

  Conditions: If you configure a minimum MTU value for path-mtu-discovery that is greater than the actual value discovered. the device should immediately fall back to the default value. However, the device waits until the ager times out (the default is 10 minutes) before it falls back to the default minimum value.

  Workaround: Enter the tunnel configuration mode, enter the **no tunnel path-mtu-discovery** command, and then enter the **tunnel path-mtu-discovery min-mtu** *mtu-value* to disable and then reenable the process, or you can wait for the ager to time out and the value will be reset.

- CSCsu01048

  Symptom: You may see high CPU utilization on the Nexus 7000 series device if the network is passing a lot of packets that require fragmentation or are hitting the TTL expiry time.

  Conditions: The device sends packets that require fragmentation or are hitting the TTL expiry time to the supervisor to forward or generate ICMP errors. Rate limiters do not take effect for this traffic. The device sends these packets to the supervisor using the copy mechanism, and so the packets are limited only by the copy rate-limiter. A high rate of such traffic can increase CPU utilization.

  Workaround: Configure the network so that the device does not receive a large number of these packets.

- CSCsu01052

  Symptom: If you configure a large number of port ACLS on a port-channel member, member port may be set to the error-disabled or suspended state.

  Conditions: When you apply a large PACL policy for the first time, some of the affected port-channel members may be put into the error-disabled or suspended state during initialization. Note that ACL policies are applied only once during the first initialization and remain persistent in the hardware. Subsequent port initializations do not trigger the device to download policies to the hardware.

  Workaround: To recover, enter the **shutdown** command and then enter the **no shutdown** command on the error-disabled or suspended ports from the Interface configuration mode.

- CSCsu01596

  Symptom: After you enable the path-mtu-discovery process, the device may fragment tunneled packets, which may lead to packet drops at the tunnel destination because of rate limiters.

  Conditions: This situation occurs when the path MTU for two or more devices in the tunnel path are configured for a lower MTU than the tunnel destination MTU.

  Workaround: Configure the tunnel interface MTU to be the lowest possible value and disable the path-mtu-discovery process.

- CSCsu01779

  Symptom: After you upgrade from the NX-OS Release 4.0.2 to the Release 4.0.3, the statistics for rate limiting may show incorrect values.

  Conditions: After you upgrade from the NX-OS Release 4.0.2 to the Release 4.0.3 and enter the **show hardware rate-limit** command, the resulting display may show incorrect values.

  Workaround: Enter the **clear hardware rate-limit** command after you upgrade the device.

- CSCsu05411

  Symptom: When there is more than one path to a prefix, the consistency checker may report an inconsistency, even though there is no inconsistency.

  Conditions: The consistency checker may report false positives for routes with ECMP.

  Workaround: You can use the consistency checker conclusively only with those routes that have a single next hop.

- CSCsu22036

  Symptom: Layer 3 multicast is not supported on port-channel subinterfaces.

  Conditions: Port-channel subinterfaces are not included in the Layer 3 multicast outgoing list.

  Workaround: No workaround.

- CSCsu45752

  Symptom: When you insert or remove the compact flash (CF) of the logflash and enter the **dir logflash** command, the Nexus 7000 Series supervisor module may reload, which results in a switchover to the standby supervisor.

  Conditions: The Nexus 7000 Series supervisor module can switch over to the standby supervisor and return the following message when you insert or remove the compact flash of the logflash and enter the **dir logflash** command:

  ```
  N7K# Raw time read from Hardware Clock: Y=2008 M=8 D=29 11:17:25 writing reset reason
  34, Service "syslogd" in vdc 1
  ```

  Workaround: To remove the CF, press the Reject Request button, for the logflash and slot0 respectively, and wait for the LED to turn off. You can then safely remove the CF. To insert the CF, you can just push it in and wait for the LED to turn on. You can also enter the **unmount logflash** command before you remove the logflash CF, and enter the **unmount slot0** command before you remove the slot 0 CF; wait for the LED to turn off.

- CSCsu87911

  Symptom: The following syslog may be seen:

```
*_2008 Sep 30 07:33:37 ch1-x7x-1b %DAEMON-2-SYSTEM_MSG: fatal:
> buffer_append_space: len 4294967295 not supported - sshd[16028]_*
```

  Conditions: This symptom occurs infrequently under all conditions.

  Workaround: This message does not affect anything on the system. No workaround.

- CSCsv35262

  Symptom: When an exporter has the source interface configured on a loopback interface in an ASCII config file, an error will occur when the ASCII configuration is applied to the running configuration.

  Conditions: This symptom may occur when an exporter has the source interface configured on a loopback interface.

  Workaround: Manually change the source interface after the configuration is applied to the running configuration.

- CSCsv35626

  Symptom: VRRP groups that are in the active state with tracking enabled can change to the backup state after the supervisor module switches over.

  Conditions: The tracking state of VRRP groups can change from up to down after the supervisor module switches over, even though the underlying interface has not changed its state. The priority of the VRRP group is lowered and can change to the backup state for that group. The priority is not affected if tracking is not configured or if the tracked interface is in the down state already.

  Workaround: Sequentially enter the **shutdown** and **no shutdown** commands for the affected VRRP group or on the tracked interface.

- CSCsv35775

  Symptom: The ACL QOS process fails on the module.

  Conditions: When there are a large number of egress RACL entries in the ACL TCAM and you attach egress NetFlow with nonatomic updates enabled, the ACL QOS may fail and the policies may not be applied.

  Workaround: No workaround.

- CSCsv40044

  Symptom: Reloading a module or the system may fail when you have configured a large number of ACLs with other features such as policy-based routing, DHCP snooping/relay, NetFlow, and so forth and you have not enabled the nonatomic update feature. This failure is due to insufficient resources.

  Conditions: This symptom may occur under all conditions.

  Workaround: Enable the nonatomic updates, save the configuration, and then reload the modules or the system.

- CSCsv40606

  Symptom: When you are interoperating with devices from vendors other than Cisco and you enter the **shutdown** and **no shutdown** commands operation on port-channel interfaces, the ports in the channel may move to the error-disabled state.

  Conditions: This symptom may occur when you are interoperating with devices from vendors other than Cisco.

  Workaround: Reenter the **shutdown** and **no shutdown** commands on the error-disabled ports.

- CSCsv47908

  Symptom: When you are configuring a VRRP group, you may see virtual MAC address addition errors. The VRRP group remains in init state.

  Conditions: The insertion of the virtual MAC address into the hardware fails when you delete a VRRP group from an interface and you configure the same VRRP group number on a different interface. This situation causes the VRRP group to fail to come up. You may also see this symptom when you disable and reenable VRRP and then reconfigure the VRRP groups.

  Workaround: Enter the **shutdown** command on the interface before you remove the VRRP group. Enter the **shutdown** command on all the interfaces on which the VRRP groups are configured before you disable the feature.

- CSCsv49677

  Symptom: When you boot up the device with autorp announce or autorp discovery configured, PIM crashes.

  Conditions: You may see this symptom when you start up the device if the startup configuration has autorp announce or autorp discovery commands.

  Workaround: Configure autorp announce and autorp discovery after you enable PIM on the specified interfaces.

- CSCsv84522

  Symptom: When the VLAN network interface or the protocol resets during an ISSU, you may see packet loss and memory or mts leaks for the FHRP (HSRP, VRRP, and GLBP) gateway MAC addresses.

  Conditions: You may see this symptom when you perform an ISSU with a VLAN network interface or an FHRP configured on the modules.

  Workaround: To recover from the packet loss, enter the shutdown and no shutdown commands successively on the VLAN network interface. There is currently no workaround for the memory or mts leak.

- CSCsw64054

  Symptom: When the device does an SNMP Set on the RMON-MIB::statistics group objects or RMON-MIB::history, the SNMP agent crashes.

  Conditions: You may see this symptom under all conditions.

  Workaround: Do not perform a SNMP SET on the following groups of MIB objects:

  - iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1).rmon(16).statistics(1)

– iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1).rmon(16).history(2).

Resolved Caveats—Cisco NX-OS Release 4.0(4)

All the caveats listed in this section are resolved in Cisco NX-OS Release 4.0(4).

- CSCsr06310

  Symptom: The output from the **show scheduler logfile** command contains the following warning message:

  ```
  VSH may not load properly (not enough memory
  ```

  Conditions: When the scheduler job configuration is executed, this message appears.

  Workaround: No workaround.

- CSCsr25585

  Symptom: The first IP fragment of a large UDP packet is dropped when the packet is switched through the device. All other packet fragments are switched without dropping.

  Conditions: When you enable the maximum UDP length check, fragmented UDP packets are switched through the device. You can verify this action by entering the **show hardware forward ip verify** command. This IDS check is disabled by default in Release 4.0(3).

  Note An upgrade from a previous release does not disable this IDS check, so an explicit configuration is required.

  Workaround: After you upgrade from NX-OS Release 4.0(2) to Release 4.0(3), you should explicitly disable the IDS check of the maximum UDP packet length by entering the **no platform ip verify length maximum udp** command for IPv4 and the **no platform ipv6 verify length maximum udp** command for IPv6.

- CSCsr30896

  Symptom: The first attempt to log into the CMP as a network-admin user using SSH fails.

  Conditions: When a user with a network-admin ID first attempts to log into the CMP using SSH, the login fails. However, a user who uses Telnet to access the CMP with the same network-admin ID can successfully log in to the CMP.

  Workaround: No workaround.

- CSCsr35499

  Symptom: You may see the following message in the syslog from the "urib" process:

  ```
  %URIB-5-GENERAL_FAILURE: urib [3607] Failed

  This message may be rate-limited, resulting in the following additional message:
  %URIB-4-SYSLOG_LOG_WARNING: URIB-5-GENERAL_FAILURE: message repeated 1 times in last
  221 sec
  ```

  Conditions: When there is a change to the IPv4 unicast RIB, you may see these messages in the syslog.

  Workaround: This message can be ignored.

- CSCsr39258

  Symptom: When you have log-neighbor-changes configured, no log messages are shown when the BGP peer goes up or down.

  Conditions: This symptom can occur in all conditions.

  Workaround: If you set the BGP logging level to Information by entering the **logging level bgp 6** command and also set the monitor/console logging level to Information, the log messages are shown when a BGP peer goes up or down.

- CSCsr52919

  Symptom: When you enter the **attach cmp** command from the console port, the system may not issue the login/switch prompt to the CMP.

  Conditions: The session appears to hang.

  Workaround: Connect to the CMP using SSH or Telnet. Alternatively, you can reload the CMP.

- CSCsr55726

  Symptom: When you enter the **show cdp neighbor detail** command, the display for the neighboring device does not show a value for the address TLV information for the CDP packets received from the Nexus 7000 Series Ethernet interface.

  Conditions: This symptom may occur when you add a new IP address or modify an existing IP address in the interface.

  Workaround: If you change the interface IP address configuration, explicitly reenable CDP on this interface by sequentially entering the **no cdp enable** and **cdp enable** commands.

- CSCsr65262

  Symptom: The system sometimes considers the secondary IP address on an interface to be the primary IP address.

  Conditions: This symptom occurs only if the primary and the secondary addresses are configured on interfaces and the Netstack process restarts.

  Workaround: No workaround.

- CSCsr69621

  Symptom: If you enter a global command, such as the **system default** or **system jumbomtu** command, while a module is powering up or initializing, the device may not recognize a module after it is reseated.

  Conditions: If you are initializing a module and you enter either the **system default** or the **system jumbomtu** command, the device may not recognize the module although it is up and functioning.

  Workaround: Do not enter either the **system default** or the s**ystem jumbomtu** command when you are powering up or initializing a module. Enter these commands only after the all the modules have reached a stable state (online or powered down).

- CSCsr71827

  Symptom: If you change the MTU value while traffic is running, you may see traffic loss on some interfaces.

  Conditions: When the device is running 100 percent traffic with random packet sizes and you change the MTU value while the traffic is running, you may see some traffic loss.

  Workaround: Do not change the MTU value while traffic is running. If the device is currently in this condition, limit traffic to 75% line rate. If the device is currently in this condition, enter the **shutdown** command and then enter the **no shutdown** command on that interface.

- CSCsr81353

  Symptom: If you reinitialize the two modules that run the Relay function, each of which contains interfaces for the same port channel, in quick succession, one of the modules will no long have the Relay function enabled.

  Conditions: You enabled the Relay functionality on the device by entering the **service dhcp** command. A port channel has member interfaces from two different modules. If you reinitialize both modules in quick succession, one of the modules will not have the Relay ACL programmed. when the modules come back up. The Relay functionality will be affected on the member interfaces in the port channel on the module on which the ACL programming is missing.

  Workaround: After you reinitialize the modules, disable the Relay ACLs by entering the **no service dhcp** command and then reenable the Relay ACLs by entering the **service dhcp** command

- CSCsr82191

  Symptom: The system may not respond to Ping6 requests.

  Conditions: The punt flag in the hardware may not be set for a IPv6 locally joined SSM group. This situation results in the data packets for the locally joined group not being punted to the supervisor module.

  Workaround: No workaround.

- CSCsr86965

  Symptom: A user with multiple roles in the AAA server cannot log into the CMP.

  Conditions: Users who have multiple roles, other than admin and that are assigned through the AAA server, cannot log into the CMP.

  Workaround: To log into the CMP, create a user in the AAA server that has only the network-admin role and log in as that user.

- CSCsr93039

  Symptom: Packets that are destined to the HSRP virtual IP address that are received on a standby HSRP router may be dropped.

  Conditions: This symptom can occur in all conditions.

  Workaround: Ping the virtual IP address from the standby HSRP device, which initiates an ARP request. Subsequent pings from other devices will work until the ARP table entry on the standby HSRP device expires.

- CSCsu01914

  Symptom: After you reload a module or shut down and bring up the ports, HSRP is stuck in Initial (Interface Reload Delay)(0 remaining).

  Conditions: This symptom appears when you have entered the **hsrp delay reload** command before you reload the module or enter the **shut/no shut** command.

  Workaround: Do not enter the **hsrp delay reload** command.

- CSCsu27743

  Symptom: If the OSPF area ID value is less than 256, the device ignores the leading zeroes in the dotted-decimal notation.

  Conditions: This symptom may occur in all conditions.

  Workaround: No workaround.

- CSCsu29293

  Symptom: If you have an IOFPGA with a version earlier than 3.19, the smart card chip, although it is present and required for CTS, is not used.

  Conditions: If you have an IOFPGA with a version earlier than 3.19, a kernel crash (KGDB) may occur while the smart card module attempts to access the smart card hardware. Access to the smart card is disabled if the IOFPGA version is earlier than 3.

  Workaround: Upgrade the current version of IOFPGA to 3.19 or higher on each supervisor module.

- CSCsu29513

  Symptom: When a Layer 2 aggregation switch is connected to a Layer 3 access router, the Layer 3 routers do not discover the Layer 2 switch as a CDP neighbor.

  Conditions: The CDP neighbor from a Nexus 7000 Series Layer 2 interface is not recognized in the Layer 3 interface side of a connected Catalyst 6500 series switch or other Cisco device.

  Workaround: On the Layer 2 side, add the first allowed VLAN from the Layer 2 port channel allowed list of VLANs. On the Layer 3 side, create a dot1Q-encapsulated subinterface with the VLAN that is the first allowed VLAN from the Layer 2 port channel. You do not need to configure an IP address on this subinterface.

- CSCsu29522

  Symptom: IGMP snooping creates a state for reserved multicast groups. The **show ip igmp snooping groups** command displays the states created by IGMP snooping for the multicast groups. However, the state is being created for packets, such as OSPF, with the multicast address 224.0.0.5.

  Conditions: This symptom may occur in all conditions.

  Workaround: Disable IGMP snooping on the VLAN on which IGMP snooping creates a state for reserved multicast groups.

- CSCsu26944

  Symptom: The system does not always properly clear virtual IP addresses used in FHRP groups. This symptom affects HSRP, GLBP, and VRRP protocols.

Conditions: This symptom can occur when all the virtual IP addresses configured for FHRP groups are removed and the Netstack service either restarts or a switchover occurs.

Workaround: This problem has been fixed in Release 4.0(4). To clean any stale entries after an upgrade to Release 4.0(4), follow these steps:

1. Copy the current FHRP configuration.
2. Disable the FHRP features by entering the **no feature hsrp** command.
3. Enter the internal **clear ip vip uuid** *uuid* command using the value for *uuid* as 406 for HSRP, 441 for GLBP, and 68 for VRRP.
4. Enable and reapply the FHRP configuration.

- CSCsu38694

  Symptom: If you repeatedly remove and reinsert the fan module, the system does not shut down when you remove one system fan after 180 seconds and when you remove both fans after 120 seconds, which is the default policy.

  Conditions: All conditions.

  Workaround: No workaround.

- CSCsu41395

  Symptom: When you configure ISIS on a 10-Gigabit Ethernet link in a shared-rate mode on VLAN interfaces where Layer 2 port channels connect to two Nexus 7000 Series devices, the devices do not bring up the correct adjacency.

  Conditions: In this configuration, the neighboring router shows the ISIS neighbor in an Init state and the first router does not see the adjacency.

  Workaround: Configure the 10-Gigabit Ethernet links for a dedicated-rate mode, or configure ISIS on the VLAN interface between the two devices and connect to the router through a Layer 2 port channel using 10-Gigabit Ethernet links.

- CSCsu42541

  Symptom: The SNMP statistics display shows incorrect results for the ifHCOutOctets value.

  Conditions: The interface output octets reported to SNMP was an incorrect value.

  Workaround: No workaround.

- CSCsu42753

  Symptom: If you have a large HSRP and/or VRRP configuration on your Nexus 7000 Series device, in Release 4.0(3) or higher, ISSU may fail.

  Conditions: When this symptom occurs, the following error message displays and the installation fails when you enter the **install all** command on Release 4.0(3) to a later release:

  ```
  2008 Sep 5 14:28:28 nexus1 %$ VDC-1 %$ %SYSMGR-2-GSYNC_SNAPSHOT_SRVFAILED: Service
  "hsrp_engine" on active supervisor failed to store its snapshot (error-id 0x40480008).
  2008 Sep 5 14:28:28 nexus1 %$ VDC-1 %$ %SYSMGR-2-GSYNC_FAILED_ACTIVE: global sync failed
  for UUID 0x196 2008 Sep 5 14:28:28 nexus1 %$ VDC-1 %$ "hsrp_engine" 2008 Sep 5 14:28:28
  nexus1 %$ VDC-1 %$ %SYSMGR-2-STANDBY_BOOT_FAILED: Standby supervisor failed to boot up.
  2008 Sep 5 14:28:31 nexus1 %$ VDC-1 %$ %PLATFORM-2-MOD_REMOVE: Module 6 removed (Serial
  number JAB1223010F).


  "hsrp_engine" 2008 Sep 5 14:28:28 nexus1 %$ VDC-1 %$ %SYSMGR-2-STANDBY_BOOT_FAILED:
  Standby supervisor failed to boot up. 2008 Sep 5 14:28:31 nexus1 %$ VDC-1 %$
  %PLATFORM-2-MOD_REMOVE: Module 6 removed (Serial number JAB1223010F).
  ```

  Workaround: If this symptom occurs, you should perform a failover of all the active HSRP and VRRP groups to the other router, disable the feature, perform the ISSU, reenable the feature, and reapply the configuration to the feature. The steps for this procedure with HSRP follow; you perform similar steps for VRRP:

  1. Save the existing HSRP configuration onto the bootflashes on both the supervisors by entering the **show running-config hsrp > bootflash:hsrp.config copy bootflash:hsrp.config bootflash://sup-standby/hsrp.config** command.

  2. Fail over all the HSRP groups to the other router. (This failover occurs automatically if the feature is disabled, and the peer router has the necessary HSRP configuration.)

  3. Disable the HSRP feature by entering the **no feature hsrp** command.

  4. Perform an ISSU.

  5. Once the ISSU is fully complete, enter the internal **clear ip vip uuid** *uuid* command, with the *uuid* value as 406 for HSRP (68 for VRRP).

  6. Reenable HSRP by entering the **feature hsrp** command.

  6. Apply the HSRP configuration by entering the **copy bootflash:hsrp.config running-config** command.

  This fix is in Release 4.0(3E1) and 4.0(4).

- CSCsu46808

  Symptom: In some configuration sequences, VRRP groups may become active-active between two routers.

  Conditions: This symptom may occur when there are multiple VRRP groups active in an interface, and one of the groups is shut down or removed. When you enter the **show ip interface** *interface* command, the resulting display shows that the VRRP multicast group (224.0.0.18) is not registered so the incoming VRRP multicast packets are dropped. All the VRRP groups in that interface are affected.

  Workaround: Sequentially enter the **shut** and **no shutdown** commands on any one of the VRRP groups in that interface.

- CSCsu50191

  Symptom: When you are interoperating with devices from vendors other than Cisco, the LACP port channels do not form if the other device is not standards-compliant.

  Conditions: This symptom may occur when you are interoperating with devices from other vendors.

  Workaround: Configure the port channels in the on mode, rather than using LACP.

- CSCsu51057

  Symptom: Link-local packets, such as an IGMP query, are ignored.

  Conditions: When you enable snooping on a VLAN, and a link-local packet such as an IGMP query is received on this VLAN, the device ignores the link-local packet.

  Workaround: Enter the **ip igmp snooping querier** command in the vlan configuration mode.

- CSCsu55633

  Symptom: The system takes a long time to completely display the **show tech** command.

  Conditions: This symptom can occur in all conditions.

  Workaround: No workaround.

- CSCsu58049

  Symptom: The memory usage is high for the unicast FIB distribution module (UFDM).

  Conditions: A port channel that flaps multiple times can result in an increase in the UFDM memory usage. The UFDM process may leak memory when a port channel flaps rapidly.

  Workaround: No workaround.

- CSCsu62384

  Symptom: When an outbound ACL is applied on an interface, you may see dropped RP-sourced control traffic (such as HSRP or PIM).

  Conditions: This symptom occurs when you configure an outbound ACL on an interface.

  Workaround: Modify the outbound ACLs to explicitly permit protocols (for example HSRP or PIM) that will be running on those interfaces.

- CSCsu67502

  Symptom: If you enter the **show tech** command when you have configured a large PACL on the management interface (management0), the Netstack process might take an exception.

  Conditions: This symptom occurs when you have a large PACL configured on the management interface.

  Workaround: Remove the PACL from the management interface before entering the **show tech** command.

- CSCsu68547

  Symptom: When you configure the MTU value on a range of subinterfaces, the port manager application may access invalid memory, which could result in a system failure or memory corruption.

  Conditions: This symptom may occur while you are configuring the MTU on a range of subinterfaces.

  Workaround: No workaround.

- CSCsu70937

  Symptom: The results of the statistics on the port channels may be incorrect.

  Conditions: This symptom may occur in all conditions.

  Workaround: No workaround.

- CSCsu79429

  Symptom: After reloading the 10-Gigabit Ethernet module on the Nexus 7000 series chassis, some of the interfaces are unavailable.

  Conditions: This symptom may occur when you have configured this module with more than 1,500 VLANs and more than 20 ports administratively up.

  Workaround: Reduce the number of VLANs.

- CSCsu89455

  Symptom: The sysmgr may crash in the ACL QoS process.

  Conditions: This symptom may occur when there are rapid link flaps on an interface.

  Workaround: Shut down the interface.

- CSCsu95333

  Symptom: The alarms and warnings, and the clearing of these alarms and warnings, for the Power Current Voltage and Temperature may not appear correctly on the syslog.

  Conditions: This symptom may occur under all conditions.

  Workaround: No workaround.

- CSCsv00128

  Symptom: When you enter the **show lacp port-channel interface port-channel** *channel-number* command, the command does not succeed. Entering this command causes the LACP process to restart gracefully.

  Conditions: This symptom may occur in all conditions.

  Workaround: To display the correct output, enter the **show interface port-channel** *channel-number* command.

- CSCsv02757

  Symptom: When HSRP is active on the Nexus Series 7000 device, the MAC table in the hardware may not contain a gateway MAC entry for the HSRP MAC address. Instead, the hardware MAC table may contain a dynamically learned MAC address, which causes the Layer 3 functionality through HSRP MAC to break.

  Conditions: This symptom can occur under the following conditions:

  - HSRP was not configured on the Nexus Series 7000 device, and Layer 2 on the device learns the HSRP MAC address.
  - Now, HSRP is configured on the device, which leaves a stale entry in the software dynamic MAC table in addition to the static gateway MAC table.
  - If the port channel on the remote end flaps, all dynamic MAC addresses (rather than only the MAC addresses related to that port channel) are deleted from the module.

  These conditions lead to the purge of the gateway MAC entries from the module, resulting in an inconsistency between the software and hardware.

  Workaround: Sequentially enter the **shutdown** and **no shutdown** commands on the VLAN interface on which HSRP is configured.

- CSCsv05722

  Symptom: On rare occasions, when your perform an ISSU upgrade from Release 4.0(2) or an older image to an image lower then Release 4.0(4), an MRIB core may occur.

  Conditions: The symptom may occur when you perform an ISSU upgrade. The problem can occur only if the upgrade image is earlier than Release 4.0(4).

  Workaround: No workaround.

- CSCsv07357

  Symptom: You may see control plane instabilities, such as a routing adjacencies loss, if the Nexus 7000 Series device is flooded continuously with IP traffic that has a time-to-live value of 1.

  Conditions: This symptom may occur under all conditions. (Modifying the copy rate-limited value does not mitigate the issue.)

  Workaround: Resolve the reason for the heavy traffic that has a time to live value of 1 by eliminating routing loops.

- CSCsv08459

  Symptom: The start time is incorrect in exported flows.

  Conditions: After a period of time when using the flow collector, the user may see that the first switched time stamp is smaller than the last switched time stamp.

  Workaround: No workaround.

- CSCsv10995

  Symptom: When the system is rebooted with VRRP/HSRP/GLBP configuration, these protocols may be stuck in the initial state.

Conditions: When a large number of VRRP, HSRP and/or GLBP groups are configured, and the system is rebooted with this configuration, these protocols may be stuck in the initial state.

Workaround: Remove the affected group configuration on the interface and reapply it.

- CSCsv11298

  Symptom: The time and frequency of the port flaps that cause an error disable on that port is changed. The port goes into the error-disabled state only after the port flaps (goes up and down) 30 times in 420 seconds.

  Conditions: Under all conditions, the port goes into the error-disabled state only after the port flaps 30 times in 420 seconds. Before Release 4.0(4), the port goes into the error-disabled state after it flaps 5 times in 10 seconds.

  Workaround: No workaround.

- CSCsv12524

  Symptom: When you reload the chassis and have configured dot1x on the interface, the port security address table may be empty.

  Conditions: This symptom may occur under all conditions.

  Workaround: No workaround.

- CSCsv21972

  Symptom: Between Release 4.0(2) and 4.0(3), we made a change to the R2D2 flow control behavior such that the backpressure is not correctly asserted back to naxos in the face on ingress congestion (that is, the forwarding engine may get oversubscribed). This results in unexpected drops in R2D2 and is contrary to the engineering system specification and public documentation of how the ingress buffering scheme is implemented.

  Conditions: This symptom may occur in all conditions.

  Workaround: No workaround.

- CSCsv23880

  Symptom: The standby supervisor module may not come online,

  Conditions: If the standby supervisor module reloads while it is coming online, it may never come online.

  Workaround: You must reload the chassis to ensure that the standby supervisor module comes up.

- CSCsv28803

  Symptom: When you perform an ISSU in Release 4.0(4), the STP topology changes. After the system performs an HA switchover, the information about the STP peer type is lost.

  Conditions: This symptom occurs only in topologies with devices that are running the legacy IEEE 802.1d-1998 STP and Rapid PVST+.

  The Nexus 7000 Series ports that are connected to devices running the 802.1d-1998 legacy STP keep track of the peer type as an STP peer and send out 802.1d-1998 configuration BPDUs. After a switchover, the system loses the STP peer information on the new active supervisor module and

sends out RSTP BPDUs even on those STP peer ports. Those PVST+ switches that are running legacy STP (802.1d-1998) do not understand RSTP BPDUs, and so the peer device drops the RSTP BPDUs and times out the received information. This triggers the STP topology change.

However, after the Nexus7000 Series device port receives an STP BPDU from peer once again, the system sets the STP peer and sends 802.1d-1998 configuration BPDUs again.

Workaround: No workaround.

- CSCsv28982

  Symptom: When you enable ACLs and NetFlow on the same interface, the collection and clearing of the ACL statistics does not work.

  Conditions: When you enable ACLs and NetFlow on the same interface, this symptom may occur.

  Workaround: No workaround.

- CSCsv33676

  Symptom: When you run Rapid PVST+ on the device, the Topology Change (TC) may not be propagated on the port that became part of the STP active topology. There is no BPDU with TC bit sent on the port that has changed its port state to forwarding. The MAC addresses are not flushed or deleted in some parts of the STP topology and the MAC address will eventually age out and be relearned.

  Conditions: This symptom can occur in all conditions.

  Workaround: No workaround.

Resolved Caveats—Cisco NX-OS Release 4.0(3)

All the caveats listed in this section are resolved in Cisco NX-OS Release 4.0(3).

- CSCsd98511

  Symptom: There is no IPv6 support for VLAN interfaces.

  Conditions: If IPv6 addresses are configured on VLAN interfaces, they do not function properly.

  Workaround: No workaround.

- CSCsl45405

  Symptom: Control Plane Policing (CoPP) policy enforcement fails with nondefault VDCs.

  Conditions: CoPP policies that refer to ACLs for match rules are not enforced if there is a module in the device that does not have an interface allocated to the default VDC.

  Note CoPP policies that do not refer to ACLs are correctly enforced even if all the interfaces for a module are allocated to nondefault VDCs.

  Workaround: Assign at least one interface from each module to the default VDC.

- CSCsm37481

  Symptom: You see the following message on the standby supervisor module syslog:

  ```
  BOOTUP_TEST-STANDBY-2-EOBC_FAIL: Module <mod_num> has failed test EOBCLoopbackTest on
  EOBC due to error Loopback packet receive timeout
  ```

  Conditions: This message displays intermittently on the standby supervisor module at bootup time, usually during switchover or during an in-service software upgrade (ISSU) when the standby supervisor module comes up using a new system image.

  Workaround: No workaround.

- CSCsm56407

  Symptom: Memory leaks occur in the L2FM process if you press **Ctrl-C** while the **show mac address-table** command is in progress.

  Conditions: An ungraceful termination of the **show mac address-table** command results in a memory leak that causes the L2FM process to crash. The problem occurs only when you press **Ctrl-C** before any MAC addresses are displayed.

  Workaround: Use one of the following workarounds:

  - Use the **quit** command to gracefully terminate the **show mac address-table** command. The **quit** command terminates the session.
  - Wait for at least one MAC address to display before pressing **Ctrl-C** to terminate the command.

- CSCsm73179

  Symptom: Certain SNMP object identifiers (IODs) provide the wrong information.

  Conditions: The SNMP get operation fails to retrieve the following OIDs correctly:

  - cpsIfViolationAction (Always shows shutdown.)
  - cpsIfViolationCount (Count is always zero.)
  - cpsIfSecureLastMacAddress (MAC display is not correct.)
  - cpsIfSecureLastMacAddrVlanId (VLAN ID corresponding to last secured MAC is wrong.)
  - cpsIfMultiVlanMaxSecureMacAddr (This is zero.)

  Workaround: No workaround.

- CSCsm81244

  Symptom: Hard coding of prestandard MST is required.

  Conditions: Cisco Nexus 7000 Series device detects prestandard MST correctly.

  Workaround: Enter the **clear spanning-tree detected-protocols interface** command on the peer device.

- CSCsm97272

  Symptom: The traffic rate is not updated in the tunnel statistics.

  Conditions: In the **show interface** command output, the 5-minute input rate for a tunnel is always zero, even with traffic present.

```
switch# show interface tunnel 10
Tunnel10 is up
    Internet address is 10.5.5.3/24
    MTU 1476 bytes
    Transport protocol is in VRF "default"
    Tunnel protocol/transport GRE/IP
    Tunnel source 10.1.1.1 (Ethernet1/4), destination 10.1.1.1
    Tx 0 packets output, 5 minute output rate 0 packets/sec
    Rx 398532117 packets input, 5 minute input rate 0 packets/sec
```

Workaround: You can manually estimate the 5-minute input rate by determining the number of packets received over an interval of 5 minutes and dividing it by 300 seconds.

- CSCsm99880

  Symptom: Rollback functions are not available on nondefault VDCs for undefined users.

  Conditions: If you switch from the default VDC to a nondefault VDC, and your user account does not exist in the nondefault VDC or you do not have the network-admin role, then the rollback operation fails in that nondefault VDC.

  Workaround: Add yourself as a user with the vdc-admin role in the nondefault VDC and try the rollback operation again.

- CSCso01961

  Symptom: DHCP relay is not supported on Layer 3 port-channel interfaces.

  Conditions: DHCP Relay does not work on Layer 3 port-channel interfaces.

  Workaround: No workaround.

- CSCso03220

  Symptom: The device may unnecessarily send multicast data packets to the control plane.

  Conditions: When the device leaves an IGMP multicast group, the multicast data packets may continue to be sent to the control plane.

  Workaround: No workaround.

- CSCso07804

  Symptom: RewriteEngine and PortLoopback on-demand GOLD tests may not work on all modules.

  Conditions: The results for the on-demand GOLD diagnostic Rewrite Engine and PortLoopback test displayed in the **show diagnostic result module** command output may indicate untested (U) for some modules under the following conditions:

  - There are more than 4 active modules.
  - The PortLoopback test is enabled to run health monitoring.

    If both the Rewrite Engine test and the PortLoopback test are enabled to run health monitoring on all the modules, the results may indicate untested (U) with three or more active modules.

  Workaround: Health monitoring can be run individually on each module.

  There is no workaround for health monitoring if there are more than four active modules.

There is no workaround if both the Rewrite Engine test and the PortLoopback test are enabled for health monitoring, and there are more than two active modules.

- CSCso09976

  Symptom: LACP system priority cannot be changed when LACP port channels are up.

  Conditions: If you attempt to change the LACP system priority when active/passive mode port channels are in the Up state, some ports may become suspended.

  Workaround: To change the LACP system priority, use the following steps:

  1. Shut down the active/passive mode port channels using the **shutdown** command.
  2. Change the system priority.
  3. Bring the active/passive mode port channels back up using the **no shutdown** command.

- CSCso16917

  Symptom: Command-line help output from the **copy license ?** command shows unsupported formats.

  Conditions: The **copy license ?** command online help output displays URI formats for remote targets that are not supported, including **scp:**, **sftp:**, and **ftp:**.

  Workaround: Use the following command sequence instead:

  1. **copy licenses bootflash:filename.tar**
  2. **copy bootflash:filename.tar scp://**...

- CSCso18248

  Symptom: Some OSPFv2 RFC4750 SNMP traps are not supported.

  Conditions: The NX-OS software supports the traps defined in the OSPF Version 2 MIB (RFC4750) except for the following:

  - ospfMaxAgeLsa
  - ospfNssaTranslatorStatusChange
  - ospfIfTxRetransmit
  - ospfVirtIfTxRetransmit
  - LsdbOverFlow
  - ApproachingLsdbOverFlow

  Workaround: No workaround.

- CSCso19495

  Symptom: The **show scheduler config** command output includes junk characters.

  Conditions: If you modify the configuration of an existing scheduler job, some junk characters can be printed in the scheduler-job output of the **show scheduler config** command.

  Workaround: Instead of modifying an existing scheduler job, delete and recreate it.

- CSCso21270

  Symptom: PBR is not supported for IPv4 multicast.

  Conditions: If the PBR policy match criteria matches the IPv4 multicast traffic, packets will not be replicated.

  Workaround: No workaround.

- CSCso22132

  Symptom: The **show running-config ntp** command is not supported.

  Conditions: Entering the **show running-config ntp** command results in an invalid command response.

  Workaround: To view the NTP configuration, use the **show running-config | include ntp** command.

- CSCso22558

  Symptom: A system message is sometimes displayed when tunnel interfaces or modules are brought up.

  Conditions: Occasionally, the NX-OS software displays the following message when tunnel interfaces are being brought up or when a module comes online:

  ```
  2008 Mar 14 19:32:05 qadc3-ind06 %$ VDC-1 %$ %KERN-2-SYSTEM_MSG: mts_sys_drop():
  invalid input q_entry 0xc8ad2920 for PID 6435 (tm). - kernel"
  ```

  Note This message is not an indication of a system malfunction.

  Workaround: No workaround.

- CSCso24616

  Symptom: When IP IGMP snooping querier functionality is enabled, queries originating on the device will have a zero source MAC address.

  Conditions: If you have enabled the IP IGMP snooping querier functionality on the device, the queries that originate on the device have a source MAC address of all zeros. If hosts ignore these packets, the multicast state is not built correctly.

  Workaround: No workaround.

- CSCso27589

  Symptom: HSRP VRMAC is still present in the MAC table after use-bia is configured.

  Conditions: After the group is configured, if you enable use-bia, then the earlier VMAC (default VMAC) remains in the Layer 2 MAC table.

  The protocol behavior is not affected, but an extra slot in the Layer 2 MAC table is used.

  Workaround: To remove the extra VMAC, follow these steps:

  1. Unconfigure use bia.
  2. Remove the group.

3. Reconfigure the same group.
4. Remove the group.
5. Configure use-bia.
6. Reconfigure the group.

- CSCso30022

  Symptom: The static binding entry is not displayed in the **show ipsg verify source** command output.

  Conditions:

  1. You create a static binding entry on a VLAN and interface using the **ip source binding vlan interface** command.
  2. You enable IP Source Guard on the interface using the **ip verify source dhcp-snooping-vlan** command.

     A dynamic binding entry is automatically created for the same VLAN on a different interface.
  3. You use the **show ipsg verify source** command.
     Although the static binding entry is created, it does not appear in the output of the **show** command.

  Note The impact is limited to the **show** command and does not affect system functions.

  Workaround: No workaround.

- CSCso30164

  Symptom: Secured addresses can be configured beyond the limit on a native VLAN.

  Conditions: On a trunk port, the maximum limit is not checked when sending addresses to a native VLAN. Static secure addresses can be configured beyond the limit.

  Workaround: No workaround.

- CSCso30349

  Symptom: The alternate Cisco Secure Access Control Server (ACS) configured is not contacted when the first Cisco Secure ACS is down.

  Conditions: The alternate Cisco Secure ACS configured is not contacted after a reauthorization timeout when the first Cisco Secure ACS server is down.

  Workaround: No workaround.

- CSCso32516

  Symptom: The DHCP feature may not work correctly after rolling back to a checkpoint.

  Conditions: Certain DHCP configurations might not work correctly when used with rollback checkpoints. For example, the DHCP configuration in the running configuration is not removed even though DHCP is disabled.

  Workaround: No workaround.

- CSCso32688

  Symptom: The MTS_SAP_PVLAN component timed out while configuring a private VLAN.

  Conditions: While configuring a private VLAN, the interface times out.

  Workaround: Enter the **shutdown/no shutdown** command sequence on the interface.

- CSCso32847

  Symptom: After a supervisor module switchover, a DOT1x_MAC type is changed to Dynamic.

  Conditions: Port security secures the MAC authenticated by 802.1X. After a switchover, the MAC type changes from DOT1x_MAC to Dynamic. This does not effect the system functions.

  Workaround: No workaround.

- CSCso32975

  Symptom: The static binding entries are lost when you disable DHCP on the device.

  Conditions: DHCP has both dynamic and static binding entries. When you disable DHCP snooping globally using the **no ip dhcp snooping** command, the static binding entries are removed. Only the dynamic entries should be removed.

  Workaround: Do not disable global DHCP snooping using the **no ip dhcp snooping** command.

- CSCso33872

  Symptom: When you suspend and reactivate a VLAN, the ACL functions related to t DHCP snooping, DAI, and ARP are lost.

  Conditions: You have configured one or more ACLs on a VLAN for DHCP snooping, DAI, or ARP. If you suspend and then reactivate the VLAN, the ACLs are lost and the features do not work.

  Workaround: Disable and reenable the feature to reprogram its ACLs. The security features are available again when the reprogramming is completed.

- CSCso35940

  Symptom: The Multicast FIB Distribution Manager (MDFM) has missing IPv6 (S,G) route information in the M6RIB after a supervisor module switchover.

  Conditions: After a supervisor module switchover, the IPv6 (S,G) route exists in the hardware and the MFIB but is missing in the M6RIB.

  Workaround: No workaround.

- CSCso36689

  Symptom: A Call Home failure event is not generated when the standby CMP is reloaded.

  Conditions: A Call Home failure event is not generated when the CMP on the standby supervisor module is reloaded. An observed Call Home failure event is generated when reloading CMP on the active supervisor module.

  Workaround: No workaround.

- CSCso37573

  Symptom: CFS-related Call Home syslog messages display.

  Conditions: If you receive any Cisco Fabric Services (CFS)-related syslog messages, you can ignore them. CFS is not supported.

  Workaround: No workaround.

- CSCso37615

  Symptom: ACLs are not programmed on the module when configured in a command script.

  Conditions: When you configure DHCP in a command script, the ACLs are not programmed on the module and DHCP and DAI features do not function correctly.

  Workaround: Wait about 5 seconds after entering the **feature dhcp** command before entering the other commands.

- CSCso37810

  Symptom: IP Source Guard does not function correctly on an interface that has DCHP snooping enabled for a private VLAN.

  Conditions:

  1. You enable DHCP snooping on a private VLAN.
     A binding entry is created on the primary VLAN.
  2. Then you enable IP Source Guard on the interface to which the host is connected (the same interface where the binding entry was created).

     The binding entry is not pushed down to the FIB so the IP Source Guard feature does not function correctly on the interface.
  3. You use the **show ip verify source** command to display the interface.
     The command output shows inactive-no-snoop-vlan status.

  Workaround: No workaround.

- CSCso37955

  Symptom: A rollback for the VLAN resource in a VDC resource template does not occur.

  Conditions: A rollback of the VLAN resource in a VDC resource template does not occur. All the other resources of the VDC resource templates are rolled back correctly.

  Workaround: Perform a manual rollback.

- CSCso38505

  Symptom: NetFlow configurations may be missing or incorrect after a multiple rollbacks.

  Conditions: While reverting NetFlow commands to previous configuration with multiple rollbacks, the NetFlow configuration may not be restored completely.

  Workaround: No workaround.

- CSCso39319

  Symptom: IPv6 neighbor directory does not find and remove neighbors that fail.

  Conditions: When a neighbor fails, the IPv6 neighbor directory does not probe to find that neighbor quickly and remove it.

  Workaround: No workaround.

- CSCso40182

  Symptom: Packets are dropped if there is an inactive time-range ACL with a single ACL rule.

  Conditions: When the time-range ACL is not active, an implicit deny is added to the ACL instead of an implicit permit. This action causes packets to be dropped and occurs only for ACLs with one ACL rule.

  Workaround: Configure more than one ACL rule in the ACL, such as an explicit permit.

- CSCso40446

  Symptom: NTP commands may hang up to 9 minutes.

  Conditions: NTP configuration and **show** commands are not acknowledged and the following message displays:

  ```
  Process did not respond within the expected timeframe
  ```

  Note Do not enter any other NTP configuration or **show** commands after the above message is displayed. Every command that you enter adds an additional delay of up to 9 minutes.

  Workaround: Use the **ntp sync-retry** command. This command is acknowledged and processed.

- CSCso41633

  Symptom: If a module goes offline, in some conditions, the Control Plane Policing (CoPP) process could restart.

  Conditions: If a module goes offline and if there is a large control plane policing, QoS, or ACL configuration, the CoPP process might restart.

  Workaround: No workaround.

- CSCso42365

  Symptom: The L2FM process crashes when the same MAC address is statically configured on a secondary and a primary VLAN after the association

  Conditions: The following actions result in the L2FM process crash:

  1. Configure a primary VLAN.
  2. Associate the secondary VLAN to the primary VLAN.
  3. Configure the same static MAC address on the secondary VLAN and primary VLAN.

  The L2FM process crashes at this point.

Workaround: Do not configure static MACs on a secondary VLAN after a primary VLAN association.

- CSCso42877

  Symptom: DHCP binding entries that are created for the clients disappear.

  Conditions: When you use a Catalyst 6500 series switch as a DHCP server and many of the DHCP clients request IP addresses, the DHCP binding entries for the clients disappear.

  Workaround: Use a Linux server as the DHCP server.

- CSCso43676

  Symptom: A false instance of OSPF may be started after a switchover.

  Conditions: If the device is very busy during a supervisor module switchover, it may create a false OSPF instance, which results in OSPF adjacencies not coming up.

  Workaround: Remove and reconfigure OSPF by copying the startup configuration to the running configuration.

- CSCso45324

  Symptom: Ingress VLAN NetFlow is not supported for trunk ports or port channels.

  Conditions: If you configure an interface in switchport trunk mode, allow all VLANs, and apply NetFlow on the VLAN, no configurations are pushed to the interface.

  Workaround: No workaround.

- CSCso46490

  Symptom: SNMP MIBs or traps for the Layer 3 protocols are supported only for the first instance default VRF.

  Conditions: There is no SNMP MIB support for nondefault VRFs in Layer 3 protocols.

  Workaround: No workaround.

- CSCso46614

  Symptom: Dynamic ARP inspection (DAI) fails following a supervisor module switchover.

  Conditions: After you have configured the DAI feature and a supervisor module switchover occurs, DAI fails.

  Workaround: No workaround.

- CSCso47093

  Symptom: CoPP of IP option exception packets does not work.

  Conditions: All IP option exception packets (sent at a high rate) are reaching the supervisor module even though a policer is configured with low values. When configured as conform drop, the I/O module and supervisor module statistics shows violated counts but all the traffic is passed to the supervisor module.

Workaround: No workaround.

- CSCso49516

  Symptom: The Rx counter for the port channel increases by an extremely large number.

  Conditions: After you add a member interface to a port channel, the Rx counter for the port channel increases by an extremely large number.

  Workaround: No workaround.

- CSCso51277

  Symptom: Disabling the aging time for a VLAN removes the default aging time from the **show running-config** command output.

  Conditions: Aging time is incorrectly displayed in the **show running-config** command output following a supervisor module switchover.

  Workaround: No workaround.

- CSCso55225

  Symptom: Only one MAC address is secured by port security.

  Conditions: An ISSU from Release 4.0(1a) to a later release occurs with one of the following configurations on the device:

  - Both port security and 802.1X are enabled on a port that is down and then brought up.
  - A port is brought up with 802.1X multihost mode enabled.

  Workaround: Disable and enable port security on the interface.

- CSCso69259

  Symptom: VLAN interface counters have extremely large values.

  Conditions: The 5-minute rate counter values appear extremely large as VLAN interface in the **show interface** command output.

  Workaround: No workaround.

- CSCso72905

  Symptom: Input and output rate for Ethernet interfaces displays as 5-minute instead of 1-minute.

  Conditions: The input and output rate counters displayed in the **show interface** output for Ethernet interfaces does not have the correct rate unit. The rate counter is should display as 1 minute instead of 5 minute. The actual rate value is calculated per 1 minute.

  Workaround: No workaround.

- CSCso72732

  Symptom: A MIB walk does not display correct values for unicast packet in and out counters for VLAN interfaces.

Conditions: The unicast packet in and out counters are updated with multicast packet in and out counter values. Since the counters are not updated correctly, MIB walk does not reflect correct values for unicast packet counts.

The **show interface counters** command and **show interface vlan counters** command display the correct counter values.

Workaround: No workaround.

- CSCso84540

  Symptom: The Ethernet interfaces do not come up after a reload when DHCP and PBR statistics are configured on the same VLAN interfaces.

  Conditions: After a system reload, the Ethernet interfaces do not come up when the startup configuration configures DHCP relay on one or more of the VLAN interfaces and policy based routing (PBR) on one or more of the VLAN interfaces.

  Workaround: Do not enable DHCP on the VLAN interfaces when you want to use PBR statistics on the VLAN interface. Specifically, do not used the **service dhcp**, **ip dhcp snooping**, and **ip dhcp snooping vlan** commands on the VLAN interfaces.

- CSCso92283

  Symptom: Route-map/VACL on an ACL ignores the deny actions defined in the ACL.

  Conditions: The route-map/VACL configuration ignores the deny actions for rules in an ACL and forwards any packet that matches either permit or deny rule within the match ACL.

  For example, you have a route map with the following route-map and VACL configuration:

```
route-map policy-foo permit 10
  match ip address foo-list
  set ip next-hop 10.20.30.40
```

  The ACL foo-list has the following rule configuration:

```
ip access-list foo-list
  20 permit tcp 10.20.220.0/23 range 1000 4000 10.30.0.0/16
  30 permit udp 10.20.220.0/23 range 1000 4000 10.30.0.0/16
  40 permit tcp 10.10.20.0/23 eq 1433 10.30.0.0/16
  50 deny tcp 10.20.30.0/23 any
```

  The packets matching the deny rule are also forwarded to the next hop configured in the route map.

  Workaround: No workaround.

- CSCso84540

  Symptom: The Ethernet interfaces do not come up after a reload when DHCP and PBR statistics are configured on the same VLAN interfaces.

  Conditions: After a system reload, the Ethernet interfaces do not come up when the startup configuration configures DHCP relay on one or more of the VLAN interfaces and PBR on one or more of the VLAN interfaces.

  Workaround: Do not enable DHCP on the VLAN interfaces when you want to use PBR statistics on the VLAN interface. Specifically, do not used the **service dhcp**, **ip dhcp snooping**, and **ip dhcp snooping vlan** commands on the VLAN interfaces.

- CSCso85622

  Symptom: Changing the summary address for EIGRP on an interface causes EIGRP to hard reset the neighbors.

  Conditions: Changing the summary address for EIGRP using the **ip summary-address eigrp** command on an interface causes EIGRP to hard reset the neighbors, rather than a soft reset.

  Workaround: No workaround.

- CSCso91087

  Symptom: The default route specified the **redistribute static route-map** command is not used.

  Conditions: You have specified the default route for EIGRP in the **redistribute static route-map** command.

  Workaround: No workaround.

- CSCso93173

  Symptom: When remove EIGRP router configuration, the NX-OS software does not send goodbye messages to the device neighbors before stopping the EIGRP process.

  Conditions: Removing the EIGRP router configuration with the **no router eigrp** command stops the EIGRP process without sending goodbye messages to the neighbors.

  Workaround: No workaround.

- CSCso93525

  Symptom: The packets received through the member ports in a Layer 2 port channel are not routed according to the PBR policy.

  Conditions: When you apply a PBR policy on a VLAN interface that has Layer 2 port channel in that VLAN, the Layer 2 member ports of the port channel do not receive the PBR policy.

  Workaround: No workaround.

- CSCso98880

  Symptom: Redistribution of BGP by EIGRP does not occur after a reload

  Conditions: The **redistribute bgp** command is in the startup configuration for EIGRP. This problem affects only BGP as the redistributing source protocol. Restarting EIGRP resumes normal redistribution of the BGP routes.

  Workaround: No workaround.

- CSCsq13120

  Symptom: EIGRP does not clear all neighbors and form new adjacencies with the new autonomous system (AS) number, when the AS number changes.

  Conditions: You have configured the **router autonomous-system** command for EIGRP. When the AS number changes, EIGRP does not clear the neighbors and form new adjacencies with the new AS number.

  Workaround: No workaround.

- CSCsq19514

  Symptom: The EIGRP process may not respond and is terminated by the system manager (sysmgr) process.

  Conditions: When you have a very large number of routes (as many as 55,000) in the EIGRP topology table, table cleanup can take very long time (more than 4 minutes), especially after a neighbor goes down.

  Workaround: No workaround.

- CSCsq20726

  Symptom: Removing a VDC may take up to 5 minutes to complete.

  Conditions: This issue can occur if the Unicast FIB distribution module (UFDM) process in the VDC has restarted either manually or due to an unexpected condition.

  Workaround: No workaround.

- CSCsq25658

  Symptom: The following error message displays in the syslog:

  ```
  EIGRP-3-IP_INTERNAL_ERROR" "Failed to get IP API VRF w
  ```

  Conditions: The message displays when a nonexisting VRF is referenced under router EIGRP mode. This message does not affect EIGRP functioning.

  Workaround: No workaround.

- CSCsq43675

  Symptom: Software PBR statistics may be incorrect.

  Conditions: After you clear the software PBR statistics with the **clear route-map** *route* **pbr-statistics** command and then send traffic, the output from the first time you use the **show route-map** *route* **pbr-statistics** command is incorrect. This problem has no functional impact on the feature.

  Workaround: Use the **show route-map** *route* **pbr-statistics** command again to display the correct software PBR statistics.

- CSCsq47196

  Symptom: In a rare occurrence, you may see that the members of a port channel are in the error-disabled state with reason Internal-Fail errDisable. The error code during this condition is Port-channel Membership Is Being Updated.

  Conditions: This situation may occur when a port-channel member is coming up at the same time that member is being removed from the port channel or that port channel is being deleted.

  Workaround: Enter the **shutdown** command and then the **no shutdown** command on the specified interfaces.

- CSCsq51647

  Symptom: All of the OIFs of (*,G) inherited to (S,G) might not be visible.

  Conditions: This problem only occurs after a switchover. When it occurs, only a few (S,G)s might be in this state.

  Workaround: Restart IGMP.

- CSCsq60142

  Symptom: Traffic might be lost in nondefault VDCs.

  Conditions: This problem can occur when you have more than one nondefault VDC with VLAN interfaces and you have created and deleted several VLAN interfaces in the nondefault VDCs.

  Workaround: No workaround.

- CSCsq71625

  Symptom: You cannot use the CISCO-CONFIG-COPY-MIB to copy a configuration file to or from a remote server if it is not reachable through the default VRF.

  Conditions: CISCO-CONFIG-COPY-MIB does not allow you to specify a VRF for transfers involving remote servers.

  Workaround: Make the server reachable through the VRF.

- CSCsr25585

  Symptom: The device drops the first fragments of large UDP packets as they are switched through the devices. All other fragments of these UDP packets are switched properly.

  Conditions: This situation occurs when you have enabled the maximum UDP length check. You display the status of this check by entering the **show hardware forward ip verify** command. The device displays whether fragmented UDP packets are being switched through the device.

  Workaround: Disable the IDS check for the maximum UDP length by entering the **no platform ip verify length maximum udp** command.

- CSCsr30773

  Symptom: When you are booting an image, the supervisor module may reset or switch over to the backup supervisor.

  Conditions: When the device is negotiating CTS sessions, the active supervisor might perform a switchover, or the system might reset when you are loading an image.

  Workaround: Do not enable CTS unless you have upgraded the IOFPGA to version 3.19. This situation does not affect the running system unless you enable CTS. However, because this bug can also cause a supervisor reset at bootup in rare circumstances, we recommend that you upgrade the IOFPGA to version 3.19 at the earliest available downtime window.

- CSCsr39659

  Symptom: An interface in a VRF is running active HSRP in both Nexus 7000 series devices. Each device does not see the other device.

Conditions: This situation occurs when you reload a device that has the command **hsrp delay reload** configured. It may also occur at the initial configuration,

Workaround: Enter the **shutdown** command and then enter the **no shutdown** command for the specified interface on the reloaded device.

- CSCsr56858

  Symptom: When you downgrade from NX-OS Release 4.0(3) to Release 4.0(2) on a device that runs port channels and remove any ports from a port channel, you will see errors when the device applies queuing policies on those ports.

  Conditions: When there is a supervisor switchover that runs the NX-OS Release 4.0(2) image or a downgrade is performed from the Release 4.0(3) image to the Release 4.0(2) image on a device that has port channels, you will see errors while the device applies queuing policies on the ports after they are removed from the port channels.

  Workaround: Use NX-OS Release 4.0(3).

- CSCsr81741

  Symptom: If you do not change the STP path-cost method to short after changing the STP mode from MST to Rapid PVST+ before an ISSU or a switchover, the device might trigger a topology change.

  Conditions: If you select the path-cost method as short, which is the default for Rapid PVST+, and you change the STP mode from Rapid PVST+ to MST and back to Rapid PVST+, the path-cost method remains as long. The device does not automatically return the path-cost method to short when you change the STP mode to Rapid PVST+ from MST. If you perform an ISSU or a switchover, STP sends BPDUs using different path-cost methods, which could trigger a topology change. (Prior to an ISSU or a switchover, STP sent BPDUs using the long path-cost method. After an ISSU or a switchover, STP sends BPDUs using the short path-cost method.)

  Workaround: Change the path-cost method to short when you switch to Rapid PVST+ before you perform an ISSU by entering the **spanning-tree pathcost method** command.

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

This section describes possible unexpected behavior by Cisco NX-OS Release 4.0(1a). All the caveats listed in this section are resolved in Cisco NX-OS Release 4.0(2).

- CSCsk51803

  Symptom: GLBP groups cannot be configured using a range of interfaces.

  Conditions: The GLBP feature does not support interface range configuration.

  Workaround: You can configure the GLBP group on a per-interface basis for the range of interfaces.

- CSCsl90405

  Symptom: When you delete a nondefault VDC, the startup configuration is not removed from the file system.

Conditions: When you create a nondefault VDC, save its running configuration to the startup configuration and then delete it, the startup configuration is not removed from the nvram: file system.

Workaround: No workaround.

- CSCsm00267

  Symptom: If you change an active port mode from access to trunk, addresses are learned on the wrong VLAN.

  Conditions: An active port has MAC addresses secured on it. If you change an active port from access to trunk, it could cause the addresses to be learned on the wrong VLAN.

  Workaround: Instead of changing an active interface from access to trunk, first shut down the interface.

  To change a interface mode from access to trunk, follow these steps:

  1. Shut down the interface by entering the **shutdown** command.
  2. Change the interface mode.
  3. Bring the interface back up using the **no shutdown** command.

- CSCsm30176

  Symptom: Frames are dropped under rare conditions on interfaces that operate at 10 Mbps or 100 Mbps.

  Conditions: The MAC continuously drops frames after an interface in the administratively down state comes up at 10 Mbps or 100 Mbps after a module reload.

  Workaround: Perform one of the following workarounds:

  - Enter the **shutdown/no shutdown** command sequence on the interface.
  - Operate the interface at 1 Gbps. This situation does not occur at 1-Gbps speeds; it occurs only at 10-Mbps or 100-Mbps speeds.

- CSCsm63787

  Symptom: Secured MAC addresses are not relearned after a supervisor module switchover.

  Conditions: After a supervisor module switchover, the MAC addresses secured by port security are aged out due to the absolute timer and are not relearned. This problem may occur during a continuous traffic condition.

  Workaround: No workaround.

- CSCso04360

  Symptom: The **no clock summer-time** command is not supported.

  Conditions: The **no clock summer-time** command results in the following message:

  ```
  EOL function syscli_config_time_zone from library libsyscli.so exited
  ```

  Workaround: No workaround.

- CSCso09108

  Symptom: CMP uses local authentication after remote authentication fails three consecutive times.

  Conditions: When AAA-based remote authentication is configured on the CP side for the Admin user, CMP uses the same for its authentication as well. However, if the remote authentication fails three consecutive times, the CMP login authentication falls back to an earlier configuration to check against the device Local Admin password.

  Authentication against the local password is automatic in the following cases:

  - AAA is not configured.
  - AAA server is not reachable.
  - CP is down, and the user tries to login to CMP.

  Workaround: No workaround.

- CSCso13876

  Symptom: The power supply serial number is missing from messages that are displayed when performing an online removal and insertion of a power supply.

  Conditions: If you insert a power supply into a running Cisco Nexus 7000 Series chassis, a PS_DETECT syslog is generated without a serial number. If the same power supply is subsequently removed, a PS_REMOVE syslog is generated without a serial number.

  Workaround: To obtain the serial number of the power supply, use the **show sprom powersupply** command.

- CSCso13909

  Symptom: The generic online diagnostic (GOLD) bootup tests may run even if they are disabled in the startup configuration.

  Conditions: The bootup diagnostic tests may run even if they are set to bypass.

  Workaround: No workaround.

- CSCso19300

  Symptom: The **reload cmp module** command may reload the local CMP instead of the standby.

  Conditions: When using the **reload cmp module** command from the standby CP, instead of reloading the CMP in the active supervisor module, it may cause the CMP in the standby slot to be reloaded instead.

  Workaround: Use the **reload cmp module** *number* command from the active supervisor module CLI.

- CSCso19498

  Symptom: The CMP clock does not synchronize with the control plane (CP)-side clock after you configure NTP.

  Conditions: The clock on the CMP side does not synchronize with the clock on the CP side when NTP is configured.

  Workaround: Do not configure NTP.

- CSCso28215

  Symptom: An ACL rule fragment flag match causes dropped packets.

  Conditions: Packets may be dropped if they have an initial fragment that matches an ACL rule with a fragment flag set.

  Workaround: Add an ACL rule before the ACL rule with the fragment flag to explicitly permit or deny initial fragments.

- CSCso29013

  Symptom: Incorrect counts are displayed in the **show interface error** command output.

  Conditions: The **show interface ethernet** *x/y* **errors** command output has the following errors:

  - Undersized packets are counted as runts and alignment errors.
  - Packets with CRC errors are counted as MAC transmission errors and display a huge value.

  Workaround: No workaround.

- CSCso43110

  Symptom: The tracking configuration uses a higher weight in GLBP during system boot.

  Conditions: Weighting is increased when loaded with the startup object tracking manager (OTM) interface.

  Workaround: No workaround.

- CSCso43110

  Symptom: The tracking configuration uses a higher weight in GLBP during system boot.

  Conditions: Weighting is increased when loaded with the startup object tracking manager (OTM) interface.

  Workaround: No workaround.

- CSCso70205

  Symptom: Multicast addresses on one or more I/O modules have become obsolete.

  Conditions: When you have configured nondefault VDCs and are downgrading from Cisco NX-OS Release 4.0(2) to Release 4.0(1a), the multicast addresses on one or more I/O modules may become obsolete.

  The supervisor module may not deliver multicast updates to the I/O module, which may result in the following symptoms:

  - RIB and FIB updates are not applied to the affected I/O modules.
  - Interface-related counters are not incremented (for example, VLAN counters).

  Workaround: Reload the affected I/O modules using the **reload module** command.

- CSCso82927

  Symptom: When you configure an ACL on the CMP with wildcard bits for the IP address, the denying or permitting of traffic based on the mask is reversed.

  Conditions: When adding an ACL on the CMP, the **permit** and **deny** command syntaxes allow either address/prefix or wildcard bits. If you type in the wildcard bits, it does not convert this to address/prefix. However, the CMP interprets the wildcard bits the same as a subnet mask and does not implement the ACL correctly.

  Workaround: Invert the mask bits in the wildcard field for IP addresses. For example, to permit traffic only from 172.18.112.x subnet, with destination ip 172.18.115.126, you would normally use the following commands:

```
switch-cmp(config-acl)# permit ip 172.18.112.0 0.0.0.255 172.18.115.126 0.0.0.0
switch-cmp(config-acl)# deny ip 0/0 0/0
switch-cmp(config-acl)# exit
switch-cmp(config)#
```

  Instead, use these commands:

```
switch-cmp(config-acl)# permit ip 172.18.112.0 255.255.255.0 172.18.115.126
255.255.255.255
switch-cmp(config-acl)# deny ip 0/0 0/0
switch-cmp(config-acl)# exit
switch-cmp(config)#
```

- CSCsq18260

  Symptom: Pause frames with protocol 0x8808 are transmitted on the port, whether traffic is occurring or not.

  Conditions: Tx pause frames are generated continuously when you bring up a port and enable flow control, whether traffic is occurring or not. The **show interface ethernet** *slot*/*port* **flowcontrol** command displays the total TX pause frames that have been transmitted so far.

  Workaround: No workaround.

- CSCsq60582

  Symptom: Multiple Cisco products contain either of two authentication vulnerabilities in the Simple Network Management Protocol version 3 (SNMPv3) feature. These vulnerabilities can be exploited when processing a malformed SNMPv3 message. These vulnerabilities could allow the disclosure of network information or may enable an attacker to perform configuration changes to vulnerable devices. The SNMP server is an optional service that is disabled by default in Cisco products. Only SNMPv3 is impacted by these vulnerabilities.

Note SNMP versions 1, 2 and 2c are not impacted by these vulnerabilities.

  The United States Computer Emergency Response Team (US-CERT) has assigned Vulnerability Note VU#878044 to these vulnerabilities.

  Common Vulnerabilities and Exposures (CVE) identifier CVE-2008-0960 has also been assigned to these vulnerabilities.

  This advisory is posted at http://www.cisco.com/warp/public/707/cisco-sa-20080610-snmpv3.shtml.

  Conditions: This applies to only SNMPv3 SNMP packets.

Workaround: Use the authPriv level of SNMPv3 security.

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

This section describes possible unexpected behavior by Cisco NX-OS Release 4.0(1). All the caveats listed in this section are resolved in Cisco NX-OS Release 4.0(1a).

- CSCsm97999

  Symptom: The terminal session may hang if a Layer 3 protocol process hangs when executing a **show** command.

  Conditions: If a Layer 3 protocol hangs while executing a **show** command, the terminal session on which the command was entered also hangs.

  Workaround: Open another terminal session to the device.

- CSCso24745

  Symptom: When you are working with more than 50 PIM interfaces, the **restart pim** command may cause PIM to fail.

  Conditions: The **restart pim** command may cause PIM to fail if you are working with more than 50 PIM interfaces.

  Workaround: No workaround. If you switch over immediately to the standby supervisor module, the PIM processes can be recovered.

- CSCso25253

  Symptom: HSRPv1 VMAC is not removed following a migration to v2.

  Conditions: When migrating an interface from HSRP V1 to HSRP V2 or from V2 to V1, and one or more groups are configured, then the VMACs that correspond to the older version are not removed from the Layer 2 MAC table.

  The protocol behavior is not affected, but an extra entry in the MAC table is used for each group configured.

  Workaround: To prevent this problem, configure the HSRP version to be used first, and then configure the groups.

  If you have already migrated and have extra VMAC entries as a result, then use the following steps to clean up the old VMAC entries. These steps apply to a VMAC leak that is due to a migration from V1 to V2.

  1. Remove V2 MACs by unconfiguring all groups.
  2. Migrate back to V1.
  3. Configure groups and unconfigure them.
  4. Migrate to V2 and configure the groups.

- CSCso32746

  Symptom: SSH login is not supported on a nondefault VDC with a public key and no password.

  Conditions: You cannot use SSH to log in to a nondefault VDC using a public key without a password.

Workaround: No workaround.

- CSCso35621

  Symptom: When using a file to configure SSH key generation, you are logged out if connected through SSH.

  Conditions: Because the SSH key generation process disables the SSH server, if you generate the keys using a file and are logged in through SSH, you will be logged out of the session during this process.

  The SSH key generation process is as follows:

  1. Internally disables the SSH server if the server is enabled.
  2. Generates the keys.
  3. Restores the SSH server to its previous state.

  If you enter the key generation commands manually, you must disable the SSH server before you can generate the key. Key generation succeeds only if an enabled SSH server is explicitly disabled.

  Workaround: No workaround.

- CSCso40114

  Symptom: HSRPv2 groups may not come up automatically after a module reload.

  Conditions: When HSRP version 2 is configured on an interface and you perform a module reload, the groups on the interface may come back up in version 1 mode.

  Workaround: Reconfigure V2 after a module reload.

  From interface configuration mode, following these steps:

  1. Enter the **no hsrp version** command.

     The version is reset to V1 in the configuration.
  2. Enter the **hsrp version 2** command.

     This action forces the entire group to V2.

- CSCso43181

  Symptom: The CoPP policy, under some conditions, is not downloaded to the module.

  Conditions: If the device is reloaded with more than two modules and large ACLs, the CoPP policies might not download to some modules.

  Workaround: Check if policies are downloaded to the module by using the **show system internal qos copp module** command. Reload the module that did not receive the policy.

- CSCso43676

  Symptom: A false instance of OSPF may be started after a switchover.

  Conditions: If the system is very busy during a system switchover, it may create a false OSPF instance. This results in OSPF adjacencies not coming up.

  Workaround: Remove and reconfigure OSPF by copying the startup-config to the running-config.

- CSCso45855

  Symptom: The drop Tail/Weighed Random Early Discard queueing configurations are incorrect

  Conditions: When configuring the class of service (CoS) thresholds for the **queue-limit cos** and **random-detect cos** commands, the hardware threshold level may not be programmed correctly if the ports are not in the administratively down (shut) state or if the queue size is not previously set. The ratios between thresholds are maintained and the traffic with lower threshold CoS values are still dropped before the traffic with higher threshold CoS values, but the physical threshold levels may not match your configuration.

  The following policy, for example, applied directly to an interface, may be programmed with the incorrect threshold values.

```
policy-map type queuing foo
    class type queuing 1p3q4t-out-pq1
      random-detect cos-based
      random-detect cos 5 minimum-threshold 500 kbytes maximum-threshold 1000 kbytes
      random-detect cos 6 minimum-threshold percent 30 maximum-threshold percent 60
    class type queuing 1p3q4t-out-q-default
      queue-limit cos 0 percent 10
      queue-limit cos 1 600 kbytes
```

  Workaround: To guarantee that a configuration is applied correctly in the device, follow these steps:

  1. Create a queuing policy with the desired name and configure just the **queue-limit** command in any appropriate class map. By default, when a policy is configured, the queue-limit size is evenly distributed among all queues.

```
switch# configure terminal
switch(config)# policy-map type queuing foo
switch(config-pmap-que)# class type queuing 1p3q4t-out-q-default
switch(config-pmap-c-que)# queue-limit percent 25
switch(config-pmap-c-que)# exit
switch(config-pmap-que)# exit
switch(config)#
```

  2. Disable the desired interface.

```
switch(config)# interface ethernet 1/1
switch(config-if)# shutdown
```

  3. Apply the newly-created policy on the desired interface/direction.

```
switch(config-if)# service-policy type queuing output foo
```

  4. Add all desired threshold actions to the applied policy.

```
switch(config-pmap-que)# class type queuing 1p3q4t-out-pq1
switch(config-pmap-c-que)# random-detect cos-based
switch(config-pmap-c-que)# random-detect cos 5 minimum-threshold 500 kbytes
maximum-threshold 1000 kbytes
switch(config-pmap-c-que)# random-detect cos 6 minimum-threshold percent 30
maximum-threshold percent 60
switch(config-pmap-c-que)# class type queuing 1p3q4t-out-q-default
switch(config-pmap-c-que)# queue-limit cos 0 percent 10
switch(config-pmap-c-que)# queue-limit cos 1 600 kbytes
switch(config-pmap-c-que)# exit
switch(config-pmap-que)# exit
switch(config)#
```

  5. Reenable the interface.

```
switch(config)# interface ethernet 1/1
```

```
switch(config-if)# no shutdown
```

The thresholds are now correctly configured for the interface.

- CSCso48993

  Some output interfaces (OIFs) may be missing after a module reload.

  Conditions: After a module reload, some (S, G) routes may be missing a few OIFs.

  Workaround: Perform a controlled restart of the route owner (IGMP or PIM) to correct this problem.

- CSCso53748

  Symptom: Enabling or disabling the beacon mode causes the port to flap.

  Conditions: When a port is up and the beacon mode is enabled or disabled, the port state is momentarily disrupted from up to down and back to up and disrupts data traffic.

  Workaround: To prevent an impact to data traffic, enable the beacon mode only when the port is in the administratively down state.

- CSCso54679

  Symptom: Under certain conditions following a device reload, a Layer 2 port-channel member port comes up as a Layer 3 port.

  Conditions: A port channel and its member ports come up with different layer information after a device reboot under the following conditions:

  - The layer mode of the port is different from that of the port channel.
  - The port is added to the Layer 2 port channel by force.
  - The configuration is saved.

  Workaround: Before adding a port to a channel group, first change its layer mode so that it is the same as that of the port channel.

- CSCso59936

  Symptom: The NetFlow Manager (NFM) process can crash when egress NetFlow is configured.

  Conditions: When you configure an egress NetFlow, the NFM process can crash.

  Workaround: Disable egress NetFlows.

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

This was the first release for the Cisco NX-OS software.

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_0/epld/release/notes/epld_rn.html

The following are related Cisco NX-OS documents:

NX-OS Configuration Guides

Cisco Nexus 7000 Series NX-OS Getting Started with Virtual Device Contexts, Release 4.0

Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS Software Upgrade Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS Licensing Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS XML Management Interface User Guide, Release 4.0

Cisco Nexus 7000 Series NX-OS System Messages Reference

Cisco Nexus 7000 Series NX-OS MIB Quick Reference

NX-OS Command References

Cisco Nexus 7000 Series NX-OS Command Reference Master Index, Release 4.0

Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference, Release 4.0

Cisco Nexus 7000 Series NX-OS Interfaces Command Reference, Release 4.0

Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference, Release 4.0

Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference, Release 4.0

Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.0

Cisco Nexus 7000 Series NX-OS Multicast Routing Command Reference, Release 4.0

Cisco Nexus 7000 Series NX-OS Security Command Reference, Release 4.0

Cisco Nexus 7000 Series NX-OS Virtual Device Context Command Reference, Release 4.0

Cisco Nexus 7000 Series NX-OS System Management Command Reference, Release 4.0

Other Software Document

Cisco Nexus 7000 Series NX-OS Troubleshooting Guide, Release 4.0

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.
