

Arista Networks EOS 4.29.6.1M Release Notes

Version: 1.0

20 March 2024

Table of Contents

EOS 4.29.6.1M Release Highlights

New Platforms and Hardware

SNMP MIBs

SNMP Traps

Supported Hardware

Reference to MLAG ISSU Upgrade/Downgrade Table

Resolved Caveats in 4.29.6.1M

CVP

Network Provisioning

Known Software Caveats in 4.29.6.1M

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

vEOS Router / CloudEOS

CVP

Network Provisioning

CVX

Enterprise routing

General

IPSEC

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

Multicast

NAT

SwitchApp

DCS-7170 Series

DCS-7170 Series

CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM and CCS-750 Series

CCS-750 and DCS-7060X4 Series

DCS-7170, DCS-7280R and DCS-7280R2 Series

DCS-7130 Series

7358, 7368, 7388, CCS-750, DCS-7300X, DCS-7500R, DCS-7800 and DCS-7800R3 Series

7500R3, DCS-7020, DCS-7130, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280R3A, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800, DCS-7800R3 and DCS-7800R3A Series

DCS-7280R and DCS-7500R Series

DCS-7280R2 and DCS-7500R2 Series

7500R3, DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series

7500R3, DCS-7130, DCS-7280CR3, DCS-7500R3 and DCS-7800R3 Series

DCS-7280R3A and DCS-7800R3A Series

DCS-7020 Series

DCS-7280R3 Series

7500R3, DCS-7500R3 and DCS-7800R3 Series

DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

7500R3, DCS-7130, DCS-7280R3, DCS-7280R3A, DCS-7500R3, DCS-7800R3 and DCS-7800R3A Series

DCS-7280R and DCS-7280R2 Series

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4,

DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7260X, DCS-7260X3 and DCS-7300X Series

CCS-750 Series

7300X3, 7320X, DCS-7060X and DCS-7260X Series

DCS-7260X3 Series

7368 and DCS-7060X4 Series

DCS-7060X5 Series

7388 and DCS-7060X5 Series

CCS-710P, CCS-720DP and CCS-720DT Series

CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series

CCS-720DF and CCS-720XP Series

CCS-750 Series

DCS-7050X3 Series

7300X3, AS7326, AS7726 and DCS-7050X3 Series

7358 and DCS-7050X4 Series

7388 and DCS-7060X5 Series

Transceiver Series

7300X3, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX and DCS-7050X3 Series

7358 and DCS-7050X4 Series

Limitations and Restrictions in 4.29.6.1M

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

Containerized EOS

vEOS Router / CloudEOS

CVX

General

IPSEC

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

MPLS

Multicast

NAT

SwitchApp

DCS-7170 Series

DCS-7170 Series

CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM and CCS-750 Series

CCS-750 and DCS-7060X4 Series

DCS-7170, DCS-7280R and DCS-7280R2 Series

DCS-7130 Series

7358, 7368, 7388, CCS-750, DCS-7300X, DCS-7500R, DCS-7800 and DCS-7800R3 Series

CCS-710P, CCS-720DP, CCS-720XP, CCS-722XPM and CCS-750 Series

7500R3, DCS-7020, DCS-7130, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280R3A, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800, DCS-7800R3 and DCS-7800R3A Series

DCS-7280R and DCS-7500R Series

DCS-7280R2 and DCS-7500R2 Series

7500R3, DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series

7500R3, DCS-7130, DCS-7280CR3, DCS-7500R3 and DCS-7800R3 Series

DCS-7280R3A and DCS-7800R3A Series

DCS-7020 Series

DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

7500R3, DCS-7130, DCS-7280R3, DCS-7280R3A, DCS-7500R3, DCS-7800R3 and DCS-7800R3A Series

DCS-7280R and DCS-7280R2 Series

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7260X, DCS-7260X3 and DCS-7300X Series

CCS-750 Series

7300X3, 7320X, DCS-7060X and DCS-7260X Series

DCS-7260X3 Series

7368 and DCS-7060X4 Series

CCS-710P, CCS-720DP and CCS-720DT Series

CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series

CCS-720DF and CCS-720XP Series

CCS-750 Series

7300X3, AS7326, AS7726 and DCS-7050X3 Series

7358 and DCS-7050X4 Series

Transceiver Series

7300X3, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX and DCS-7050X3 Series

7358 and DCS-7050X4 Series

Conditions and Impacts

EOS 4.29.6.1M Release Highlights

New Platforms and Hardware

- None

SNMP MIBs

- ARISTA-ACL-MIB
- ARISTA-ASIC-COUNTERS-MIB
- ARISTA-BGP4V2-MIB
- ARISTA-BRIDGE-EXT-MIB
- ARISTA-CONFIG-COPY-MIB
- ARISTA-CONFIG-MAN-MIB
- ARISTA-DAEMON-MIB
- ARISTA-ENTITY-SENSOR-MIB
- ARISTA-FIB-STATS-MIB
- ARISTA-GENERAL-MIB
- ARISTA-HARDWARE-UTILIZATION-MIB
- ARISTA-IF-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-MAU-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-NEXTHOP-GROUP-MIB
- ARISTA-PFC-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-PRODUCTS-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-QOS-MIB
- ARISTA-QUEUE-MIB
 - Excluding 7150
- ARISTA-REDUNDANCY-MIB
 - 7500R, 7500E, 7300X, 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-SMI-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-SNMP-TRANSPORTS-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-SW-IP-FORWARD-MIB
 - IPv4 only
- ARISTA-SW-IP-FORWARDING-MIB
- ARISTA-TEST-MIB
- ARISTA-VRF-MIB
- ARISTA-VXLAN-MIB
- ARISTA-XCVR-DWDM-MIB
- ARISTA-XGS-MIB

- ENTITY-MIB
- ENTITY-SENSOR-MIB
- ENTITY-STATE-MIB
- HC-RMON-MIB
 - etherStatsHighCapacityGroup
- HOST-RESOURCES-MIB
- IEEE8021-PFC-MIB
- IEEE8021-SECY-MIB
- IEEE8023-LAG-MIB
- IGMP-MIB
- IPMCAST-MIB
- IPMROUTE-STD-MIB
- LLDP-EXT-DOT1-MIB
- LLDP-EXT-DOT3-MIB
- LLDP-MIB
- MPLS-LDP-STD-MIB
- MPLS-TE-STD-MIB
- MSDP-MIB
- NOTIFICATION-LOG-MIB
- OSPF-TRAP-MIB
- OSPFV3-MIB
- PIM-MIB
- PTPBASE-MIB
- RFC 2863 IF-MIB
 - obsoletes RFCs 1229, 1573, 2233
- RFC 2864 IF-INVERTED-STACK-MIB
- RFC 3418 SNMPv2-MIB
 - obsoletes RFCs 1450, 1907
- RFC 3635 EtherLike-MIB
 - obsoletes RFCs 1650, 2358, 2665
- RFC 3636 MAU-MIB
 - ifMauDefaultType and ifMauAutoNegStatus are writeable
- RFC 4022 TCP-MIB
 - obsoletes RFC 1213
- RFC 4113 UDP-MIB
 - obsoletes RFC 1213
- RFC 4188 BRIDGE-MIB
- RFC 4273 BGP4-MIB
- RFC 4292 IP-FORWARD-MIB
 - obsoletes RFC 1354
- RFC 4293 IP-MIB
 - obsoletes RFC 1213
- RFC 4363 Q-BRIDGE-MIB
- RFC 4750 OSPF-MIB
- RMON-MIB
 - rmonEtherStatsGroup

- RMON2-MIB
 - rmon1EthernetEnhancementGroup
- SNMP-TLS-TM-MIB
 - RFC 6353
- SNMP-TSM-MIB
 - RFC 5591
- SNMPv2, SNMPv3
- VRRPV2-MIB

SNMP Traps

- ARISTA-BRIDGE-EXT-MIB
 - aristaMacMove, aristaMacLearn, aristaMacAge
- ARISTA-CONFIG-MAN-MIB
 - aristaConfigManEvent
- ARISTA-REDUNDANCY-MIB
 - aristaRedundancySwitchOverNotif only for: 7500R, 7500E, 7300X, 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- BGP4-MIB
 - bgpEstablished, bgpBackwardTransition
- ENTITY-MIB
 - entConfigChange
- ENTITY-STATE-MIB
 - entStateOperEnabled, entStateOperDisabled
- ISIS-MIB
 - isisDatabaseOverload, isisAttemptToExceedMaxSequence, isisOwnLSPPurge, isisSequenceNumberSkip, isisAuthenticationTypeFailure, isisAreaMismatch, isisRejectedAdjacency, isisAdjacencyChange
- LLDP-MIB
 - lldpRemTablesChange
- MPLS-LDP-STD-MIB
 - mplsLdpSessionUp, mplsLdpSessionDown
- MSDP-MIB
 - msdpEstablished, msdpBackwardTransition
- NET-SNMP-AGENT-MIB
 - nsNotifyRestart
- OSPF-MIB
 - ospfNbrStateChange, ospfIfConfigError, ospfIfAuthFailure, ospfIfStateChange
- PIM-MIB
 - pimNeighborLoss
- RFC 2863 IF-MIB
 - linkUp, linkDown
- RFC 3418 SNMPv2-MIB
 - coldStart
- SNMPv2-MIB
 - authenticationFailure

- VRRPv2-MIB
 - vrrpTrapNewMaster

Supported Hardware

Fixed System

* AS7326-56X	* DCS-7130-48LS-R (HwRev:B2)	* DCS-7280SR-48C6-F
* AS7726-32X	* DCS-7130-48LS-R (HwRev:B3)	* DCS-7280SR-48C6-M-F
* CCS-710P-12	* DCS-7130-96	* DCS-7280SR-48C6-M-R
* CCS-710P-16P	* DCS-7130-96-F (HwRev:A7)	* DCS-7280SR-48C6-R
* CCS-720DF-48Y	* DCS-7130-96-F (HwRev:A8)	* DCS-7280SR2-48YC6
* CCS-720DF-48Y-2	* DCS-7130-96-R (HwRev:A7)	* DCS-7280SR2-48YC6-F
* CCS-720DF-48Y-2F	* DCS-7130-96-R (HwRev:A8)	* DCS-7280SR2-48YC6-M-F
* CCS-720DF-48Y-F	* DCS-7130-96E	* DCS-7280SR2-48YC6-M-R
* CCS-720DP-24S	* DCS-7130-96E-F (HwRev:A7)	* DCS-7280SR2-48YC6-R
* CCS-720DP-24S-2	* DCS-7130-96E-R (HwRev:A7)	* DCS-7280SR2A-48YC6
* CCS-720DP-24S-2F	* DCS-7130-96L	* DCS-7280SR2A-48YC6-F
* CCS-720DP-24S-F	* DCS-7130-96L-F (HwRev:A2)	* DCS-7280SR2A-48YC6-M-F
* CCS-720DP-48S	* DCS-7130-96L-F (HwRev:A3)	* DCS-7280SR2A-48YC6-M-R
* CCS-720DP-48S-2	* DCS-7130-96L-R (HwRev:A2)	* DCS-7280SR2A-48YC6-R
* CCS-720DP-48S-2F	* DCS-7130-96L-R (HwRev:A3)	* DCS-7280SR2K-48C6
* CCS-720DP-48S-F	* DCS-7130-96LA	* DCS-7280SR2K-48C6-M-F
* CCS-720DT-24S	* DCS-7130-96LA-F (HwRev:A2)	* DCS-7280SR2K-48C6-M-R
* CCS-720DT-24S-2	* DCS-7130-96LA-F (HwRev:A3)	* DCS-7280SR3-40YC6
* CCS-720DT-24S-2F	* DCS-7130-96LA-R (HwRev:A2)	* DCS-7280SR3-40YC6-F
* CCS-720DT-24S-F	* DCS-7130-96LA-R (HwRev:A3)	* DCS-7280SR3-40YC6-R
* CCS-720DT-48S-2	* DCS-7130-96LAS	* DCS-7280SR3-48YC8
* CCS-720DT-48S-2F	* DCS-7130-96LAS-F (HwRev:B1)	* DCS-7280SR3-48YC8-F
* CCS-720XP-24Y6	* DCS-7130-96LAS-F (HwRev:B2)	* DCS-7280SR3-48YC8-R
* CCS-720XP-24Y6-F	* DCS-7130-96LAS-R (HwRev:B1)	* DCS-7280SR3E-40YC6
* CCS-720XP-24Y6-R	* DCS-7130-96LAS-R (HwRev:B2)	* DCS-7280SR3E-40YC6-F
* CCS-720XP-24ZY4	* DCS-7130-96LB	* DCS-7280SR3E-40YC6-R
* CCS-720XP-24ZY4-F	* DCS-7130-96LB-F (HwRev:A2)	* DCS-7280SR3K-48YC8
* CCS-720XP-48Y6	* DCS-7130-96LB-F (HwRev:A3)	* DCS-7280SR3K-48YC8-F
* CCS-720XP-48Y6-F	* DCS-7130-96LB-R (HwRev:A2)	* DCS-7280SR3K-48YC8-R
* CCS-720XP-48Y6-R	* DCS-7130-96LB-R (HwRev:A3)	* DCS-7280SR3K-48YC8A
* CCS-720XP-48ZC2	* DCS-7130-96LBA	* DCS-7280SR3K-48YC8A-F
* CCS-720XP-48ZC2-F	* DCS-7130-96LBA-F (HwRev:A2)	* DCS-7280SR3K-48YC8A-R
* CCS-720XP-48ZC2-R	* DCS-7130-96LBA-F (HwRev:A3)	* DCS-7280SR3M-48YC8
* CCS-720XP-96ZC2	* DCS-7130-96LBA-R (HwRev:A2)	* DCS-7280SR3M-48YC8-F
* CCS-720XP-96ZC2-F	* DCS-7130-96LBA-R (HwRev:A3)	* DCS-7280SR3M-48YC8-R

* CCS-722XPM-48Y4	* DCS-7130-96LBAS	* DCS-7280SRA-48C6
* CCS-722XPM-48Y4-F	* DCS-7130-96LBAS-F (HwRev:B1)	* DCS-7280SRA-48C6-F
* CCS-722XPM-48ZY8	* DCS-7130-96LBAS-R (HwRev:B1)	* DCS-7280SRA-48C6-M-F
* CCS-722XPM-48ZY8-F	* DCS-7130-96LBS	* DCS-7280SRA-48C6-M-R
* DCS-7010TX-48	* DCS-7130-96LBS-F (HwRev:B1)	* DCS-7280SRA-48C6-R
* DCS-7010TX-48-DC	* DCS-7130-96LBS-F (HwRev:B2)	* DCS-7280SRAM-48C6
* DCS-7010TX-48-DC-F	* DCS-7130-96LBS-R (HwRev:B1)	* DCS-7280SRAM-48C6-F
* DCS-7010TX-48-DC-R	* DCS-7130-96LBS-R (HwRev:B2)	* DCS-7280SRAM-48C6-R
* DCS-7010TX-48-F	* DCS-7130-96LS	* DCS-7280SRM-40CX2
* DCS-7010TX-48-R	* DCS-7130-96LS-F (HwRev:B1)	* DCS-7280SRM-40CX2-F
* DCS-7020SR-24C2	* DCS-7130-96LS-F (HwRev:B2)	* DCS-7280SRM-40CX2-R
* DCS-7020SR-24C2-F	* DCS-7130-96LS-R (HwRev:B1)	* DCS-7280TR-48C6
* DCS-7020SR-24C2-R	* DCS-7130-96LS-R (HwRev:B2)	* DCS-7280TR-48C6-F
* DCS-7020SR-32C2	* DCS-7130-96S	* DCS-7280TR-48C6-M-F
* DCS-7020SR-32C2-F	* DCS-7130-96S-F (HwRev:B1)	* DCS-7280TR-48C6-M-R
* DCS-7020SR-32C2-R	* DCS-7130-96S-R (HwRev:B1)	* DCS-7280TR-48C6-R
* DCS-7020SRG-24C2	* DCS-7130LBR-48S6QD	* DCS-7280TRA-48C6
* DCS-7020SRG-24C2-F	* DCS-7130LBR-48S6QD-F	* DCS-7280TRA-48C6-F
* DCS-7020SRG-24C2-R	* DCS-7130LBR-48S6QD-R	* DCS-7280TRA-48C6-M-F
* DCS-7020TR-48	* DCS-7170-32C	* DCS-7280TRA-48C6-M-R
* DCS-7020TR-48-F	* DCS-7170-32C-F	* DCS-7280TRA-48C6-R
* DCS-7020TR-48-R	* DCS-7170-32C-M-F	* FAN-7000
* DCS-7020TRA-48	* DCS-7170-32C-M-R	* FAN-7000-F
* DCS-7020TRA-48-F	* DCS-7170-32C-R	* FAN-7000-R
* DCS-7020TRA-48-R	* DCS-7170-32CD	* FAN-7000H-R
* DCS-7050CX3-32S	* DCS-7170-32CD-F	* FAN-7001C
* DCS-7050CX3-32S-F	* DCS-7170-32CD-R	* FAN-7001C-F
* DCS-7050CX3-32S-R	* DCS-7170-64C	* FAN-7001C-R
* DCS-7050CX3-32S-SSD-F	* DCS-7170-64C-F	* FAN-7001D
* DCS-7050CX3-32S-SSD-R	* DCS-7170-64C-M-F	* FAN-7001D-F
* DCS-7050CX3M-32S	* DCS-7170-64C-M-R	* FAN-7001DH
* DCS-7050CX3M-32S-F	* DCS-7170-64C-R	* FAN-7001DH-F
* DCS-7050CX3M-32S-R	* DCS-7170B-64C	* FAN-7002H-R
* DCS-7050DX4-32S	* DCS-7170B-64C-F	* FAN-7010
* DCS-7050DX4-32S-F	* DCS-7260CX-64	* FAN-7010X
* DCS-7050PX4-32S	* DCS-7260CX-64-F	* FAN-7011H
* DCS-7050PX4-32S-F	* DCS-7260CX-64-R	* FAN-7011H-F
* DCS-7050SX3-48C8	* DCS-7260CX-64-SSD-F	* FAN-7011H-R
* DCS-7050SX3-48C8-F	* DCS-7260CX-64-SSD-R	* FAN-7011M
* DCS-7050SX3-48C8-R	* DCS-7260CX3-64	* FAN-7011M-F
* DCS-7050SX3-48YC12	* DCS-7260CX3-64-F	* FAN-7011M-R

* DCS-7050SX3-48YC12-F	* DCS-7260CX3-64-R	* FAN-7012H
* DCS-7050SX3-48YC8	* DCS-7260CX3-64E	* FAN-7012H-BLUE
* DCS-7050SX3-48YC8-F	* DCS-7260CX3-64E-F	* FAN-7012H-RED
* DCS-7050SX3-48YC8-R	* DCS-7260CX3-64E-R	* FAN-7012M
* DCS-7050SX3-96YC8	* DCS-7260CX3-64LQ	* FAN-7012M-BLUE
* DCS-7050SX3-96YC8-F	* DCS-7260CX3-64LQ-F	* FAN-7012M-RED
* DCS-7050SX3-96YC8-R	* DCS-7260CX3-64LQ-R	* FAN-7130-1U
* DCS-7050TX3-48C8	* DCS-7260QX-64	* FAN-7130-1U-F
* DCS-7050TX3-48C8-F	* DCS-7260QX-64-F	* FAN-7130-1U-R
* DCS-7050TX3-48C8-R	* DCS-7260QX-64-R	* FAN-7130-2U
* DCS-7060CX-32C	* DCS-7260QX-64-SSD-F	* FAN-7130-2U-F
* DCS-7060CX-32C-F	* DCS-7260QX-64-SSD-R	* FAN-7130-2U-R
* DCS-7060CX-32C-R	* DCS-7280CR-48	* PWR-1011-AC
* DCS-7060CX-32S	* DCS-7280CR-48-F	* PWR-1011-AC-BLUE
* DCS-7060CX-32S-F	* DCS-7280CR2-60	* PWR-1011-AC-RED
* DCS-7060CX-32S-F (HwRev:13.20)	* DCS-7280CR2-60-F	* PWR-1011-DC
* DCS-7060CX-32S-R	* DCS-7280CR2A-30	* PWR-1011-DC-BLUE
* DCS-7060CX-32S-R (HwRev:13.20)	* DCS-7280CR2A-30-F	* PWR-1011-DC-RED
* DCS-7060CX2-32S	* DCS-7280CR2A-60	* PWR-1021-AC
* DCS-7060CX2-32S-F	* DCS-7280CR2A-60-F	* PWR-1021-AC-RED
* DCS-7060CX2-32S-R	* DCS-7280CR2K-30	* PWR-1100AC
* DCS-7060DX4-32	* DCS-7280CR2K-30-F	* PWR-1100AC-F
* DCS-7060DX4-32-F	* DCS-7280CR2K-60	* PWR-1100AC-R
* DCS-7060DX4-32S	* DCS-7280CR2K-60-F	* PWR-1511-AC
* DCS-7060DX4-32S-F	* DCS-7280CR2M-30	* PWR-1511-AC-BLUE
* DCS-7060DX4-32S-MON	* DCS-7280CR2M-30-F	* PWR-1511-AC-RED
* DCS-7060DX4-32S-MON2	* DCS-7280CR3-32D4	* PWR-1511-DC
* DCS-7060DX4-32S-PEM	* DCS-7280CR3-32D4-F	* PWR-1511-DC-BLUE
* DCS-7060DX4-32SON	* DCS-7280CR3-32D4-M-F	* PWR-1511-DC-RED
* DCS-7060DX4-32SON2	* DCS-7280CR3-32D4-M-R	* PWR-1600AC
* DCS-7060DX5-64	* DCS-7280CR3-32D4-R	* PWR-1600AC-F
* DCS-7060DX5-64-F	* DCS-7280CR3-32P4	* PWR-1611-AC
* DCS-7060DX5-64S	* DCS-7280CR3-32P4-F	* PWR-1611-AC-RED
* DCS-7060DX5-64S-F	* DCS-7280CR3-32P4-M-F	* PWR-1611-DC
* DCS-7060DX5-64S-R	* DCS-7280CR3-32P4-M-R	* PWR-1611-DC-RED
* DCS-7060PX4-32	* DCS-7280CR3-32P4-R	* PWR-1900-DC
* DCS-7060PX4-32-F	* DCS-7280CR3-36S	* PWR-1900-DC-F
* DCS-7060PX5-64E	* DCS-7280CR3-36S-F	* PWR-1900-DC-R
* DCS-7060PX5-64E-F	* DCS-7280CR3-36S-R	* PWR-1900AC
* DCS-7060SX2-48YC6	* DCS-7280CR3-96	* PWR-1900AC-F
* DCS-7060SX2-48YC6-F	* DCS-7280CR3-96-F	* PWR-1900AC-R

* DCS-7060SX2-48YC6-R	* DCS-7280CR3K-32D4	* PWR-2021-AC
* DCS-7130-16G3	* DCS-7280CR3K-32D4-F	* PWR-2021-AC-RED
* DCS-7130-16G3-F (HwRev:C1)	* DCS-7280CR3K-32D4-R	* PWR-2411-AC
* DCS-7130-16G3-R (HwRev:C1)	* DCS-7280CR3K-32D4A	* PWR-2411-AC-BLUE
* DCS-7130-16G3S	* DCS-7280CR3K-32D4A-F	* PWR-2411-AC-RED
* DCS-7130-16G3S-F (HwRev:D1)	* DCS-7280CR3K-32D4A-R	* PWR-2411-DC
* DCS-7130-16G3S-R (HwRev:D1)	* DCS-7280CR3K-32P4	* PWR-2411-DC-BLUE
* DCS-7130-48E	* DCS-7280CR3K-32P4-F	* PWR-2411-DC-RED
* DCS-7130-48E-F (HwRev:A4)	* DCS-7280CR3K-32P4-R	* PWR-3001-AC
* DCS-7130-48E-R (HwRev:A4)	* DCS-7280CR3K-32P4A	* PWR-3001-AC-BLUE
* DCS-7130-48EH	* DCS-7280CR3K-32P4A-F	* PWR-3001-AC-RED
* DCS-7130-48EH-F (HwRev:A4)	* DCS-7280CR3K-32P4A-R	* PWR-3001-DC
* DCS-7130-48EH-R (HwRev:A4)	* DCS-7280CR3K-36A	* PWR-3001-DC-BLUE
* DCS-7130-48EHS	* DCS-7280CR3K-36A-F	* PWR-3001-DC-RED
* DCS-7130-48EHS-F (HwRev:B1)	* DCS-7280CR3K-36A-R	* PWR-400-DC
* DCS-7130-48EHS-F (HwRev:B2)	* DCS-7280CR3K-36S	* PWR-400-DC-BLUE
* DCS-7130-48EHS-R (HwRev:B1)	* DCS-7280CR3K-36S-F	* PWR-400-DC-RED
* DCS-7130-48EHS-R (HwRev:B2)	* DCS-7280CR3K-36S-R	* PWR-400AC
* DCS-7130-48G3	* DCS-7280CR3K-96	* PWR-400AC-BLUE
* DCS-7130-48G3-F (HwRev:B1)	* DCS-7280CR3K-96-F	* PWR-400AC-RED
* DCS-7130-48G3-F (HwRev:B2)	* DCS-7280CR3MK-32D4	* PWR-460AC
* DCS-7130-48G3-R (HwRev:B1)	* DCS-7280CR3MK-32D4-F	* PWR-460AC-F
* DCS-7130-48G3-R (HwRev:B2)	* DCS-7280CR3MK-32D4-R	* PWR-460AC-R
* DCS-7130-48G3S	* DCS-7280CR3MK-32D4S	* PWR-460DC
* DCS-7130-48G3S-F (HwRev:C1)	* DCS-7280CR3MK-32D4S-F	* PWR-460DC-F
* DCS-7130-48G3S-R (HwRev:C1)	* DCS-7280CR3MK-32D4S-R	* PWR-460DC-R
* DCS-7130-48L	* DCS-7280CR3MK-32P4	* PWR-500-DC
* DCS-7130-48L-F (HwRev:A2)	* DCS-7280CR3MK-32P4-F	* PWR-500-DC-F
* DCS-7130-48L-F (HwRev:A3)	* DCS-7280CR3MK-32P4-R	* PWR-500-DC-R
* DCS-7130-48L-R (HwRev:A2)	* DCS-7280CR3MK-32P4S	* PWR-500AC
* DCS-7130-48L-R (HwRev:A3)	* DCS-7280CR3MK-32P4S-F	* PWR-500AC-F
* DCS-7130-48LA	* DCS-7280CR3MK-32P4S-R	* PWR-500AC-R
* DCS-7130-48LA-F (HwRev:A2)	* DCS-7280DR3-24	* PWR-501AC
* DCS-7130-48LA-F (HwRev:A3)	* DCS-7280DR3-24-F	* PWR-501AC-F
* DCS-7130-48LA-R (HwRev:A2)	* DCS-7280DR3-24-M-F	* PWR-501AC-R
* DCS-7130-48LA-R (HwRev:A3)	* DCS-7280DR3A-54	* PWR-511-AC
* DCS-7130-48LAS	* DCS-7280DR3A-54-F	* PWR-511-AC-BLUE
* DCS-7130-48LAS-F (HwRev:B2)	* DCS-7280DR3AK-54	* PWR-511-AC-RED
* DCS-7130-48LAS-F (HwRev:B3)	* DCS-7280DR3AK-54-F	* PWR-511-DC
* DCS-7130-48LAS-R (HwRev:B2)	* DCS-7280DR3AM-54	* PWR-511-DC-BLUE
* DCS-7130-48LAS-R (HwRev:B3)	* DCS-7280DR3AM-54-F	* PWR-511-DC-RED

* DCS-7130-48LB	* DCS-7280DR3K-24	* PWR-621-AC
* DCS-7130-48LB-F (HwRev:A2)	* DCS-7280DR3K-24-F	* PWR-621-AC-BLUE
* DCS-7130-48LB-F (HwRev:A3)	* DCS-7280PR3-24	* PWR-621-AC-RED
* DCS-7130-48LB-R (HwRev:A2)	* DCS-7280PR3-24-F	* PWR-721-DC
* DCS-7130-48LB-R (HwRev:A3)	* DCS-7280PR3-24-M-F	* PWR-721-DC-RED
* DCS-7130-48LBA	* DCS-7280PR3K-24	* PWR-745AC
* DCS-7130-48LBA-F (HwRev:A2)	* DCS-7280PR3K-24-F	* PWR-745AC-F
* DCS-7130-48LBA-F (HwRev:A3)	* DCS-7280QR-C36	* PWR-745AC-R
* DCS-7130-48LBA-R (HwRev:A2)	* DCS-7280QR-C36-F	* PWR-747AC
* DCS-7130-48LBA-R (HwRev:A3)	* DCS-7280QR-C36-M-F	* PWR-747AC-F
* DCS-7130-48LBAS	* DCS-7280QR-C36-M-R	* PWR-747AC-R
* DCS-7130-48LBAS-F (HwRev:B2)	* DCS-7280QR-C36-R	* PWR-750AC
* DCS-7130-48LBAS-F (HwRev:B3)	* DCS-7280QR-C72	* PWR-750AC-F
* DCS-7130-48LBAS-R (HwRev:B2)	* DCS-7280QR-C72-F	* PWR-750AC-R
* DCS-7130-48LBAS-R (HwRev:B3)	* DCS-7280QR-C72-M-F	* FAN-7000H
* DCS-7130-48LBS	* DCS-7280QR-C72-M-R	* FAN-7000H-F
* DCS-7130-48LBS-F (HwRev:B2)	* DCS-7280QR-C72-R	* FAN-7002
* DCS-7130-48LBS-F (HwRev:B3)	* DCS-7280QRA-C36S	* FAN-7002-F
* DCS-7130-48LBS-R (HwRev:B2)	* DCS-7280QRA-C36S-F	* FAN-7002-R
* DCS-7130-48LBS-R (HwRev:B3)	* DCS-7280QRA-C36S-M-F	* FAN-7002H
* DCS-7130-48LS	* DCS-7280QRA-C36S-M-R	* FAN-7002H-F
* DCS-7130-48LS-F (HwRev:B2)	* DCS-7280QRA-C36S-R	* FAN-7012MP
* DCS-7130-48LS-F (HwRev:B3)	* DCS-7280SR-48C6	* FAN-7012MP-RED

Modular

* FAN-7000H	* 7500R-8CFPX-LC	* CCS-750X-48ZP-LC
* FAN-7000H-F	* 7500R2-18CQ-LC	* CCS-750X-48ZXP-LC
* FAN-7002	* 7500R2-36CQ-LC	* CCS-755-CH
* FAN-7002-F	* 7500R2A-36CQ-LC	* CCS-755-X3-SC
* FAN-7002-R	* 7500R2AK-36CQ-LC	* CCS-758-CH
* FAN-7002H	* 7500R2AK-48YCQ-LC	* CCS-758-X3-SC
* FAN-7002H-F	* 7500R2AM-36CQ-LC	* DCS-7304
* FAN-7012MP	* 7500R2M-36CQ-LC	* DCS-7308
* FAN-7012MP-RED	* 7500R3-24D-LC	* DCS-7500-SUP2
* 7289	* 7500R3-24P-LC	* DCS-7500-SUP2-D
* 7289-SUP	* 7500R3-36CQ-LC	* DCS-7504
* 7289-SUP-D	* 7500R3-48Y4D-LC	* DCS-7504N
* 7289-SUP-S	* 7500R3K-36CQ-LC	* DCS-7508
* 7289-SUP-S-D	* 7500R3K-48Y4D-LC	* DCS-7508N
* 7289R3A-SC	* 7500RM-36CQ-LC	* DCS-7512N
* 7289R3AK-SC	* 7504R-FM	* DCS-7516-SUP2
* 7289R3AM-SC	* 7504R3-FM	* DCS-7516-SUP2-D
* 7300-SUP	* 7508R-FM	* DCS-7516N
* 7300-SUP-D	* 7508R3-FM	* DCS-7800-SUP
* 7300-SUP2	* 7512R-FM	* DCS-7800-SUP1A
* 7300-SUP2-D	* 7512R3-FM	* DCS-7800-SUP1S
* 7300X3-32C-LC	* 7516R-FM	* DCS-7804-CH
* 7300X3-48TC4-LC	* 7800R3-36D-LC	* DCS-7808-CH
* 7300X3-48YC4-LC	* 7800R3-36P-LC	* DCS-7812-CH
* 7304X3-FM	* 7800R3-48CQ-LC	* DCS-7816-CH
* 7304X3-FM2	* 7800R3-48CQ2-LC	* DCS-7816-SUP
* 7308X3-FM	* 7800R3-48CQM-LC	* DCS-7816-SUP1S
* 7308X3-FM2	* 7800R3-48CQM2-LC	* FAN-752M
* 7320X-32C-LC	* 7800R3-48CQMS-LC	* FAN-752M-RED
* 7324X-FM	* 7800R3A-36D-LC	* PWR-2421-HV
* 7328X-FM	* 7800R3A-36D2-LC	* PWR-2421-HV-RED
* 7358-16C	* 7800R3A-36DM-LC	* PWR-2700-DC
* 7358X4-SC	* 7800R3A-36DM2-LC	* PWR-2700-DC-F
* 7368	* 7800R3A-36P-LC	* PWR-2700-DC-R
* 7368-16C	* 7800R3A-36PM-LC	* PWR-2900AC
* 7368-16S	* 7800R3AK-36DM-LC	* PWR-3351-AC
* 7368-4D	* 7800R3AK-36DM2-LC	* PWR-3351-AC-RED
* 7368-4P	* 7800R3AK-36PM-LC	* PWR-3K-AC
* 7368-SUP	* 7800R3K-36DM-LC	* PWR-3K-AC-F

* 7368-SUP-D	* 7800R3K-48CQ-LC	* PWR-3K-AC-R
* 7368X4-SC	* 7800R3K-72Y-LC	* PWR-3K-DC
* 7388	* 7804R3-FM	* PWR-3K-DC-BLUE
* 7388-16CD	* 7808R3-FM	* PWR-3K-DC-RED
* 7388-16CD2	* 7808R3-FM2	* PWR-3KT-AC
* 7388-8D	* 7812R3-FM	* PWR-3KT-AC-BLUE
* 7388-SUP	* 7816R3-FM	* PWR-3KT-AC-RED
* 7388-SUP-D	* CCS-750-SUP100	* PWR-D1-3041-AC
* 7388X5-SC	* CCS-750-SUP25	* PWR-D1-3041-AC-BLUE
* 7500R-36CQ-LC	* CCS-750X-48THP-LC	* PWR-D4-3041-AC
* 7500R-36Q-LC	* CCS-750X-48TP-LC	* PWR-D4-3041-AC-BLUE
* 7500R-48S2CQ-LC		

Transceiver

* 1000BASE-BX10	* 25GBASE-SR	* OSFP-400G-LR4
* 1000BASE-BX10-D	* 25GBASE-XSR	* OSFP-400G-LR8
* 1000BASE-BX10-U	* 40GBASE-AOC	* OSFP-400G-PLR4
* 1000BASE-LX	* 40GBASE-CR4	* OSFP-400G-PLR4-S
* 1000BASE-SX	* 40GBASE-ER4	* OSFP-400G-SR8
* 1000BASE-T	* 40GBASE-LR4	* OSFP-400G-XDR4
* 100BASE-FX	* 40GBASE-LRL4	* OSFP-400G-XDR4-S
* 100GBASE-AOC	* 40GBASE-PLR4	* OSFP-400G-ZR
* 100GBASE-CR4	* 40GBASE-PLRL4	* OSFP-400G-ZRP
* 100GBASE-CWDM4	* 40GBASE-SR4	* OSFP-800G-2FR4
* 100GBASE-DR	* 40GBASE-SRBD	* OSFP-800G-2LR4
* 100GBASE-ERL4	* 40GBASE-UNIV	* OSFP-800G-2PLR4
* 100GBASE-FR	* 40GBASE-XSR4	* OSFP-800G-2XDR4
* 100GBASE-LR	* 50GBASE-CR	* OSFP-AMP-ZR
* 100GBASE-LR4	* ADPT-O-Q-100G	* QDD-200G-2LR4
* 100GBASE-LRL4	* AOC-D-D-400G	* QDD-400G-DR4
* 100GBASE-PLR4	* AOC-O-O-400G	* QDD-400G-FR4
* 100GBASE-PLRL4	* CAB-D-2Q-200G	* QDD-400G-LR4
* 100GBASE-PSM4	* CAB-D-2Q-400G	* QDD-400G-LR8
* 100GBASE-SR4	* CAB-D-4Q-200G	* QDD-400G-PLR4
* 100GBASE-SRBD	* CAB-D-4Q-400G	* QDD-400G-PLR4-S
* 100GBASE-SWDM4	* CAB-D-8S-200G	* QDD-400G-SR8
* 100GBASE-XCWM4	* CAB-D-8S-400G	* QDD-400G-SR8-C
* 100GBASE-XSR4	* CAB-D-D-400G	* QDD-400G-XDR4
* 100GBASE-ZR4	* CAB-O-2Q-200G	* QDD-400G-XDR4-S
* 100GE-DWDM2	* CAB-O-2Q-400G	* QDD-400G-ZR
* 10GBASE-AOC	* CAB-O-4Q-200G	* QDD-400G-ZRP
* 10GBASE-CR	* CAB-O-4Q-400G	* QDD-800G-2FR4
* 10GBASE-DWDM	* CAB-O-8S-200G	* QDD-800G-2LR4
* 10GBASE-DWDM-T	* CAB-O-8S-400G	* QDD-800G-2PLR4
* 10GBASE-DWDM-ZR	* CAB-O-O-400G	* QDD-800G-2XDR4
* 10GBASE-ER	* CAB-Q-2Q-100G	* QSFP-200G-AR4
* 10GBASE-ERBD	* CFP2-100G-DWDM-DCO	* QSFP-200G-FR4
* 10GBASE-ERBD-D	* CFP2-100GBASE-ER4	* QSFP-AMP-ZR
* 10GBASE-ERBD-U	* CFP2-100GBASE-LR4	* SFP-10G-MRA-T
* 10GBASE-ERLBD	* CFP2-100GBASE-XSR10	* SFP-10G-RA-1G-LX
* 10GBASE-ERLBD-D	* CFPX-100G-DWDM	* SFP-10G-RA-1G-SX
* 10GBASE-ERLBD-U	* CFPX-200G-DWDM	* SFP-10G-T
* 10GBASE-LR	* E-D800-D800	* SFP-10GRA-T

* 10GBASE-LRL	* E-O800-O800	* SFP-25G-LR-E
* 10GBASE-SR	* H-D400-4Q100	* SFP-25G-MR-LR
* 10GBASE-SRL	* H-O400-4Q100	* SFP-25G-MR-SR
* 10GBASE-ZR	* OSFP-200G-2LR4	* SFP-25G-MR-XSR
* 200GBASE-CR4	* OSFP-400-2FR4-LC	* SFP-25G-SR-E
* 200GBASE-SR4	* OSFP-400G-2FR4	* SFP-25GR-LR
* 25GBASE-AOC	* OSFP-400G-DR4	* SFP-25GR-SR
* 25GBASE-CR	* OSFP-400G-FR4	* SFP-25GR-XSR
* 25GBASE-LR		

Reference to MLAG ISSU Upgrade/Downgrade Table

<https://www.arista.com/en/support/mlag-portal>

Resolved Caveats in 4.29.6.1M

CVP

Network Provisioning

- On chassis-based switches, even if the command 'show hardware eeprom detail' shows a ZtpToken value is present, the embedded token ZTPaaS workflow may fail due to the location that EOS reads the values from.

For these platforms, devices can be onboarded to CVaaS via either the local bootstrap or USB ZTPaaS workflows. (920616)

Known Software Caveats in 4.29.6.1M

Accelerated System Update

- Performing SSU when /mnt/flash/persist/persistentRestartLog doesn't exist may result in extended traffic outage due to a watchdog reset that can occur. (158117)
- When performing ASU2 on an MLAG setup with VXLAN configuration, ASU2 can result in a fatal error resulting in a cold boot being performed resulting in large traffic loss. (245555)
Series Affected: DCS-7060X and DCS-7060X2 Series
- SSU with Splunk enabled may result in extended boot times, which could lead to protocol timeouts or switch cold rebooting. (343629)
- When performing SSU on an MLAG setup, SSU can cause peer links on port-channel to flap, resulting in large traffic loss. (405788)
- SSU can cause a momentary change in ECMP hashing, leading to extra traffic loss (738594)
- In certain EVPN deployments where the VTEPs share the same virtual IPV6 address, it is possible that the devices after undergoing hitless reboot are unable to claim the virtual interface address due to an otherwise avoidable IPV6 Duplicate Address Detection.

As a workaround to this issue, the SVIs with Duplicate Address Detection failure can be flapped to successfully register the virtual IPV6 address for the SVI. (832854)

- Electrical tuning values for QSFP-100G-CWDM4 transceivers may not be recovered and applied correctly after an ASU reload. This could result in link flaps or the link not coming up on affected interfaces. Output for 'show

interface transceiver tuning detail' command may be used to check the operational values of the tuning coefficients.

Affected interfaces can be recovered by issuing the interface configuration command

'transceiver diag simulate removed' followed by 'no transceiver diag simulate removed' (838186)

SKUs Affected: 100GBASE-CWDM4

- During an SSU, remote EVPN dynamic MACs might flap, resulting in temporary traffic loss. (838547)
- A system may reboot during an ASU if at least one CMIS transceiver (i.e. QSFP200, OSFP, QSFP-DD) is unresponsive. (845097)
- Performing SSU when OSPF router ID is not explicitly configured may result in extended traffic outage.

A workaround is to configure router ID before performing SSU (856714)

- Performing SSU with MSTP configuration may result in extended traffic outage (857363)
- QSFP optical transceivers may experience a link flap after configuring "reload fastboot". (866709)
- Connections using tunable SFP transceivers like the 10GBASE-DWDM-T and 10GBASE-DWDM-ZR may flap during 'reload fastboot'. There is no workaround, but after the flap(s), the link should be operational again. (878971)
SKUs Affected: 10GBASE-DWDM-T and 10GBASE-DWDM-ZR
- Performing SSU may lead to 100% CPU utilization for the StrataL3 agent post SSU.

Workaround is to restart StrataL3 agent by issuing "agent StrataL3 terminate". (917676)

SKUs Affected: DCS-7060CX-32S, DCS-7060CX2-32S, DCS-7060DX5-64S, DCS-7060SX2-48YC6, DCS-7260CX3-64, DCS-7260CX3-64E and DCS-7260CX3-64LQ

BGP EVPN L2/L3 VXLAN/MPLS

- A policy change (such as changing the color extended communities) which causes an L2 EVPN MAC-IP or IMET route to resolve over a different SR-TE policy or tunnel may cause a brief traffic loss. (521090)
- A remote VTEP reached through unnumbered BGP over IPv6 will be incorrectly displayed as a configuration error with "FAIL" state due to "No route" (525842)

- Some MAC addresses may get blacklisted when the BGP agent is restarted on a multihomed peer. The workaround is to clear the blacklist using the "clear bgp evpn host-flap" command. (538177)
- In an EVPN multi-homing deployment, designated forwarder election for a given Ethernet segment and EVPN instance pair is based on the complete set of VLANs within the EVPN instance. This differs from the standard in which only the instance VLANs that are in the Ethernet segment interface's allowed VLAN list are used in the election algorithm. This deviation may cause interoperability issues with other vendors.

As a workaround, a single VLAN-aware bundle EVPN instance may need to be split into multiple instances, such that for every instance and Ethernet segment either every instance VLAN is in the Ethernet segment interface's allowed VLAN list, or none are. (583126)

- On Multi-Domain EVPN VXLAN deployment using anycast gateways, administratively enabling and disabling the VXLAN interface may cause and ARP entry to be lost and trigger traffic lost. (632955)
- In an L2 EVPN MPLS configuration, disabling MAC address learning on a VLAN that is part of the EVPN bridging domain does not take effect. (677182)
 Series Affected: 7500R3, DCS-7280CR2, DCS-7280CR3, DCS-7280DR3, DCS-7280PR3, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 Series
- In a multi-domain EVPN VXLAN to MPLS deployment, when VLAN-aware bundled MAC-VRFs are configured, the MPLS ETID and VXLAN VNI must be the same value in order to support multi-domain MAC moves.

A workaround is to explicitly configure the MPLS side's VLAN-aware bundled MAC-VRFs using the follow configuration.

```

vlan-aware-bundle vlan_38_39
    ! mpls etid for vlan 38 is 38
    vlan 38
    ! mpls etid for vlan 39 is 1039
    vlan 39 etid 1039 (709750)

```

- When configuring "set_nexthop_self()" using RCF in the EVPN outbound context, the BGP agent can restart. (717957)
- In an EVPN VXLAN deployment using gateway overlay index, flapping configuration for "route import overlay-index gateway" may cause vias for overlay-index routes to become empty, as seen in the "show ip route <prefix>" command. This situation can be avoided by waiting for ip route convergence between the "no route import overlay-index gateway" and "route import overlay-index gateway" commands.

If the situation is encountered, as a workaround the routes can be reset to the functioning state by resetting the BGP session with "clear ip bgp". (723279)

- In an EVPN VXLAN deployment, the BGP process can leak memory when the VLAN to VNI mapping is continuously provisioned and unprovisioned. The amount of memory leaked does not scale with the number of routes but instead with the number of MAC-VRFs. (774066)
- In a EVPN VXLAN deployment with L2 ARP or ND proxy enabled ("arp learning bridged" for IPv4 hosts and/or "nd learning bridged" for IPv6 hosts), when a GARP is generated by a host to update its MAC in an ARP or ND binding, the multi-homed PEs may advertise stale and conflicting MAC to ARP/ND binding. (838924)
- If a DCI gateway changed from connecting VXLAN and VXLAN domains to connecting VXLAN and MPLS domains, EVPN MAC-IP and IMET routes sent by this gateway may be missing on the remote DCI gateways or route reflectors.

Use "clear ip bgp <neighbor> soft in" on the remote gateways or route reflectors to recover the missing EVPN routes. (841750)

- On an EVPN router, the system might run out of dynamic VLANs available, in the event of a VXLAN VRF VNI mapping configurations churn.

To workaround the problem, restart the BGP agent. (846021)

- In a EVPN VXLAN deployment where L2 ARP and ND proxy is enabled, executing the "clear arp|nd bridged" or the "no arp|nd learning bridged" command might not clear all the locally learned L2 ARP entries in the "show arp bridged" or "show ipv6 neighbors bridged" output. (882754)
- An EVPN L3 gateway may not install IP-Prefix routes from the gateways of the remote domain.
A work around is to restart the BGP agent. (885900)
- In MAC-VRFs that are configured to redistribute the router MAC or system MAC for associated SVIs, re-configuring the VLAN to VNI map can lead to the router MAC and/or system MAC routes to not be properly redistributed. Specifically, if a VNI that was previously mapped to one VLAN is reconfigured to map to a different VLAN, it's possible for the router MAC and system MAC routes for the new mapping to be missing.

Flapping the associated SVI will cause the routes to be advertised as expected. (886496)

vEOS Router / CloudEOS

- The OpenFlow agent fails to handle packet-in messages (145001)

- When several IPsec tunnels are up and traffic is passing through them, removal and addition of a IPsec license may cause the the Ipsec agent to restart.

The workaround is to shutdown all the interfaces before the IPsec license is removed and reapplied and then do a "no shutdown" on the tunnel interfaces. (248213)

- A reboot due to a kernel panic can happen during first 10 minutes of vEOS bootup on Microsoft Azure. The kernel panic occurs because of a task blocking without being scheduled for more than 300 seconds in "D" state. (249618)
- Interface may be allowed to be configured with un-supported speed of NIC using "speed <speed>", which will cause an unexpected PhyEthtool agent restart. Do not use "speed <speed>" command to configure interface speed that NIC does not support. (264395)
- When running vEOS devices in Sfe mode with interfaces having an MTU of 9214, users may encounter a scenario where not all routes are propagated through the network. The workaround is to have interfaces use an MTU of 9202 or less. (289886)
- On vEOS Router in Sfe/DPDK mode with many GRE/IPSec tunnels, changes to configuration may cause a long time to be applied to the datapath. (349552)
- When loading a vEOS instance in an Azure cloud, it is possible for the device to experience a kernel panic just after bootup, causing the device to reboot a second time. (354411)
- In an EVPN IRB deployment on vEOS, changing the VXLAN source-interface interface may result in EVPN ip-prefix routes not getting installed into one or more VRFs for one or more remote VTEPs. (369816)
- If DNS resolution for GCP instance fails, SFE agent may reload unexpectedly. (403456)
- CloudEOS supports upto 600K BGP routes with 8G of system memory; If 800K BGP routes are programmed, out-of-memory condition will kill random processes.

Out-of-memory issue can be mitigated by increasing the system RAM to 12G to program 800K BGP routes. (478214)

- When an IPsec connection is established with a primary IP address of an interface and a secondary IP address is configured on an interface and that secondary IP address is used to establish an IPsec connection, the IPsec connection with the secondary IP address doesn't come up.
The workaround is to restart the IPsec agent. (572808)

- When running CloudEOS on ESXI hypervisor, while using Intel Ethernet Controller X710 NIC in SRIOV mode, the TX traffic out of the interface, may stop working after either the router reload and/or SFE agent (Software Forwarding Engine) restart for any reason.
The workround is to set "trust-mode" ON and "spoof-check" OFF for the SRIOV interface on the Vmware hypervisor.
Another possible workaroud is to log into the ESXI VCenter console and change the SRIOV port's assigned VLAN to any other VLAN and then revert it back to the original VLAN. (647410)
- The /var/log/ filesystem on AWS CloudEOS VMs may get filled up with awslogs.*backup files. The workaroud is to delete the /etc/cron.d/awslogs_log_rotate file. This workaroud is better applied using an on-boot event handler that deletes this file on every reboot of the VM. (713562)
- CloudHA uses older Boto SDK version than some of the AWS regions allow. CloudHA may not work in such AWS regions.

The workaroud is to create a boto config file as indicated below and follow it with shut/no shut action on the CloudHA feature. This file needs to be created in the designated path upon every restart. An event-handler can be used to achieve it.

Boto config file content with the exact path:

```
-----
[ec2-user@localhost ~]$ cat /etc/boto.cfg
[Boto]
use_endpoint_heuristics = True
-----
```

Event handler configuration:

```
-----
event-handler AWS
    trigger on-boot
    action bash sudo echo -e "[Boto]\nuse_endpoint_heuristics = True" > /etc/
boto.cfg
    delay 0
----- (758301)
```

- When auto-discovery of path MTU is enabled in a DPS path-group, and an IPsec profile with tunnel mode is configured in the same path-group, Sfe agent may restart unexpectedly.
Workaroud is to avoid such a configuration by configuring a static path MTU value in the path-group or changing the IPsec profile from Tunnel mode to Transport mode. (811452)

CVP

Network Provisioning

- ZeroTouch agent may keep restarting continuously during Zero Touch Provisioning (ZTP) if the DHCP option 67 (boot file name) is a hostname that needs to be resolved as opposed to an IPv4 or IPv6 address. This will result in ZTP getting stuck. The workaround is to use an IPv4 or IPv6 address in the boot file URL. (718251)
- During Zero Touch Provisioning(ZTP), repeated failures to execute a bootstrap script that forks other scripts may result in the switch running out of memory.
The workaround is to rewrite the bootstrap script to either avoid forking or properly cleanup forked child processes when the script exits, and then reloading the switch to restart ZTP. (820844)

CVX

- If a VmTracer session configures all the available VLANs on the switch, then the switch may errdisable the routed ports by freeing up the internal VLANs. (139294)
- On all switch series configured as a MLAG pair where CVX is used as the VXLAN control plane, after a MAC address move between switches the MAC address might be advertised to the CVX server by two different VTEPS. This can cause packets destined to the MAC address to be blackholed. The workaround is to clear the mac address on both the advertising VTEPS so that the mac address will be relearned properly. (158130)
- The MacroSegmentation Service (MSS), working with L2 transparent firewalls, requires the firewall interface be statically identified on CVX via configuration commands (as opposed to using LLDP to detect them dynamically). This prevents issues such as traffic loss in the event that a member port of a aggregated link goes down. (275384)
- Macro-segmentation Service (MSS) might stop intercepting traffic when the last member of the Near service LAG goes down.

To work around this failure, use the interface mapping from the MSS dynamic device-set configuration to statically map service device interfaces to their corresponding switch port channels. (275656)

- When a security policy on a CheckPoint firewall includes unsupported services, the Macro-Segmentation Service (MSS) intercepts traffic for all services for the specified end-points in the policy. (433067)
- When the Policy Control Service is enabled on CVX, under certain circumstances such as re-numbering IP on a host connected to a tagged port,

the groups received from NSX controller may not be resolved into IP addresses. This may result in incorrect programming of ACLs on the switches.

Workaround is to disable and re-enable 'service pcs' on CVX. (538537)

- When MSS is configured in a VXLAN routing deployment with VXLAN controller service (VCS), an MLAG ISSU upgrade from a pre 4.22.1 release to a post 4.22.1 will be hitful and MLAG roles will not be re-established until both MLAG peers have been upgraded.

The workaround is to disable MSS from the VCS prior to the ISSU upgrade. Note that disabling MSS will cause ARP entries on the switch to be relearned and is disruptive. (694514)

- In a VCS deployment with MLAG, a rapid MAC move may result in a MAC being learned locally on multiple VTEPs. The workaround is to run "clear mac address-table dynamic vlan VLAN" on the MLAG secondary which should not have the MAC learned locally. (702518)
- The MCS agent might restart after multiple addSenders POST requests for the same senders that exist in the setup. (767011)
- In a VXLAN deployment using VXLAN Controller Service (VCS) to distribute MAC addresses, the MAC address of an orphan host behind a secondary MLAG switch may not get published to VCS. (770762)
- If MCS client agent restarts, it could affect multicast traffic flows for the secondary cvx. (783230)
- The Media Control Service (MCS) can oversubscribe an interface when it configures a multicast flow with a reservation percentage of 0 on the interface while other multicast flows are configured on the interface. (876232)
- Media Control Service (MCS) does not show multicast flows in response to GET mcs/flows-active API When ingress and egress interface of the multicast flow are in the same VLAN and on the same switch (876240)
- Media Control Service (MCS) may fail to process HTTP requests for /mcs/senders and /mcs/multicast/senders API when upgrading CVX with incompatible EOS versions that MCS does not support.

As a workaround clear multicast senders by using HTTP /mcs/clearRoutes API (882295)

Enterprise routing

- When IPsec is configured on DPS path groups, applying the same configuration using "configure replace" CLI can cause DPS paths to flap. (814121)

General

- When a policy map of type PBR is attached to an interface that is either in down state or a switchport (or both), and if the policy map is too large to fit into available hardware resources, the CLI will not report the error and roll back the configuration. Later, when the interface becomes an operational routed interface, the policy map will not be programmed into hardware. However, "show policy-map type pbr" shows the policy applied to the interface.

To fix the discrepancy, remove extra rules from the policy map to make it fit into the hardware. (89931)

- The OpenFlow agent fails to return counters associated with a flow, when queried by the OpenFlow controller using protocol version 1.0 (132812)
- On switches using the tg3 driver for the management interface a kernel panic can happen resulting in a switch reload. This is due to a Single Event Upset (SEU) affecting the management interface ethernet controller, and is a rare occurrence. (136944)

Series Affected: 7368 Series

SKUs Affected: DCS-7020SR-24C2-F, DCS-7020SR-24C2-R, DCS-7020SR-32C2-F, DCS-7020TR-48-F, DCS-7020TR-48-R, DCS-7020TRA-48-F, DCS-7020TRA-48-R, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050SX3-48YC12-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-F, DCS-7060SX2-48YC6-R, DCS-7170-32C-F, DCS-7170-32C-M-F, DCS-7170-32C-M-R, DCS-7170-32C-R, DCS-7170-32CD-F, DCS-7170-32CD-R, DCS-7170-64C-F, DCS-7170-64C-M-F, DCS-7170-64C-M-R, DCS-7170-64C-R, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7280CR2A-30-F, DCS-7280CR2K-30-F, DCS-7280CR2M-30-F, DCS-7280QR-C36-F, DCS-7280QR-C36-M-F, DCS-7280QR-C36-M-R, DCS-7280QR-C36-R, DCS-7280QRA-C36S-F, DCS-7280QRA-C36S-M-F, DCS-7280QRA-C36S-M-R, DCS-7280QRA-C36S-R, DCS-7280SR-48C6-F, DCS-7280SR-48C6-M-F, DCS-7280SR-48C6-M-R, DCS-7280SR-48C6-R, DCS-7280SR2-48YC6-F, DCS-7280SR2-48YC6-M-F, DCS-7280SR2-48YC6-M-R, DCS-7280SR2-48YC6-R, DCS-7280SR2A-48YC6-F, DCS-7280SR2A-48YC6-M-F, DCS-7280SR2A-48YC6-M-R, DCS-7280SR2A-48YC6-R, DCS-7280SR2K-48C6-M-F, DCS-7280SR2K-48C6-M-R, DCS-7280SRA-48C6-F, DCS-7280SRA-48C6-M-F, DCS-7280SRA-48C6-M-R, DCS-7280SRA-48C6-R, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F, DCS-7280SRM-40CX2-R, DCS-7280TR-48C6-F, DCS-7280TR-48C6-M-F, DCS-7280TR-48C6-M-R, DCS-7280TR-48C6-R, DCS-7280TRA-48C6-F, DCS-7280TRA-48C6-M-F, DCS-7280TRA-48C6-M-R,

DCS-7280TRA-48C6-R, DCS-7304, DCS-7308, DCS-7504N, DCS-7508N, DCS-7512N and DCS-7516N

- When configuration of a service-policy of type pdp on an interface fails due to lack of hardware resources, the interface may not be protected by any PDP policy. The workaround is to remove the configuration of the PDP service-policy and to configure the PDP service-policy default-pdp-policy on the interface. (216365)
- The CLI command "route-target import auto <ASNumber>" for VLAN VRF does not work and hence should not be used. Alternatively, the suggested configuration is "route-target import auto". In this case the AS Number is picked up from the BGP configuration.

Example:

```
router bgp 301
  vlan 300
    # This does not work as AS Number for generation of route target is
    explicitly
    # configured
    route-target import auto 302
```

Suggested Usage

```
router bgp 301
  vlan 300
    # In this case 301 is used as the AS number for generation of route
    target
    route-target import auto (255062)
```

- The MacroSegmentation Service (MSS) feature doesn't work in conjunction with the VARP VTEP IP configuration required for VXLAN routing where a subset of VTEPs are L2 VTEPs. (263532)
- When a MLAG peer (configured as a service VTEP for MSS) reloads, MSS re-installs intercept flows which may cause traffic to drop for a short duration. (349254)
- If multiple front-panel ethernet interfaces are configured as reflector interfaces for the same flow of traffic to be reflected, loops can be formed and hosts can flap from one interface to another when the reflector configuration is subsequently changed. The workaround is to only configure one reflector interface to handle traffic reflection. (360869)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- When configuring VXLANs from the EOS CLI, "ip address virtual" is supported and mapped to YANG, but "ip virtual-router address" is not supported. (372193)

- If the number of PTP vlan's enabled on a switch using 'ptp vlan' command exceeds 65536, PTP agent will become unresponsive and eventually restart. Reducing the VLANs below the limit will avoid getting in this state. (408199)
- When in transparent two-step mode PTP packet sequenceId collision may cause transient jump in offsetFromMaster. Workaround is to use one-step transparent clock mode. (408202)
- DirectFlow 'action output nexthop <ip addr>' ignores TTL of incoming packet and forwards the packet even if the TTL has expired. (417994)
- In rare conditions, the kernel may panic with the message containing "BUG at ../mm/slub.c:3334" (476225)
- On system reload, the kernel may crash with the "PANIC: double fault" message. (493516)
- Management interfaces may flap with kernel message "transmit timed out, resetting" and should come back up automatically. On rare occasions an interface will remain down after such flap until the interface PCI device is reset or the system reloaded. (500322)
- Modular switch may restart in case of Fabric or Linecard failure. (554503)
- Radsec does not honour CRLs configured in the corresponding SSL profiles while establishing TLS connections with servers. (557935)
- Local (admin) user is allowed to login with empty password when RADIUS over TLS is enabled (569426)
- "reload peer" configured with a delay such as "reload peer in 1 force" results in both supervisors being reloaded after the specified delay. (621572)
- MLAG agent may be stuck in a "Connecting" negotiation status under certain scenarios on the peer MLAG (e.g., incompatible EOS images). Rebooting the peer switch and restart local MLAG agent may resolve this problem. (630735)
- The OpenConfig & Octa agents can restart unexpectedly if a subscriber stops consuming data (or slows consuming data) for a long period of time. (647535)
- The Octa or OpenConfig agent may restart if there are a large number of sample subscriptions, or if sample subscriptions are made to a large number of paths. (648223)
- ConnectivityMonitor may restart unexpectedly if more than 250 http probes are configured. (672269)

- If number of probes exceed 100 with a minimum of 5 second interval configured, the ConnectivityMonitor agent can restart unexpectedly, causing all of the probes to go down momentarily. (685392)
- BOM change on PCA-01601-03 and change of riser pol image from SFT-01392-01 to SFT-01392-02 resolves issue. (732299)
SKUs Affected: DCS-7280DR3A-54-F, DCS-7280DR3AK-54 and DCS-7280DR3AM-54-F
- If all sampled flow-tracker flows time out at the same time, SftAgent deletes them very slowly and consumes a lot of cpu until it is finished. (773943)
- When an event handler has the following 3 conditions:
 - The "trigger on-config disabled" command is configured
 - A delay greater than 0 seconds
 - The first event is triggered sooner than the delay time since the handler is configured

In this case, the handler's action is triggered after the delay seconds since the time the handler is configured, not since the actual event is detected.

Workaround:

Configure the delay to be 0 seconds and add a "sleep" command to the action bash command (778181)

- On reload, it is possible for the VXLAN datapath learning setting to become out of sync with its state prior to the reload. This could lead to hosts and VTEPs not being learned from dataplane traffic when they're expected to be, or being learned from dataplane traffic when they're not expected to be.

As a workaround, flapping the VTI should re-synchronize the setting and fix the issue (789611)

- Removing an interface and reassigning its IP addresses to another interface is not possible via OpenConfig in a single transaction. A workaround is to split the transaction into two or use a mixed-mode (CLI + YANG tree) OpenConfig transaction (791299)

- Repeatedly sampling/subscribing to certain OpenConfig paths causes the AirStream agent to leak memory.

The following OpenConfig paths are affected:

```
/stp/global
/stp/mstp/mst-instances/mst-instance/state
/stp/interfaces/interface/state
/qos/scheduler-policies
/qos/queues
/qos/queue-management-profiles
/qos/interfaces
/qos/forwarding-groups
```

```
/qos/classifiers
/arista/pseudowire/status
/arista/ip/nat/status
/arista/redundancy/management-active
/components/component/state
/components/component/transceiver/thresholds/threshold
/components/component/transceiver/physical-channels/channel
/components/component/transceiver/state
/components/component/optical-channel/state
/terminal-device/logical-channels/channel
/terminal-device/logical-channels/channel/logical-channel-assignments/
assignment
/arista/varp/status
```

Workaround is to restart AirStream agent and sample paths at most every 10 or 30 seconds. (801020)

- In an EVPN VXLAN multihoming setup, administratively shutting down and re-enabling the VXLAN source loopback interface on a remote VTEP can result in MAC-IP routes of some hosts, received from the multihoming VTEPs, not being installed causing the traffic to fail to those hosts, in the absence of other covering routes. (804491)
- An EOS agent may restart unexpectedly with a log message such as:
%FWK-3-SOCKET_CLOSE_LOCAL: Closing connection to Sysdb (pid:3033) at tbl://sysdb/+n (Received an unexpected Tac::Exception)

Additionally, the agent log should contain a message such as:
Exception thrown is Tac::AttrLogEntityException("not found eid 0xd675, sAdapt: NULL")

There is no workaround but the agent is expected to recover normally after a restart in most cases. (807110)

- Units may be unable to complete exiting from maintenance mode if entering maintenance mode is cancelled before reaching Under Maintenance state. (813375)
- Using an event-handler config with multi-line action will fail in running the action on the newly active supervisor in an SSO setup when the switchover happen.

Workaround: Try to convert the the multi-line action to a single line if possible (821674)

- Multiple restarts of platform forwarding agents, e.g. Strata-FixedSystem, Strata-Linecard, SandFap-FixedSystem, SandFapNi-Linecard, etc., can result in an out of memory condition.

To recover from this situation, power cycle the switch. (840871)

- OpenConfig/Octa agent may restart while a large number of routes are added using EOS SDK (843671)
- If a PTP boundary clock receives its own messages on a different port, it will fail to sync. (852891)
- With CPU traffic policy configured, a number of L3 subinterfaces flapping simultaneously (e.g. from their parent interface being shut) may cause the Acl agent to unexpectedly restart. This may cause control-plane protocols with aggressive timers (e.g. BFD) to also reset. (853905)
- Multi threaded EosSdk application may crash due to non-thread safe access to data structures (860836)
- When editing traffic-policy match rules, if one tries to edit a rule for a different match type it is possible to configure an unsupported address family source/destination prefix. For example, if you have 'match myRule ipv4' and try to create 'match myRule ipv6' a warning is thrown, but you are able to add v6 prefixes to the rule which incorrectly lead to the v6 prefixes being added to the 'match myRule ipv4'. The result is a continuous SandAegis crash. To stop the crashes, either delete the rule or remove the v6 prefixes from the v4 match rule. (862051)
- If we have a LAG spanning multiple chips and an interface traffic-policy applied to the LAG and to interfaces on each chip the LAG has members, removing all the member for a given chip from the LAG causes SandAegis to restart. (867018)
- On VXLAN routing setup using "ipv6 address virtual", neighbor entries for front panel hosts for an entire VRF may not get learned right after switch reload.

The workaround is to restart the VxlanSwFwd agent. (873888)

- Any third-party software running iptables (eg: dockerd) can race with AclAgent corrupting iptables rules in the kernel and can lead to traffic loss. This can be detected by the %ACL-3-KERNEL_TABLE_MANAGEMENT_FAILURE syslog.
Restarting Acl agent without any other background iptables commands being run can restore the EOS specific iptables state. (875512)
- In the MDNS gateway solution, if there are service with the same record name on different response links, responses from the MDNS gateway for that service may include additional records from services from different links. (875716)

- When a new SBFD session is initialized, before the session comes up, there may be a false "Init" to "Down" state transition within a few seconds of session bringup. Configure a SBFD interval greater than one second to work around the issue. (896821)
- A communication problem with the TPM may prevent the user from installing an extension. A reload of the system will usually fix the underlying issue with the TPM, allowing extensions to be installed again. (904810)
- Configuring an untagged L3 subinterface, I', under some parent interface I, in a non-default VRF will prevent sampling of packets received on any subinterface under the given parent interface, I, for features like sFlow and IPFIX. (912274)
- If the /mnt/flash/debug/reload_history file is very large or becomes corrupted, Sysdb may restart repeatedly.

Delete /mnt/flash/debug/reload_history to stop this behaviour. Doing this will erase the device's reload history. (913793)

- The AVA sensor may not reconnect to the nucleus upon connection failure. Restart the FlowWatcher agent. (920688)

IPSEC

- Under some circumstances IPsec rekey time may show up as 'N/A'. (561575)
- In Jericho2 IPsec tunnel mode, both peers should either enable or disable replay protection. Asymmetric configuration will lead to complete packet loss over the tunnel. (605866)
- On a topology configured with IPsec AutoVPN, doing a shut/no shut several times on interfaces can cause the IPsec agent to restart unexpectedly. (634820)
- Few packet drops are observed intermittently during IPsec Rekey (847831)
- After performing a configuration replace, TunnelIntf agent may restart with a large scale of tunnels if the hostname is used as the tunnel destination. To resolve the agent restart issue for a large scale of tunnels, the IP address can be used as the tunnel destination. (868618)
 SKUs Affected: DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

Layer 1

- Babbage Firmware v1.0.11: Links may not come up with 40G mode due to firmware data structures becoming invalid. The workaround is to restart training by flapping the TH3 side. This will be fixed in an upcoming firmware version. UPDATE 10/11: EOS mid-Oct firmware bundle expected to fix this. UPDATE 10/18/2018: EOS mid-Oct EFT (EFT 6) firmware does not fix this. The firmware is expected in the next EOS EFT. UPDATE 10/31/2018: early-Nov EFT with Firmware 1.13.0 is expected to resolve this. UPDATE 11/07/2018: Babbage Firmware v1.13.0 resolves this (292147)
- The forwarding agent may unexpectedly restart after a sequence of speed changes involving speeds of 10G, 25G, 40G, 50G and 100G. (358317)
Series Affected: DCS-7060X, DCS-7060X2, DCS-7260CX, DCS-7260QX, DCS-7260X and DCS-7260X3 Series
- When a QSFP port is set to 4x25G rate, inserting a 40G QSFP transceiver can trigger a single or continuous forwarding agent restart.
A workaround is to configure the port for 40G or 4x10G rate before inserting the transceiver. (405185)
Series Affected: 7300X3, DCS-7050X3 and DCS-7260CX3 Series
- Sequences of speed changes involving speeds of 10G, 25G, 40G, 50G and 100G may result in a Strata forwarding agent restart. (456835)
Series Affected: DCS-7260CX3 Series
- Configuring speed, forward error-correction or speed-groups on a large number of interfaces at once may cause the forwarding agent to restart and all interfaces to flap. The workaround is to configure interfaces in smaller batches. (494284)
Series Affected: 7368, 7388 and DCS-7060X4 Series
- On QSFP-DD/OSFP port EthernetX, if EthernetX/1 and EthernetX/3 are configured for two lane auto-negotiation, EthernetX/1 linking down can cause EthernetX/3 to link down. The above is true for EthernetX/5 and EthernetX/7 as well.

One workaround is to remove auto-negotiation configuration on EthernetX/3 and reapply it.

Another workaround is to perform "platform trident reset" if it is available. The forwarding agent will be restarted and links will flap but all ports will be in a good state.

Another workaround is to save the configuration to startup-config and restart. (571883)

Series Affected: DCS-7060X4 Series

SKUs Affected: 7500R3-24D-LC, 7500R3-24P-LC, 7500R3-48Y4D-LC, 7500R3K-48Y4D-LC, 7800R3-36P-LC, 7800R3-48CQ2-LC, 7800R3-48CQMS-LC, 7800R3K-72Y-LC, DCS-7280CR3K-32P4-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F,

DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F, DCS-7280PR3K-24-F, DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R, DCS-7280SR3K-48YC8-F and DCS-7280SR3K-48YC8-R

- 10GBASE-T SFP adapters will always advertise 100M/1G/10G link modes during BASE-T auto-negotiation, regardless of the CLI configuration. With the default speed configuration on an interface, they might fail to link up or might report linkdown while the peer has a linkup.

Workaround is to configure both peers with the desired linkup speed explicitly. For Arista switches, ``speed auto SPEED`` CLI command can be used. (675238)

Series Affected: CCS-710P, CCS-720DP and CCS-720DT Series

SKUs Affected: CCS-720DF-48Y and CCS-720DP-48S

- BASE-T ports might link up but blackhole traffic when they are configured to not use BASE-T auto-negotiation and force speed to 100M using ``speed 100full`` CLI command.

Workaround is to run ``shutdown``, wait for a minute, and then run ``no shutdown`` on the affected interfaces. (681902)

Series Affected: CCS-710P and DCS-7010TX Series

SKUs Affected: CCS-720DP-24S, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-48S, CCS-720DP-48S-2F, CCS-720DP-48S-F and CCS-722XPM-48Y4

- The /1 interface of a QSFP28 port may get stuck in multi-lane mode when changing speed from ``speed 100g`` to ``speed 10g`` after inserting a 40G transceiver. Some of the other interfaces for the port may remain inactive rather than coming up as expected.

Workaround is to configure the /1 interface to 40g and then 10g again, or configure 10g before inserting the transceiver. (694587)

Series Affected: DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7260CX, DCS-7260CX3 and DCS-7260QX Series

- Interfaces with BASE-T SFP transceivers and link-debounce configuration may experience a link down debounce period double than the configured period. The workaround is to configure "phy link detection aggressive" on affected interfaces (this may lead to an interface flap at the time of applying this configuration). (702920)
- Inserting a 1Gbps SFP transceiver or 10Gbps BASE-T SFP transceiver configured at 1G into a QSFP port using a QSFP-SFP Adapter (QSA) causes 1 logical port to be allocated. If the interface is configured with "speed forced 1000full" or "speed auto 1000full", transceiver removal causes 3 additional logical ports to be allocated. These logical ports may be reallocated to other active interfaces and cause them to become inactive.

The workaround is to configure "default speed" prior to the removal of affected transceivers. (706063)

Series Affected: DCS-7050SX3 and DCS-7060CX2 Series

- Inserting a 40Gbps transceiver into a port with a default speed configuration may cause the speed to be auto-adjusted to 4x10G. This may lead to reallocation of 3 logical ports and inactivation of other interfaces.

The workaround is to explicitly configure "speed forced 40gfull" on affected interfaces prior to inserting the transceiver, or configure "transceiver qsfp default-mode 4x10G" globally. (706067)

Series Affected: DCS-7050SX3, DCS-7050TX3 and DCS-7260X3 Series

- When inserting a linecard into a running system there is a chance that the forwarding agent on the switchcard may unexpectedly restart. The same issue may occur on the linecards' forwarding agent if and when a second switchcard is inserted. After restarting all traffic should forward normally.

During cold boot the forwarding agent on one or more of the cards may also restart, but will stabilize before starting to pass traffic. (747464)

SKUs Affected: CCS-755-CH and CCS-758-CH

- L1 Profiles configuration is not supported with OpenConfig. (789141)
- Executing "shutdown" interface configuration command after SSU may result in an unexpected forwarding agent restart and a brief traffic outage on all ports. (792791)

SKUs Affected: CCS-720DP-48S, CCS-722XPM-48Y4 and DCS-7010TX-48

- Configuring speed-group followed by configuring "speed auto 25gfull" may cause the link to remain down.
To recover, execute "platform trident reset". This will cause all links to flap and a brief traffic outage to occur.
To workaround the issue, issue the speed configuration first before issuing the speed-group configuration. (795944)

Series Affected: DCS-7060SX2 Series

- ZTP may not complete if all of the ports within a speed-group on the system are populated with transceivers that do not support the default speed-group configuration. If these interfaces are the interfaces needed to connect to the DHCP/configuration server this will result in ZTP failing to complete.

Possible workarounds are using a different interface for DHCP/server connectivity, temporarily removing one of the unneeded transceivers from the group until ZTP succeeds, or changing the wiring on the system so that the speed-group has at least one transceiver that supports the default configuration. (797607)

- In some very rare cases the XcvrAgent will fail to program BASE-T transceivers, leading to a failure to bring the link up. The workaround is to shut/no shut the interface. (798262)
 SKUs Affected: DCS-7130-16G3-F (HwRev:C1), DCS-7130-16G3-R (HwRev:C1), DCS-7130-16G3S-F (HwRev:D1), DCS-7130-16G3S-R (HwRev:D1), DCS-7130-48E-F (HwRev:A4), DCS-7130-48E-R (HwRev:A4), DCS-7130-48EH-F (HwRev:A4), DCS-7130-48EH-R (HwRev:A4), DCS-7130-48EHS-F (HwRev:B1), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48EHS-R (HwRev:B1), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48G3-F (HwRev:B1), DCS-7130-48G3-F (HwRev:B2), DCS-7130-48G3-R (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48G3S-F (HwRev:C1), DCS-7130-48G3S-R (HwRev:C1), DCS-7130-48L-F (HwRev:A2), DCS-7130-48L-F (HwRev:A3), DCS-7130-48L-R (HwRev:A2), DCS-7130-48L-R (HwRev:A3), DCS-7130-48LA-F (HwRev:A2), DCS-7130-48LA-F (HwRev:A3), DCS-7130-48LA-R (HwRev:A2), DCS-7130-48LA-R (HwRev:A3), DCS-7130-48LAS-F (HwRev:B2), DCS-7130-48LAS-F (HwRev:B3), DCS-7130-48LAS-R (HwRev:B2), DCS-7130-48LAS-R (HwRev:B3), DCS-7130-48LB-F (HwRev:A2), DCS-7130-48LB-F (HwRev:A3), DCS-7130-48LB-R (HwRev:A2), DCS-7130-48LB-R (HwRev:A3), DCS-7130-48LBA-F (HwRev:A2), DCS-7130-48LBA-F (HwRev:A3), DCS-7130-48LBA-R (HwRev:A2), DCS-7130-48LBA-R (HwRev:A3), DCS-7130-48LBAS-F (HwRev:B2), DCS-7130-48LBAS-F (HwRev:B3), DCS-7130-48LBAS-R (HwRev:B2), DCS-7130-48LBAS-R (HwRev:B3), DCS-7130-48LBS-F (HwRev:B2), DCS-7130-48LBS-F (HwRev:B3), DCS-7130-48LBS-R (HwRev:B2), DCS-7130-48LBS-R (HwRev:B3), DCS-7130-48LS-F (HwRev:B2), DCS-7130-48LS-F (HwRev:B3), DCS-7130-48LS-R (HwRev:B2), DCS-7130-48LS-R (HwRev:B3), DCS-7130-96-F (HwRev:A7), DCS-7130-96-R (HwRev:A7), DCS-7130-96-R (HwRev:A8), DCS-7130-96L-F (HwRev:A2), DCS-7130-96L-F (HwRev:A3), DCS-7130-96L-R (HwRev:A2), DCS-7130-96L-R (HwRev:A3), DCS-7130-96LA-F (HwRev:A2), DCS-7130-96LA-F (HwRev:A3), DCS-7130-96LA-R (HwRev:A2), DCS-7130-96LA-R (HwRev:A3), DCS-7130-96LAS-F (HwRev:B1), DCS-7130-96LAS-F (HwRev:B2), DCS-7130-96LAS-R (HwRev:B1), DCS-7130-96LAS-R (HwRev:B2), DCS-7130-96LB-F (HwRev:A2), DCS-7130-96LB-F (HwRev:A3), DCS-7130-96LB-R (HwRev:A2), DCS-7130-96LB-R (HwRev:A3), DCS-7130-96LBA-F (HwRev:A2), DCS-7130-96LBA-F (HwRev:A3), DCS-7130-96LBA-R (HwRev:A2), DCS-7130-96LBA-R (HwRev:A3), DCS-7130-96LBAS-F (HwRev:B1), DCS-7130-96LBAS-R (HwRev:B1), DCS-7130-96LBS-F (HwRev:B1), DCS-7130-96LBS-F (HwRev:B2), DCS-7130-96LBS-R (HwRev:B1), DCS-7130-96LBS-R (HwRev:B2), DCS-7130-96LS-F (HwRev:B1), DCS-7130-96LS-F (HwRev:B2), DCS-7130-96LS-R (HwRev:B1), DCS-7130-96LS-R (HwRev:B2), DCS-7130-96S-F (HwRev:B1) and DCS-7130-96S-R (HwRev:B1)
- When the CortinaApiProxy agent restarts, newly inserted transceivers are not detected and links won't come up.

A Workaround is to restart the XcvrAgent as well, after CortinaApiProxy. (833464)

SKUs Affected: DCS-7130-16G3-F (HwRev:C1), DCS-7130-16G3-R (HwRev:C1), DCS-7130-16G3S-F (HwRev:D1), DCS-7130-16G3S-R (HwRev:D1), DCS-7130-48E-F (HwRev:A4), DCS-7130-48E-R (HwRev:A4), DCS-7130-48EH-F (HwRev:A4), DCS-7130-48EH-R (HwRev:A4), DCS-7130-48EHS-F (HwRev:B1), DCS-7130-48EHS-F

(HwRev:B2), DCS-7130-48EHS-R (HwRev:B1), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48G3-F (HwRev:B1), DCS-7130-48G3-F (HwRev:B2), DCS-7130-48G3-R (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48G3S-F (HwRev:C1), DCS-7130-48G3S-R (HwRev:C1), DCS-7130-48L-F (HwRev:A2), DCS-7130-48L-F (HwRev:A3), DCS-7130-48L-R (HwRev:A2), DCS-7130-48L-R (HwRev:A3), DCS-7130-48LA-F (HwRev:A2), DCS-7130-48LA-F (HwRev:A3), DCS-7130-48LA-R (HwRev:A2), DCS-7130-48LA-R (HwRev:A3), DCS-7130-48LAS-F (HwRev:B2), DCS-7130-48LAS-F (HwRev:B3), DCS-7130-48LAS-R (HwRev:B2), DCS-7130-48LAS-R (HwRev:B3), DCS-7130-48LB-F (HwRev:A2), DCS-7130-48LB-F (HwRev:A3), DCS-7130-48LB-R (HwRev:A2), DCS-7130-48LB-R (HwRev:A3), DCS-7130-48LBA-F (HwRev:A2), DCS-7130-48LBA-F (HwRev:A3), DCS-7130-48LBA-R (HwRev:A2), DCS-7130-48LBA-R (HwRev:A3), DCS-7130-48LBAS-F (HwRev:B2), DCS-7130-48LBAS-F (HwRev:B3), DCS-7130-48LBAS-R (HwRev:B2), DCS-7130-48LBAS-R (HwRev:B3), DCS-7130-48LBS-F (HwRev:B2), DCS-7130-48LBS-F (HwRev:B3), DCS-7130-48LBS-R (HwRev:B2), DCS-7130-48LBS-R (HwRev:B3), DCS-7130-48LS-F (HwRev:B2), DCS-7130-48LS-F (HwRev:B3), DCS-7130-48LS-R (HwRev:B2), DCS-7130-48LS-R (HwRev:B3), DCS-7130-96-F (HwRev:A7), DCS-7130-96-R (HwRev:A7), DCS-7130-96-R (HwRev:A8), DCS-7130-96L-F (HwRev:A2), DCS-7130-96L-F (HwRev:A3), DCS-7130-96L-R (HwRev:A2), DCS-7130-96L-R (HwRev:A3), DCS-7130-96LA-F (HwRev:A2), DCS-7130-96LA-F (HwRev:A3), DCS-7130-96LA-R (HwRev:A2), DCS-7130-96LA-R (HwRev:A3), DCS-7130-96LAS-F (HwRev:B1), DCS-7130-96LAS-F (HwRev:B2), DCS-7130-96LAS-R (HwRev:B1), DCS-7130-96LAS-R (HwRev:B2), DCS-7130-96LB-F (HwRev:A2), DCS-7130-96LB-F (HwRev:A3), DCS-7130-96LB-R (HwRev:A2), DCS-7130-96LB-R (HwRev:A3), DCS-7130-96LBA-F (HwRev:A2), DCS-7130-96LBA-F (HwRev:A3), DCS-7130-96LBA-R (HwRev:A2), DCS-7130-96LBA-R (HwRev:A3), DCS-7130-96LBAS-F (HwRev:B1), DCS-7130-96LBAS-R (HwRev:B1), DCS-7130-96LBS-F (HwRev:B1), DCS-7130-96LBS-F (HwRev:B2), DCS-7130-96LBS-R (HwRev:B1), DCS-7130-96LBS-R (HwRev:B2), DCS-7130-96LS-F (HwRev:B1), DCS-7130-96LS-F (HwRev:B2), DCS-7130-96LS-R (HwRev:B1), DCS-7130-96LS-R (HwRev:B2), DCS-7130-96S-F (HwRev:B1) and DCS-7130-96S-R (HwRev:B1)

- When changing the speed of an interface managed by a CRT55321 PHY (et1-2,31-34,63-64), the Babbage agent may restart once causing all of these interfaces to flap. (837387)
SKUs Affected: DCS-7060DX5-64S
- Using the rc.eos script to add port numbering to the boot config will result in the port numbering not being applied.

Workaround: Please add port numbering through the CLI using 'boot port numbering octal interfaces INTERFACE_COUNT' commands and then reboot the switch. (843173)

SKUs Affected: 7800R3-36D-LC, 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3A-36P-LC, 7800R3A-36PM-LC, 7800R3AK-36DM-LC, 7800R3AK-36DM2-LC, 7800R3AK-36PM-LC and 7800R3K-36DM-LC

- Interfaces with 1GBASE-T SFPs installed may not link up. Workaround is to shut/no shut the port to recover. (847035)

Series Affected: CCS-720DF Series

SKUs Affected: CCS-720DF-48Y, CCS-720DF-48Y-2, CCS-720DF-48Y-2F and CCS-720DF-48Y-F

- CRT55321 interfaces that are not /1 might not be able to get configured at all. (917897)

Layer 2

- MACsec delay protection feature is not working. (282186)
- When running rapid-PVST at high scales of interfaces and vlans, the device may occasionally not transmit a BPDU, which can lead to a BPDU timeout on the neighbor. (442663)
- Performing an SSO switchover on an MLAG primary peer running Multiple Spanning Tree Protocol may result in STP topology reconverging. (470962)
- Restarting STP agent with a high number of interfaces and vlans in PVST mode may result in STP topology reconverging. (472675)
- On EVPN setup using VXLAN and Multihomed active hosts, when the ownership of the host MAC address changes to a new peer, the MAC/IP route may go missing since the new peer is unable to install the MAC address as a dynamic learned entry. (598730)
- Dot1x agent may restart unexpectedly when there are a large number of supplicants present on a port with "dot1x aaa unresponsive action traffic allow vlan <>" and "dot1x mac based authentication always" configured. (731691)
- Pseudowire agent may restart unexpectedly when a large number of flexible cross-connect pseudowires are configured/unconfigured. (732883)
- The following CLI is used to set the source IP address of all outgoing RADIUS packets to that of a particular interface.

```
"[no] ip radius [ VRF ] source-interface <interface-name>"
```

Currently proactive probing of RADIUS servers does not honour this command and will not use the source IP according to this command, the probes which are sent out will still continue to have the source IP of an interface that the switch chooses based on an internal algorithm. (798136)

- Configuring port channels with number greater than 65535 and "spanning-tree port-id allocation port-channel range" feature enabled might cause STP agent to restart continuously.

To work around, reconfigure port channels with number less than or equal to 65535. (832076)

- If Port Security and 802.1X are configured on an interface with supplicants authenticated on the interface, post SSU, it is seen that the supplicants are logged off and are made to re-authenticate resulting in traffic loss for such supplicants. A workaround to avoid this traffic loss would be to disable portsec on the interface. (839987)
- 802.1x authenticator in a switch is sending the RADIUS accounting stop messages with the same identifier value in the RADIUS header message. Due to this a RADIUS proxy server may drop those messages, resulting in the 802.1x authenticator not receiving response for those RADIUS requests. (873700)
- In an EVPN VXLAN IRB deployment, the MTU used by the dynamic VLAN interface created for the tenant VRF might be incorrect.

A workaround is to restart the VXLAN agent. (893443)

- Churn in multihomed EVPN type2 routes may result in some of the multihomed MAC addresses not programmed in hardware when forwarding plane L3 agent restarts at the same time resulting in flooding the packets destined to those MAC addresses.

Workaround is to flap the VXLAN interface. (900375)

Layer 3

- With scaled VRF config, in the order of 1K VRFs, if the Arp agent restarts, it may fail to come up. (275514)
- In certain scale situations upon a sysdb restart, arp entries can be learned in kernel but not show up under 'show arp'. When this happens, the work around is to restart the Arp agent. (275749)
- Configuring 'max-metric router-lsa' on an OSPF router may cause summary LSAs to be generated with maximum metric. (276629)
- Auto discovery of multiple IPv6 link-local BGP neighbors over a single interface does not work and will result in only one of the neighbors transitioning to established state. (289875)
- BFD configurations with large numbers of sessions may see BFD session flaps during initial session bringup shortly after boot. (375773)
- When graceful restart is enabled under "router isis" and other protocol routes are redistributed into IS-IS then on ASU or SSO, some traffic that is forwarded using redistributed routes might be lost. (413706)
- Removing SVI configuration at scale via a config-replace operation may cause the BFD agent to unexpectedly restart. (414702)

- The MLAG agent may exit unexpectedly after upgrading from 4.22.0 or an older release if the MLAG peer switch is still running the older software version and DHCPv6 prefix delegation routes are being sync'd across the peers. (425810)
- The output of 'show ip nat sync peer' may incorrectly indicate a connected state, even though an incorrect peer IP address has been configured. (464011)
- DHCP relay agent might exit unexpectedly if several VRFs are continuously deleted and re-created. (471169)
- For short-lived TCP connections, NAT may not tear down the address mapping when it should. The workaround is to lower the TCP timeout using the `ip nat translation tcp-timeout` CLI. (526152)
- When large number of distribute lists are configured in RIP then it may result in route flaps and the protocol agent CPU utilization can increase significantly. (536500)
- KernelFib agent may restart on interface flap, if there are a large number of route resolving through that interface (539757)
- KernelFib agent may restart, if EVPN config is toggled in quick succession. (557358)
- RSVP tunnels might not get established or may use non-best paths when IS-IS system ID or OSPF router ID of any two neighbouring nodes in the IGP network is changed back-to-back and the best path goes via that node. (572080)
- KernelFib agent may restart when interface flaps lead to a large number of routes (entire routing table) being withdrawn (597557)
- Moving the managementactive(ma0) interface to a VRF and restarting Sysdb would lead to managementactive interface staying down in the default VRF after Sysdb(and other agents) come back up. (613014)
- During Zero Touch Provisioning(ZTP) in an IPv6 environment, even if the Router Advertisement(RA) packets do not have the M and the O bits set, the Arista switch still tries to contact a DHCPv6 server for information other than the IPv6 address. (691387)
- Cspf agent may sometimes restart when there are two or more routers with duplicate TE router ID in the network (700468)
- With multiple IS-IS instances configured in the same VRF, the IS-IS agent may sometimes restart again when it is restarting with GR enabled. (742959)
- Nodes marked with the IS-IS overload bit will not be avoided if they are specified in an include-hop constraint (754609)

- Upon the Arp agent starting up in setups with a large number of routes, the Arp agent may restart. (764196)
- The Sflow agent selects a random port within the ephemeral port range, 32768 - 60999, to bind to. Since the VxlanSwFwd agent is statically bound to a port within that range, 51023, the Sflow agent might select that same port, causing the VxlanSwFwd agent to repeatedly restart until the Sflow agent chooses another port.

As a workaround, restart the Sflow agent. (780707)

- After reload, static ARP or ND entries may not get correctly restored.

Workaround is to restart the Arp agent. (790370)

- VRF names with "." result in repeated restart of the Launcher agent while starting a process with a name like "Rib-vrf-foo.bar".

To restore the system, first remove the VRF with "." in its name with "no vrf instance <vrf.name>". Then, delete any file in /etc/ProcMgr.d/run/ ownedByLauncher/ of the form "<agent>-vrf.<vrf.name>". For an example VRF named "foo.bar", the file to delete might be "Rib-vrf.foo.bar" or, when using the multi-agent routing protocol model, a file for a protocol-specific agent such as "Ospf-vrf.foo.bar" or "Isis-vrf.foo.bar". (792987)

- Router Solicitation (RS) packets with the source address that is not the unspecified address and no source link-layer address option, causes a Router Advertisement (RA) to be sent with the destination address as unicast address and the destination link-layer address as all-nodes multicast address. (802490)
- RSVP tunnels can fail to come up if an include-hop is specified on a path that causes the shortest path to the destination to contain a loop (827083)
- Kernel routes with special IPv6 next hops (IPv4-over-IPv6, BGP-LU, SBFD SR-TE) require hardware forwarding and may blackhole traffic after an L3 forwarding agent restart because a kernel network device was not created. A workaround is to configure a Loopback interface with an IPv6 address and "hardware forwarding id". (858684)
- IS-IS process may restart if the SRGB range received in the IS-IS SR-Capabilities sub-TLV is out of MPLS label space (0 to 1048575) (865940)
- When the Refresh Overhead Reduction extension is enabled, receiving a Path message through a bypass (during Fast-Reroute) without an Explicit Route Object might cause the agent to be terminated (884191)
- Bgp agent restart may result in positions that were in use for Optimal Route Reflection (ORR) becoming unavailable, and failing to be displayed by

the 'show bgp optimal-route-reflection positions' command.

An affected ORR position may be recovered by removing all configuration that assigns that position to a route reflector client (in a supported address family mode), waiting for 2 minutes, then restoring those configuration commands. (886899)

- IS-IS agent may restart with v4+v6 multi-topology if traffic-engineering is enabled and if with non multi-topology peer we exchange global v6 interface address and v4 interface address and the corresponding non multi-topology link is shutdown. (890036)
- When Extended Reachability TLV #22 for an IS-IS neighbour is present in multiple LSP fragments, post IS-IS restart, IS-IS agent may again restart when the same Extended Reachability TLV is withdrawn. (890794)
- If local adjacency segment protection is enabled and IS-IS agent restarts at scale, MPLS agent may restart sometimes. (893203)
- After the Cspf agent restarts, tunnels using IS-IS flexible algorithm can experience momentary traffic loss. (894389)
- After the RsvpTunnel is terminated, if the system is slow or the RsvpProtocol agent is not running, the RsvpTunnel agent might restart unexpectedly. Graceful Restart or Hitless Restart can be enabled as a workaround. (899784)
- The OSPFv3 agent may restart unexpectedly during a SSU if there is an IPv6 link local address configured on any OSPFv3 interface (901632)
- When BGP is configured to populate a field-set with peers using 'field-set ...' and 'peers ...' under 'traffic-policy field-set mappings' and BGP activates or deactivates in a non-default VRF, some or all field-set members may be removed erroneously.
To fix the exported field-set, enter the 'field-set (ipv4|ipv6) ...' mode and remove and re-add the config, eg. 'no peers ibgp' then 'peers ibgp'. (906216)

Multi-agent

- Performing a stateful switchover in a switch with a large number of VRFs configured, with BGP config in all VRFs and IGP configured in a large number of VRFs, may cause the system to restart unexpectedly, leading to traffic loss. (376773)
- SSO on ASBR with MPLS L3 VPN Inter-AS Option B configured may result in temporary traffic loss due to some VPN labels being reallocated (498848)
- SSO with MPLS L3 VPN configured may result in temporary loss of traffic flowing from MPLS core towards the tenant sites (507353)

- IS-IS agent will restart if a BGP-LU route is redistributed into IS-IS with a prefix segment configured for the corresponding BGP LU prefix and if the same prefix is also configured on a loopback with a different node segment. (744361)
- Under rare circumstances, Bgp Agent may restart when VRF configuration is removed (845945)
- The IpRib agent might restart unexpectedly during system restart. (849261)
- Setting community no-advertise, no-export, and local-as inline in route-map results in integer values in the YANG path /routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/set-community/inline instead of NO_ADVERTISE, NO_EXPORT, and NO_EXPORT_SUBCONFED (858821)
- Under rare circumstances, StaticRoute agent may not install any routes in some VRFs on agent startup.
As a workaroud, try restarting the StaticRoute agent. (863907)
- Ignore received BGP attributes configuration is sometimes not honoured by some dynamic BGP peers after 'clear ip bgp *'.
Workaroud by issuing 'clear ip bgp * soft in' after session reset to send route-refresh message. (867128)
- Connected routes may disappear in some rare cases during switch reboot, configure replace or switchover from standby to active if the attached interfaces to the routes belong to non-default VRF. (881936)
- Leaking routes between VRFs which have a BGP community which matches a BGP field-set community mapping (either traffic-policy or VRF selection policy) will cause BGP to restart unexpectedly.

As a workaroud the route (or BGP community) can be filtered via route-map or RCF policy when leaking. (886394)

- Configuring 'neighbor <> ttl maximum-hops 0' command, under 'router bgp' mode, will not take effect after device reload.

The only workaroud is to reconfigure the command after the Bgp agent is running, or to use another value for the 'maximum-hops' such as '1'. (890018)

- BGP agent can crash if BGP monitoring (BMP) is configured to export 6PE routes (897060)
- A malformed Prefix-SID BGP path attribute will result in a session reset rather than attribute discard. (899981)

- A malformed BGP Link State path attribute will result in a session reset rather than attribute discard. (900147)
- When the config contains only interface peer groups (for IPv6 link-local) and no static or dynamic peers or peer-groups are configured, SSU reload of switch (using 'reload fast-boot' command) may result in sub-seconds of traffic drop.
This is because BGP AfiSafi convergence is declared prematurely even before BGP global convergence.
One workaround is to provision at least one static peer - can be a dummy peer with non-existing IP address. The presence of dummy peer helps by prolonging AfiSafi convergence timeout to at least slow peer timeout value (default 90 sec). (907573)
- Bgp agent restart may result in positions that were in use for Optimal Route Reflection (ORR) becoming unused though available as displayed by the 'show bgp optimal-route-reflection positions' command. The same problem could happen after device reload or a major BGP configuration change. An affected ORR position may be recovered by removing all configuration that assigns that position to a route reflector client (in a supported address family mode), waiting for 2 minutes, then restoring those configuration commands. (908676)

Rib

- In the ribd routing protocol model, prior to EOS 4.21.3F release, the command
"neighbor <peer|peer-group> send-community extended" (under "router bgp") will result in both standard and extended communities being sent (in outbound route advertisements) to the corresponding peer.

Starting with EOS release 4.21.3F this behavior has changed. The command "neighbor <peer|peer-group> send-community extended" will result in *only* extended communities sent to the peer.

If the previous behavior is desirable, then the command "neighbor <peer|peer-group> send-community" without additional keywords can be used. This command will ensure that all applicable community types are included in outbound route advertisements for the peer.

This change can be done in any release (running EOS version 4.18.1F or newer) prior to the upgrade to EOS 4.21.3F (or newer) release.

EOS 4.21.3F or newer releases allow much more granular control of what community types (standard, extended, large communities etc.) are included in the

outbound
route advertisements. (384629)

- In an EVPN VXLAN deployment where EVPN Type-2 or Type-5 routes have been received, configuring a static EVPN route that reuses a nexthop VTEP received from EVPN can cause received routes to become unresolved.

A workaround is to unconfigure the static EVPN route. (443964)

- The Rib agent may restart when large number of interfaces go down at the same time and there is a large scale of BGP routes reachable over these interfaces. (549816)
- Next hops for static and BGP routes undergo proactive ARP resolution, and routes for which the next hop is not resolved behave as drop routes. When a subnet (associated with a local IP address) is moved from one interface to another, next hops reachable via that network change to reflect the new interface, but the proactive ARP resolution may not change accordingly, resulting in traffic loss despite the affected routes appearing correctly in "show ip(v6) route" output.

This will principally affect static and iBGP routes because by default eBGP routes' next hop is the address used for peering, so ARP resolution is completed before the route is learned. The straightforward workaround is to restart the Rib agent for the affected VRF. If all next hops affected by an address move can be identified, actively connecting to them from the switch, such as with ping, will also result in ARP resolution and restore forwarding. (804746)

MLAG

- If non-MLAG reload-delay is configured to be lower than the MLAG reload-delay when LACP standby is used, traffic entering the non-MLAG interfaces destined towards hosts on the MLAG interfaces will be lost during the LACP standby period. (83330)
- MLAG ISSU breaks on a switch running a version of EOS < 4.20.5 if it receives ARP entries from the VXLAN Control Service (VCS) while the peer has been upgraded to EOS version >= 4.20.5 (497776)
- When 'switchport mode trunk phone' is enabled on one peer of an MLAG switch, it must also be manually configured on the other MLAG peer to allow for the switchport configuration to take effect. (877198)

Multicast

- An S,G route may not properly cleanup previously sent S,G joins. This would prevent S,Gs from being deleted in the upstream router. A work around is to flap the interface in which the join is sent on. (291756)
- If large number of IGMP reports are received in a very short interval, some of them might get dropped resulting in the multicast route not being created. The routes will be re-learned on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIgmp", consider increasing the bandwidth for this Copp class. (356675)
- show ip igmp vrf <name> groups command does not filter the IGMP interfaces as expected. (368071)
- High CPU utilization might be observed for the forwarding agent when PIM Bidir routes are present. (479119)
- Removing router pim bsr also remove router pim sparse mode configuration. Workaround is to just remove specific configuration under router pim bsr mode. (516188)
- In a PIM-SM network when an RP is also the first hop router, it may take upto a minute to advertise source to MSDP if the source goes from inactive state to active state. (584263)
- Configuring "ip igmp host-proxy <ACL-NAME>" where the access list has a large multicast group prefix length may cause the IgmpHostProxy to be repeatedly restart resulting in traffic loss. Use smaller prefix lengths to work around the problem. (586865)
- Igmp host proxy configuration using overlapping prefixes where a permit is followed by a deny rule, the groups in the permit range are not proxied. As a workaround, remove `ip igmp host proxy access-list <acl>` and use `ip igmp host proxy <group>` to explicitly configure the permitted groups. (619688)
- Running MFA feature on x86_64 can lead to memory contention. (662221)
Series Affected: CCS-710P, CCS-720XP and DCS-7010TX Series
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- MVPN related multicast route might be programmed incorrectly after BGP restart on a scaled MVPN PE (677280)
- On devices running EVPN multicast over a multicast underlay with a high scale of VXLAN enabled VLANs, triggers that cause churn in the mapping of VLANs to underlay multicast groups such as restarting BGP via configuration change can cause BessMgr agent to restart.

This affects software forwarded multicast traffic and new multicast traffic. BessMgr will recover without need for intervention. (694001)

- On devices running multicast with `software-forwarding sfe` with a larger multicast interface scale and/or VRF scale. BessMgr agent may get restarted by ProcMgr. (700023)
- On devices running EVPN VXLAN multicast OISM, Overlay multicast traffic with TTL1 arriving on VXLAN enabled VLAN in a VRF with `evpn multicast` enabled is not VXLAN encapsulated when there is a receiver in the same VLAN on remote VTEP.

Workaround

Use TTL>1 or enable `redistribute igmp` under router BGP for the VLAN on the receiver VTEP. (703609)

Series Affected: CCS-720XP, DCS-7050CX3, DCS-7050SX3, DCS-7050TX3 and DCS-7300X3 Series

- If KernelMfib agent is/gets terminated the device will be unable to discover new multicast sources and/or detect traffic duplication. This can result persistent multicast traffic loss the first case and persistent traffic duplication in the latter.

Workaround is to restart Pimsm and PimReg agent by issuing "agent pimsm terminate", "agent pimreg terminate" command. (707691)

- When the VXLAN interface is shutdown and immediately unshut using the "no shutdown" CLI, for some of the multicast tenant VRFs, the SBD IMET routes may be withdrawn and never re-advertised. This might impact any EVPN multicast traffic flowing through that device. A workaround would be to shutdown the VXLAN interface again, wait for all SBD IMET routes to be withdrawn and then unshut the interface. (722075)
- On MLAGed devices running multicast, `shut` followed by `no shut` under `mlag` submode can result in stale multicast routes (ones existing before the shut was executed) being programmed in the hardware resulting in traffic loss.

Workaround: Toggle `routing` under `router multicast` for the affected VRF (727482)

- In an MVPN setup, traffic at ingress PE might get dropped if multiple flaps occur on the source interface. (777298)
- On devices running multicast with MRIB lookup enabled(rpf resolution ribs multicast-rib), when the MRIB lookup yields a connected route it is not used as the RPF and results in traffic not being forwarded. (783549)

- On devices running multicast with "software forwarding sfe," BessMgr agent may restart unexpectedly when performing configure replace. (828727)
- On devices running multicast with 'software forwarding-sfe', running 'configure replace' may result in persistent multicast traffic loss in the default VRF.

Use 'agent BessMgr terminate' to workaround the issue. (829341)

- On devices running PIM sparse mode with ECMP, a multicast route might get stuck in switching state and continuously draw traffic on more than one link. In case this happens in an EVPN setup, this can result in persistent traffic duplication. In all other cases, no other traffic impact would be noticed.

Workaround: `clear ip mroute <group> <source>` to reset the route. (838201)

- On devices running EVPN L2 multicast with a PIM -ASM underlay, persistent multicast traffic duplication may be seen.
To avoid this issue, consider having PIM-SSM in underlay by configuring `vxlan multicast protocol pim ssm`. (845134)
- On devices running EVPN multicast OISM, when a PIM External Gateway (PEG) device is configured as RP, receivers in the external network that join the EVPN network via the PEG-nonDR RP may experience traffic loss.

`clear pim ipvX vrf <VRF> sparse-mode route <S> <G> ` on the PEG-non-DR will recover the flow. (848542)

- Msdp may choose the wrong Rpf Peer for receiving MSDP messages when AS Path rules are considered. Workaround by using MSDP Peer groups or statically configuring MSDP RPF Peer. (881844)
- On devices running multicast with `software-forwarding sfe`, running configure replace could result in no new flows from being discovered resulting in traffic loss.

Workaround is to perform `agent BessMgr terminate` to recover (898856)

NAT

- On MLAG set-up using NAT synchronization between the two peers, TCP connection might reset while the connection is being established.

A workaround is to make sure that all the control packets for a specific connection arrive on the same peer. Careful network design can be used to achieve this goal. (681428)

Series Affected: DCS-7170 Series

- When "ip nat tcp flow creation trigger syn" is set, if dynamic NAT configuration is modified or if the admin shutdown/no shutdown performed on an interface configured with dynamic NAT while NAT traffic is flowing, NAT connections may be left untranslated.

Workaround is to manually flush all the entries using the command "conntrack -F" in the corresponding network namespace bash mode. If configured on a non-default VRF (vrfName, for example), run the command "ip netns exec ns-vrfName conntrack -F".0 (748487)

- NAT translations may not happen if either the ingress or egress interface is repeatedly flapped. (777877)
Series Affected: DCS-7130 Series
- Hardware translations may not properly programmed if Dynamic NAT profile is changed while high amount of NAT entries are created. (792730)
- NAT agent might restart, if there are large number of address-only NAT flows from the same host. This is caused due to an inefficient loop that updates the last activity time for each of these flows causing NAT agent to fill the attrlog buffer. (798605)

SwitchApp

- Heavy CPU bound traffic may cause BGP flaps despite Copp being configured with ingressCpuDrop counts being observed under "show fpga switch counters". (748835)
- MLAG peer link does not come up post reboot until the non-mlag reload delay timer expires (758022)
- A Security ACL can not be applied to the same VLAN as a NAT ACL (772289)
- If the MakoCentral agent is restarted, the MakoSharedResources can restart unexpectedly, and some configured NAT sessions may not be correctly programmed in hardware. (773579)
- Switch interfaces may discard transmitted traffic for up to 30s after the corresponding front panel port is brought up despite having an UP status. If being UP causes the interface to attract traffic that would otherwise take a different path (e.g. if the port is a routed port it will install its connected prefix and this may attract traffic from neighboring routers) then this may cause traffic loss in the network. (778047)
Series Affected: DCS-7130 Series
- When a member is added to a port-channel, reverse NAT translation may not take effect for the new member.

The workaround is to remove and re-add the relevant NAT configuration.

(778996)

Series Affected: DCS-7130 Series

- After MakoL3Unicast and/or MakoSharedResources agent restarts routes for connected hosts may remain programmed in the FPGA (and be visible via the 'show fpga switch ip route') even though the route and ARP entry have been removed (no longer shown in 'show ip route' and 'show arp').

Restart SwitchApp using 'switch default' followed by 'disabled' followed by 'no disabled' to clear this state. (800353)

- ACLs configured with a lt match on ports will not match on packets with port equal to 0. The workaround is to configure an explicit rule to match on port 0 (809597)
- If a SwitchApp device has configuration that matches following condition:
 1. There are trunk interfaces configured on the switch
 2. ACL applied on an SVI that is included in the trunk VLAN range
 3. The ACL involved has counters enabled
 4. There is at least one rule in ACL involved uses L4 port range
 5. There are multiple SVIs with ACL configured and all in trunk VLAN range
 6. The involved SVIs must have members sharing the same pipeline
 7. When ACL applied on the SVI, members of SVI involved needs to be active

When links of the SVI involved are down, MakoSharedResource agent might (not always) crash with the following failure: `assert(!hasCounter || conterRef)`

There are many workarounds to avoid the issues by breaking any of the conditions above, for example:

1. disable counters in ACL involved
 2. Do not use L4 port range in ACL rules, split it into individual rules with single L4 port filter
 3. Split members of SVI involved into different pipelines
 4. Narrow trunk interface VLAN range using CLI 'switchport trunk allowed vlan xxx'
- ... (869876)
- When switchApp is used and ACL is configured, on MakoSharedResources agent restarting alone (either caused by bug or manually), the internal ACL states is lost, which causes inconsistency between hardware and software.

The workaround is to reboot the device. (869902)

- Under heavy load SwitchApp may fail to update IP routes in the FPGA forwarding plane. (908335)

DCS-7170 Series

- NAT sync does not work reliably if the number of NAT connections are much higher
(in the order of 50K or more NAT connections). (506507)

DCS-7170 Series

- PTP transparent mode is currently not supported. Attempting to configure PTP transparent mode with the ``ptp mode e2ettransparent`` or ``ptp mode p2pttransparent`` command will cause the BfnSlice agent to unexpectedly restart. (609143)
Series Affected: DCS-7170 Series

CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM and CCS-750 Series

- If link speed is changed on a 2.5G interface, Strata forwarding plane agent may unexpectedly restart with "Process died with signal 3 \((SIGQUIT\)\" signature. Subsequent restart of the Strata agent will resolve the issue. (368424)
- The forwarding slice agent may restart unexpectedly when a SEU is encountered in the BSC_AG_AGE_TABLE during initialization. This may result in brief traffic outage. Switch is expected to start normal operation after the agent restart. (417588)
- On phone trunk ports which are not controlled by Dot1x, phones do not honour port-security max-address limit and get installed in the L2 table even after the port-security limit is breached (494095)
- Interfaces configured for MAC loopback are limited to 10 Mbps instead of the configured speed. (585395)
SKUs Affected: CCS-720XP-24Y6, CCS-720XP-48Y6 and DCS-7010TX-48
- PhyIsland agent might restart unexpectedly multiple times, while configuring 2.5GBASE-T interfaces or during a reload. This will cause link flaps and traffic disruption on all front-panel BASE-T ports. (586920)
SKUs Affected: CCS-720XP-96ZC2, CCS-722XPM-48ZY8 and CCS-750X-48ZP-LC
- During hitless restart of forwarding plane agent(s), packets which were denied by security ACLs may leak through the chip(s) being managed by the agent(s) for a second. The deny security ACL rules start acting on the traffic once the forwarding plane slice agent(s) is warm.

There is no workaround for this bug. (589311)

Series Affected: CCS-750 Series

- 2.5GBASE-T interfaces may fail to pass traffic in the ingress direction after a reload, linecard insertion (for modular systems) or reconfiguration, if they are configured with ``speed 100full`` command.

Workaround is to run ``shutdown`` followed by ``no shutdown`` on the affected interfaces. (598487)

SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-48ZC2, CCS-720XP-96ZC2, CCS-722XPM-48ZY8 and CCS-750X-48ZP-LC

- When MACsec is disabled on an interface which is used for VXLAN underlay, non-unicast traffic may not seamlessly transition to clear text. The work around is to use "shutdown" followed by "no shutdown" after disabling MACsec on an interface. (637199)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- An SEU (Single Event Upset) in the L3_ENTRY_LP table results in an unexpected forwarding agent restart causing a brief traffic disruption. (641120)

Series Affected: CCS-750 Series

- EVPN VXLAN A-A multihoming is not supported when MACSec is enabled on the core port. Configuring it may result in undesirable double delivery of BUM traffic. (642181)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- When MACsec is disabled on an interface, multicast traffic may not seamlessly transition to clear text. The work around is to use "shutdown" followed by "no shutdown" after disabling MACsec on an interface. (645584)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- On switchcard switchover, there may be extended traffic outage. (682181)

Series Affected: CCS-750 Series

- 2.5GBASE-T ports configured with ``speed 100full`` (100M forced speed) configuration may fail to pass traffic even if they report linkup after a failing cable test run with the ``test 11 cable base-t interfaces INTFS`` CLI command.

Workaround is to run ``shutdown`` followed by a ``no shutdown`` on the affected interfaces. (800901)

SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-48ZC2, CCS-720XP-96ZC2, CCS-722XPM-48ZY8 and CCS-750X-48ZP-LC

- 5GBASE-T/10GBASE-T ports may stop passing traffic or start causing packet corruption with errors. (e.g. FCS/Symbol/Rx/Runt packet errors).

Workaround is to run `shutdown` followed by a `no shutdown` on the affected interfaces. (801926)

SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-48ZC2 and CCS-720XP-96ZC2

- Under rare circumstances, after the system reloads, certain switch chips may fail to initialize correctly, resulting in Ethernet links being established but network traffic being dropped.

The workaround is to reset the affected switch chip using `platform trident <x> reset`, replacing <x> with the name of the switch chip as shown by `show platform trident system port`. Alternatively, `platform trident \* reset` can be used to reset all switch chips in the system. (808236)

Series Affected: CCS-750 Series

- Excessive link flaps can cause high CPU usage, potentially causing agents to restart. Agent restarts may cause brief traffic disruptions.

Workaround is to configure the desired linkup speed using `speed forced <SPEED>` on the affected interfaces. (815477)

SKUs Affected: CCS-750X-48ZXP-LC and DCS-7050TX3-48C8

- BASE-T ports may stop passing traffic or start causing packet corruption with errors. (e.g. FCS/Symbol/Rx/Runt packet errors).

Workaround is to run `shutdown` followed by a `no shutdown` on the affected interfaces. (816890)

SKUs Affected: 7300X3-48TC4-LC, CCS-750X-48ZXP-LC and DCS-7050TX3-48C8

- 1000BASE-T ports may fail to link up after a reload or forwarding agent restart.

Workaround is to restart the forwarding agent with `platform trident reset` CLI command. This will cause link flaps and brief traffic disruption on all front panel ports.

Affected ports:

DCS-7010TX-48: 25-48

DCS-7010TX-48-DC: 25-48

DCS-7010TX-48C: 25-48

CCS-710P-12: 1-12

CCS-710P-16P: 1-12

CCS-720DP-24S: 1-24

CCS-720DP-48S: 25-48

CCS-720DT-24S: 1-24

CCS-720DT-48S: 25-48

CCS-722XPM-48Y4: 25-48 (818224)

SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DP-24S, CCS-720DP-48S, CCS-720DT-24S, CCS-722XPM-48Y4, DCS-7010TX-48 and DCS-7010TX-48-DC

- 2.5GBASE-T ports configured with ``speed auto 100full`/`no speed`` configurations and linked up at 100M speed may fail to pass traffic after an SSU reload (i.e. reload fast-boot) if they link down while the SSU reload is in progress, and come back up after the SSU reload is complete.

Workaround is to run ``shutdown`` followed by a ``no shutdown`` on the affected interfaces. The issue can also be avoided by configuring the susceptible interfaces with ``speed forced 100full`` configuration prior to the SSU reload. Please note that the configuration change will flap the interfaces. (840024)

SKUs Affected: CCS-720XP-24ZY4 and CCS-720XP-48ZC2

- PhyIsland agent restart after an SSO (stateful switchover) may be hitful, causing link flaps and brief traffic disruption on all front panel BASE-T ports of the affected linecard(s). (841719)

SKUs Affected: 7300X3-48TC4-LC, CCS-750X-48THP-LC, CCS-750X-48TP-LC, CCS-750X-48ZP-LC and CCS-750X-48ZXP-LC

- BASE-T interfaces may stay down indefinitely after forwarding plane agents restart if the interfaces are configured with the default speed or the ``speed auto`` config.

Workaround is to configure the BASE-T interfaces with ``speed auto <SPEED>`` command where the `<SPEED>` is the desired linkup speed in order to avoid the issue, or to run ``shutdown`` followed by ``no shutdown`` on the affected interfaces in order to recover. (852418)

SKUs Affected: 7300X3-48TC4-LC, CCS-720XP-24Y6, CCS-720XP-24ZY4, CCS-720XP-48Y6, CCS-720XP-48ZC2, CCS-720XP-96ZC2, CCS-722XPM-48ZY8, CCS-750X-48THP-LC, CCS-750X-48TP-LC, CCS-750X-48ZP-LC, CCS-750X-48ZXP-LC and DCS-7050TX3-48C8

- BASE-T interfaces in the range Et1-Et24 may stay down indefinitely after forwarding plane agents restart if the interfaces are configured with the default speed or the ``speed auto`` config.

Workaround is to configure the BASE-T interfaces with ``speed auto <SPEED>`` command where the `<SPEED>` is the desired linkup speed in order to avoid the issue, or to run ``shutdown`` followed by ``no shutdown`` on the affected interfaces in order to recover. (852419)

SKUs Affected: CCS-720DP-48S, CCS-722XPM-48Y4, DCS-7010TX-48 and DCS-7010TX-48-DC

- 5GBASE-T interfaces may stay down indefinitely after forwarding plane agents restart if the interfaces are configured with the default speed or the ``speed auto`` config.

Workaround is to configure the BASE-T interfaces with ``speed auto <SPEED>`` command where the `<SPEED>` is the desired linkup speed in order to avoid the

issue, or to run `shutdown` followed by `no shutdown` on the affected interfaces in order to recover. (852421)

SKUs Affected: CCS-710P-16P

- Transceiver power will be limited to 1W. Modules requiring higher power will be disabled. (899095)

SKUs Affected: CCS-720DP-24S, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-48S, CCS-720DP-48S-2F, CCS-720DP-48S-F, CCS-720DT-24S, CCS-720DT-24S-2F, CCS-720DT-24S-F and CCS-720DT-48S-2F

CCS-750 and DCS-7060X4 Series

- In a chassis with redundant supervisors configured with the SSO redundancy protocol, after a switchover has been performed, a forwarding agent restart may cause the active supervisor to experience a fatal CPU error and restart, which will cause another switchover. (586096)

Series Affected: CCS-750, DCS-7300X and DCS-7300X3 Series

DCS-7170, DCS-7280R and DCS-7280R2 Series

- The AlgoMatch forwarding agent will restart when config replace is used to configure access-lists with large number of rules. (240540)
- Following a physical error on link between Jericho and Algomatch FPGA the SandHalo agent may restart, causing ACLs to be removed and reinstalled. (283917)
- When an acl configuration is applied to a large number of SVIs in a configuration session, the acl configuration fails. (367513)
- A SEU can occur on the Jericho side of the ELK interface to the AlgoMatch FPGA. These error interrupts can happen for other reasons like oversubscribing the ELK interface, but if they persist, they are usually the result of a SEU. This results in dropped control packets from SandHalo and SandHalo will unexpectedly restart continuously. A reboot or linecard reset is required to clear this condition. (486369)

Series Affected: DCS-7280R2 and DCS-7500R2 Series

DCS-7130 Series

- When configured to an illegal speed, interfaces will go to inactive rather than errdisabled. This will cause the interface to seemingly disappear from 'show interface status' unless 'service interface inactive expose' setting is used to display it.

Workaround: change the speed configuration of the interface to a valid speed. To check which speeds are valid, please use 'show interface hardware'. (457130)

SKUs Affected: DCS-7130LBR-48S6QD, DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R

- Some transceivers sometimes report an invalid EEPROM checksum.

Workaround is to unplug/replug the module. (670128)

SKUs Affected: DCS-7130-16G3-F (HwRev:C1), DCS-7130-16G3-R (HwRev:C1), DCS-7130-16G3S-F (HwRev:D1), DCS-7130-16G3S-R (HwRev:D1), DCS-7130-48E-F (HwRev:A4), DCS-7130-48E-R (HwRev:A4), DCS-7130-48EH-F (HwRev:A4), DCS-7130-48EH-R (HwRev:A4), DCS-7130-48EHS-F (HwRev:B1), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48G3-F (HwRev:B1), DCS-7130-48G3-F (HwRev:B2), DCS-7130-48G3-R (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48G3S-F (HwRev:C1), DCS-7130-48G3S-R (HwRev:C1), DCS-7130-48L-F (HwRev:A2), DCS-7130-48L-F (HwRev:A3), DCS-7130-48L-R (HwRev:A2), DCS-7130-48L-R (HwRev:A3), DCS-7130-48LA-F (HwRev:A2), DCS-7130-48LA-F (HwRev:A3), DCS-7130-48LA-R (HwRev:A2), DCS-7130-48LA-R (HwRev:A3), DCS-7130-48LAS-F (HwRev:B2), DCS-7130-48LAS-F (HwRev:B3), DCS-7130-48LAS-R (HwRev:B2), DCS-7130-48LAS-R (HwRev:B3), DCS-7130-48LB-F (HwRev:A2), DCS-7130-48LB-F (HwRev:A3), DCS-7130-48LB-R (HwRev:A2), DCS-7130-48LB-R (HwRev:A3), DCS-7130-48LBA-F (HwRev:A2), DCS-7130-48LBA-F (HwRev:A3), DCS-7130-48LBA-R (HwRev:A2), DCS-7130-48LBA-R (HwRev:A3), DCS-7130-48LBAS-F (HwRev:B2), DCS-7130-48LBAS-F (HwRev:B3), DCS-7130-48LBAS-R (HwRev:B2), DCS-7130-48LBAS-R (HwRev:B3), DCS-7130-48LBS-F (HwRev:B2), DCS-7130-48LBS-F (HwRev:B3), DCS-7130-48LBS-R (HwRev:B2), DCS-7130-48LBS-R (HwRev:B3), DCS-7130-48LS-F (HwRev:B2), DCS-7130-48LS-F (HwRev:B3), DCS-7130-48LS-R (HwRev:B2), DCS-7130-48LS-R (HwRev:B3), DCS-7130-96-F (HwRev:A7), DCS-7130-96-R (HwRev:A7), DCS-7130-96-R (HwRev:A8), DCS-7130-96L-F (HwRev:A2), DCS-7130-96L-F (HwRev:A3), DCS-7130-96L-R (HwRev:A2), DCS-7130-96L-R (HwRev:A3), DCS-7130-96LA-F (HwRev:A2), DCS-7130-96LA-F (HwRev:A3), DCS-7130-96LA-R (HwRev:A2), DCS-7130-96LA-R (HwRev:A3), DCS-7130-96LAS-F (HwRev:B1), DCS-7130-96LAS-F (HwRev:B2), DCS-7130-96LAS-R (HwRev:B1), DCS-7130-96LAS-R (HwRev:B2), DCS-7130-96LB-F (HwRev:A2), DCS-7130-96LB-F (HwRev:A3), DCS-7130-96LB-R (HwRev:A2), DCS-7130-96LB-R (HwRev:A3), DCS-7130-96LBA-F (HwRev:A2), DCS-7130-96LBA-F (HwRev:A3), DCS-7130-96LBA-R (HwRev:A2), DCS-7130-96LBA-R (HwRev:A3), DCS-7130-96LBAS-F (HwRev:B1), DCS-7130-96LBAS-R (HwRev:B1), DCS-7130-96LBS-F (HwRev:B1), DCS-7130-96LBS-F (HwRev:B2), DCS-7130-96LBS-R (HwRev:B1), DCS-7130-96LBS-R (HwRev:B2), DCS-7130-96LS-F (HwRev:B1), DCS-7130-96LS-F (HwRev:B2), DCS-7130-96LS-R (HwRev:B1), DCS-7130-96LS-R (HwRev:B2), DCS-7130-96S-F (HwRev:B1) and DCS-7130-96S-R (HwRev:B1)

- Removing the cable from a BASE-T transceiver or shutting down the peer interface may leave the FPGA interface behind the transceiver in a connected state. (797593)

Series Affected: DCS-7130 Series

- In some cases when there is simultaneously a routed and a dynamically source-NAT-ted traffic, the rate of the NAT-ted one might be limited.

Workaround: don't configure a behind-the-NAT subnet and a routed subnet within the same port group (ports are grouped in fours). (884065)

7358, 7368, 7388, CCS-750, DCS-7300X, DCS-7500R, DCS-7800 and DCS-7800R3 Series

- When many IPv6 routes are configured and the BFD protocol in use, BFD sessions may flap when a linecard is removed or shut down. (279129)
- ManagementActive Agent repeatedly restarts and Management0 interface is not reachable in Modular systems. Doing switchover will address the issue. (369565)
- After linecard removal and re-insertion, lag member links may remain inactive due to "waiting for hardware to program port to RX (or receive)".

Restarting the StrataLag agent recovers the port channels. (501496)
Series Affected: DCS-7300X Series

- After LedPolicy agent restart or SSO, all LEDs are turned off. Workaround is to do "shut/no shut" on all interfaces. (647115)
- A Linecard with link issues on the control plane may cause the Supervisor to reboot (670003)
- Correctable PCIe errors may be seen on a forwarding chip control plane link. (772195)
- If ManagementActive interface is configured with redundancy enabled, then ManagementActive agent can restart unexpectedly during supervisor switchover. (805315)

7500R3, DCS-7020, DCS-7130, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280R3A, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800, DCS-7800R3 and DCS-7800R3A Series

- When the SandMcast and the SandFap agents are restarted, the VLAN floodset may not contain all the ports for several minutes. (67439)
- When in MLAG mode without ip routing enabled, under some control plane oversubscription scenarios the MLAG peer link might come down. The workaround is to either enable IP routing or have a static ARP entry on the MLAG peers for each other. (90247)
- An IPv6oGRE packet terminated using decap-group with a host route destination IPv6 address will be dropped.

The workaround is to disable exact match host routes using "no platform sand ipv6 host-route exact-match". (95873)

- On DCS-7280E, DCS-7500E series, DCS-7050X, DCS-7250X, and DCS-7300 series switches, the vlan tag on IGMP messages(report, query, leave) generated by the switch is not translated in the egress direction by the switchport vlan mapping command. (113104)
- Some packet loss may be observed when IP multicast traffic has high fanout on a single physical port (e.g. trunk port with many receivers for the same multicast group on different VLANs), even when the total egress packet rate is well below linerate. The default egress buffer settings may not perform well under this scenario. This issue can be resolved by adding "platform arad buffering egress port multicast max-queues-full 1" to the switch config and rebooting the switch. (115343)
- MAC mirroring ACLs do not match IPv6 packets. (137537)
- On a device with with a large number of active VLANs, doing "shutdown" and "no shutdown" on many interfaces may cause the SandMcast agent to restart. (175529)
- Mirrored traffic is not subject to egress security ACLs which are applied to mirroring destination ports. (190637)
- PBR is not supported along with IPv6 in IPv6 packets. (193914)
- On 100M interfaces acting as IEEE 1588 slave ports, restarting IEEE 1588 may cause ingress timestamps to spike by as much as 100s of ms.

A workaround would be to reset the interface after restarting IEEE 1588. (218386)

- Acl and related features that use TCAM for rule matching do not work correctly on packets with IPv6 extension headers if the rules have layer 4 match criteria such as tcp flags and source/destination ports. (221161)
- PBR policies applied on interfaces do not work if the interface has some subinterfaces that also have PBR policies applied. (243286)
Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- VTEP cannot learn a remote host's MAC address from VXLAN encapsulated IPv6 Neighbor Advertisement packets received from that host. This might be an issue in a pure IPv6 deployment where the VTEP does not receive any bridged traffic from the host. In such a scenario, the traffic toward the host will be flooded in the VLAN.

The work around for this is to force Neighbor Solicitation (NS) packets

from the host for its gateway or send bridging data traffic from the host.
(246676)

- ACL/PBR/QOS will not take effect on ports with VLAN Translation configured.
(246982)
- A TCAM profile not compatible with all linecards in a system will not be installed on any linecard in that system. However, hardware tcam profile status may show erroneously that the profile is installed on compatible linecards.

Removing the incompatible linecard(s) or use a TCAM profile that is supported by all linecards in the system. (247842)

- VXLAN bridging traffic sent out of LAG interfaces may be temporarily dropped after an SSO event. (251398)
Series Affected: DCS-7500R Series

- When a DirectFlow flow is created with an action to drop the outer tag, the packet is still leaving the interface tagged. (264336)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

- When a physical port (e.g. Et1) is removed from a Port-Channel (e.g. Po1) that has subinterfaces configured (e.g. Po1.1), and there are subinterfaces configured under that physical port (e.g. Et1.1), the subinterfaces under the physical port can be incorrectly programmed. This results in traffic being dropped for the incorrectly programmed subinterfaces.

The workaround is to flap the subinterfaces affected. (269823)

- If one or more linecards in a modular system is running in tap aggregation mode, a small percentage of multicast traffic on linecards running in normal mode may be dropped, following a change to membership of some tap aggregation destination groups. (275286)
- When the primary and fallback policy are both changed in the same config session, traffic may stop matching on either primary or fallback policy.
(287146)
- When primary and fallback policies are both changed in the same config session, some packets will not be subject to either policy for a brief period during the configuration change. (287154)
- Power cycling a 7500E linecard in a chassis with 7500E fabric modules can cause packet loss on 7500R linecards. (295825)
- If VXLAN is configured, PBR and QOS policies won't apply to traffic ingressing on a MLAG peer-link. (390804)

- Sometimes Macsec Proxy interface gets into error disabled state on reload when fips is enabled.

To workaround the issue configure errdisable max flap value; i.e., "errdisable flap-setting cause link-flap max-flaps 24 time 10". Note the setting is global and affects non Macsec Proxy interfaces as well. (396589)

- A switch reload with qos policy configured may result in SandAcl agent restart. (417935)
- Storm control maximum burst size is fixed to 10kB which can lead to storm control drops for bursty incoming traffic. (423231)
- With WRR scheduling on many ports and linerate across multiple traffic classes there can be a small amount of egress buffer drops. (425123)
- While a line card is being power cycled, corrupt VXLAN encapsulated packets may be sent. (441141)
- Performing many updates to a PBR policy in a short amount of time may cause the SandAcl agent to restart unexpectedly. (441912)
- VXLAN over a MACsec proxy port as underlay may stop forwarding traffic on reload. (451061)
- VXLAN routing over MACsec proxy port stops working when a proxy port that is not bound to an external interface is updated to be bound to an external interface. (455759)
- "match nexthop-group" filter is not supported in PBR rules when the corresponding nexthop-group is referenced by ECMP routes. (458189)
- SflowAccel does not send flow samples to collectors routed via BGP routes with additional backup paths. (476553)
- Once the maximum number of supported unique PBR policies per chip is exceeded, removing the PBR policy on an interface, converting the interface to a vlan member and applying a RACL on the SVI does not work. (477753)
- With a tcam profile configuration containing "feature tunnel vxlan routing", SandAcl agent restart may lead to another unexpected restart before the agent recovers. (492320)
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- With MPLS VPN, during stateful switchover, some VPN traffic after MPLS decapsulation may be transiently forwarded in the default VRF instead of the appropriate tenant VRF. (541206)

- An over-voltage condition on the POS1V2_HBM_IO_JE0 power rail can cause the system to power-cycle. (544659)
SKUs Affected: DCS-7280CR3-96 and DCS-7280CR3K-96
- 'Persistent Internal Drop DeqDeletePktCnt' error messages are seen due to incorrect/missing system ACEs. Hitful restart of SandAcl agent is needed to resolve the issue. The trigger for this issue is unknown. (548524)
- When 'ip-length' key field is missing from the "ipsec" TCAM feature in the operational TCAM profile, all traffic on the IPsec tunnels may be trapped to the CPU, leading to low tunnel throughput.

The workaround is to enable the 'ip-length' key field in the "ipsec" TCAM feature. (588127)

SKUs Affected: DCS-7020SRG-24C2, DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

- On a shaped subinterface tx-queue, configuring the bandwidth doesn't take effect unless no-priority is also configured. (588133)
- If 'cpu traffic-policy <policyName> vrf all' is configured, where the traffic-policy contains rules with the policer action, and the operating TCAM profile does not support the CPU traffic-policy feature, then configuring any routed port in the switch may lead to an unexpected SandAcl agent restart. The workaround is to remove the 'cpu traffic-policy <policyName> vrf all' configuration. (619998)
- Postcard sampling doesn't work for traffic that is forwarded out of a sub-interface that has shaping applied. (621932)
- Postcard Telemetry sampling doesn't work for IPv6 packets if the egress interface has an IPv6 egress ACL applied on it. (624010)
- Continuous shut/no shut on a Port-Channel with 1000s of subinterfaces can cause the SandTunnel agent to restart unexpectedly. (718924)
- When using a large scale of PBR rules, removal and addition of member links used as nexthop groups destinations of PBR policies may cause the SandAcl agent to restart. (759846)
- On a config-replace, a change in the PBR config that qualifies for in-place update may still cause traffic loss. (760884)
- Performing many updates to a traffic policy in a short amount of time may cause the SandAcl and/or SandAegis agent to restart unexpectedly. (770478)
- Added support for QoS policy-map names with more than 64 characters. (791679)
- When a LAG and a member interface of that LAG have identical VLAN translation configuration, the LAG VLAN translation config may not get

applied, resulting in packets forwarding in an incorrect VLAN.

To recover, disable VLAN translation config on the LAG member, then flap the VLAN translation config on the LAG. (802669)

- Flapping the interface that is the destination for a large number of redirect actions in a traffic-policy may result in a long programming delay to updating redirect actions in hardware and to any future changes in the traffic-policy config. (814453)
- If multiple traffic policies are applied such that the TCAM usage for traffic-policy is close to scale (~100%), then the traffic-policy actions may not work if the Sand agent is restarted.

The workaround is to remove these policies from applied interfaces, and re-apply each of them. (821984)

- In rare instances, entries in a TCAM bank may fail to be deleted before the TCAM bank is released by one feature and then allocated to another feature, resulting in non-deterministic behavior in the new feature. (826437)
- Configuring ACs at scale (L2/L3 Subinterfaces, VLAN Translation, VXLAN, etc.) can result in Routed Ports getting stuck in an unprogrammed state.

To work around this issue, modify your configuration to reduce scale, then toggle the affected routed ports to "switchport", then back to "no switchport", in order to recover. (847496)

- In rare cases, it is possible when making egress interface traffic-policy changes for SandTcam to continuously crash upon policy application or removal. The workaround for this is to removal policy application from all interfaces and then re-apply the policies. (849779)
- Micro BFD packets may not hit CPU traffic-policy matches for interfaces that are put in RFC-7130 interop mode. (851466)
- Incorrectly configuring an interface without a transceiver module inserted in that interface to 800G speed using "speed 800g-8" will result in forwarding agent restarts and all links flapping. The workaround is to remove the "speed 800g-8" command from the configuration. (870787)
- Ingress VLAN translation configured on LAGs may stop working if the following conditions are met:
 1. LAG parent interface flaps or LAG member interface flaps.
 2. The VLAN has MAC learning disabled.

To recover missing VLAN translation configuration, flap member interfaces of the impacted LAG. (875253)

- When a traffic policy is applied on interfaces in different VRFs, changing the action of a match rule from unnamed counter to named counter, or changing the name of the counter causes the SandAegis agent to unexpectedly restart. The agent recovers after restart. (880583)
- In EVPN VXLAN deployments, if the SandMact agent is restarted in the presence of high traffic data rate, then after a subsequent agent restart, it may not initialize to completion. This in turn could lead to MACs advertised over EVPN not being programmed in hardware and also lead to control plane issues where interface resolution for L3 nexthops might fail. (881911)
- If the SandMact agent restarts during a flush operation, hardware L2 learning may be disabled until after the agent initializes to completion. (882394)
- When the GRE endpoint of a monitor session is reachable through multiple next hops and one of them is a port-channel, it is possible that the mirrored packet is sent to the nexthop with a wrong destination MAC address.

The workaround is to configure the port channels so that their members are in different FAPs. (886745)

- If MAC Learning is disabled on a VLAN prior to SandTunnel restart, MAC Learning may end up stuck in a disabled state on that VLAN once SandTunnel comes back up.

To recover, if MAC learning is stuck in a disabled state, flap MAC Learning configuration for the impacted VLANs (disable MAC Learning and re-enable MAC Learning). (892343)

- Under rare circumstances when a downstream LSR RsvpProtocol agent is restarted, the stale routes may not be updated resulting in mis-forwarding by the upstream LSRs.

Restarting AleL3Agent-primary will fix the issue. (898353)

- If the same L2 protocol forwarding profile is applied on two or more sub-interfaces across multiple chips, then there might be traffic loss when a stateful switchover or a restart of SandAcl agent happens. In some rare cases, this traffic loss might be persistent, in which case a subsequent restart of the SandAcl agent can be performed to recover from the problem state. (901795)

DCS-7280R and DCS-7500R Series

- In MPLS L2EVPN configuration, traffic requiring 3 or more LU labels to be pushed, may not be forwarded correctly. (275834)

- On the 7500R-8CFPX-LC linecard modules, rapid removal and reinsertion of the transceiver may result in the Denali agent restarting.

Inserting transceiver 10 seconds after removal from the same slot is the workaround. (293594)

SKUs Affected: 7500R-8CFPX-LC

- When an ingress Port Acl includes TCP or UDP port range matches and vxlan is configured on the switch, any vxlan transit packet ingressing into the switch may incorrectly match on such rules. To prevent this, key field l4ops must be excluded from the tcam profile. (477195)
- When MACsec is enabled and a SSO failover occurs, the Evora agent may restart due to "LMI Access timeout", which may cause links to flap resulting in extended traffic loss. (494072)
SKUs Affected: 7500RM-36CQ-LC
- Private VLAN may not work correctly after a stateful switchover. The workaround is to restart the SandTunnel agent. (505258)
- IQM_RST_USCNT_ERR, IQM_FR_FIMC_DB_ROLL_OVER, IQM_FR_MNMC_DB_ROLL_OVER interrupts of Jericho chip indicate traffic to and from ports of the affected chip to be dropped. A full reset of the chip could workaround this issue. (540881)

DCS-7280R2 and DCS-7500R2 Series

- The default forwarding behavior for IS-IS packets received on subinterface or vlan-tag-based pseudowire connectors is now "trap" rather than "forward".

To achieve forwarding of IS-IS packets received on these connector types, apply an L2 Forwarding Profile, as described in the TOI "L2 Protocol Forwarding". This capability for IS-IS is supported in newer releases only. (611787)

- Multicast traffic on non-default VRF will be lost for as long as 15 minutes when default VRF is multicast enabled without any L3 interfaces. Workaround is to disable mulitcast on default VRF or configure a dummy L3 interface. (619693)

7500R3, DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series

- During restart of SandAcl agent, packets may not be subjected to ACL rules for a short period of time. (372971)
- ECN bits may not be correctly copied from outer header when packets are routed after VXLAN decapsulation. (489998)

- If tap aggregation and a CPU traffic policy are configured together, either the CPU traffic policy will not apply to any traffic or the SandAcl agent will continuously restart. (492549)
- When an interface traffic-policy is changed such that hardware resources are exhausted by the change, SandAegis agent may restart. After the agent restart, the switch will recover after traffic-policy rollback kicks in. (574639)
- On applying/removing a Egress TC-to-CoS map on L2 LAG Subinterfaces, the packets going out of the Subinterfaces can have wrong dot1q tag. Workaround is to shut the subinterface, apply the map and then unshut it (582063)
- ECN bits are not preserved when packets are bridged over a VXLAN tunnel. (587486)
- When speed of a port is changed while it is carrying traffic, speed change may fail and cause interface to become inactive.

One workaround is to change the speed-group configuration of problematic port and revert it back to the previously applied speed-group configuration. Another workaround is to change the speed of primary port to 400g-8 and revert it back to the previously applied speed. (597294)

- Following a Stateful Switchover, disabling power to either supervisor using "no power enable module Supervisor" may cause some control plane agents to restart unexpectedly and repeatedly.

To resolve the agent restarts, restart the Linecard agents one at a time. (606781)

- Priority-based flow control fail to work after SSO on some instances. Workaround is to restart SandFapNi agent for the linecard after scheduling downtime. (606784)
- The system could end up in an unrecoverable power cycle loop. (773897)
- Configuring 256 Ipsec or Secure Vxlan tunnels while also configuring MACsec can result in continuous B52 agent restarts. Workaround is to reduce tunnel scale or to remove MACsec configuration. (871256)

7500R3, DCS-7130, DCS-7280CR3, DCS-7500R3 and DCS-7800R3 Series

- Systems shipped early in the life cycle of this product have been programmed with an incorrect ISL chip image for the voltage regulator monitor (VRM), which may cause HBM CRC errors and lead to packet drops.

A swix is available with an updated ISL image to fix this issue. (778115)
 SKUs Affected: DCS-7280SR3-48YC8 and DCS-7280SR3K-48YC8

- Congestion on CPU egress queue may lead to DeqDeletes on VOQs corresponding to traffic class 7 of interface mapped to BaseQPair 512. BaseQPair of an interface can be retrieved via "show platform fap mapping <interfaceName>" (901857)

DCS-7280R3A and DCS-7800R3A Series

- When MACsec is enabled in hardware on a switch, then if the switch ingress experiences congestion due to a large number of small packets at line-rate, the priority drop feature may drop packets with higher priority before packets with lower priority. (685317)
- Continuous L2 EVPN MAC churn and interfaces flaps may lead to restarts of the SandMact agent. (882608)

DCS-7020 Series

- When multiple Ipsec tunnel interfaces are shutdown, then "no shut" in a short period of time, routes associated with some tunnel interfaces may not be installed after "no shut". The workaround is to do shut/no-shut again on the individual tunnels where routes not installed. (377984)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- When traffic is flowing from a 10G interface to a 1G interface and PFC is enabled on the ingress 10G interface, if the administrator changes the COS to TC mapping all the incoming traffic on the 10G interface is dropped. The only way to recover is by restarting SandFap. (400653)

DCS-7280R3 Series

- With Jumbo MTU packets we could have some FCS errors on the peer. (747305)

7500R3, DCS-7500R3 and DCS-7800R3 Series

- In some situations, some fabric links going to a fabric chip are not used to carry fabric traffic. (450214)

DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

- A linecard addition or hotswap operation on a switch configured with "port-channel load-balance sand replication egress" may lead to continuous RQP discards and STP flaps.
A workaround is to unconfigure "port-channel load-balance sand replication egress". (353249)
Series Affected: DCS-7500R2 Series
- PBR, QOS, and BGP Flowspec policies applied on routed interfaces will not take effect if sub-interfaces are configured on that routed interface. PBR,

QOS, and BGP Flowspec policies applied on SVIs will also not take affect on traffic ingressing SVI member interfaces if VLAN mapping configuration is present. (372390)

- Ethernet over MPLS packets being terminated and decapped on a LDP pseudowire will be hashed as if the inner header is IPv4 rather than Ethernet (382763)
- With VXLAN, and PIM configuration on overlay SVI, VXLAN packets with inner IPv6 multicast packet consisting of Hop-by-Hop option may not be terminated and forwarded correctly. (484825)
- Traffic on a subinterface will be subject to BGP Flowspec matches applied to the parent interface, unless the TCAM profile has the flowspec feature configured with 'port qualifier size N', where N > 1. (541950)
- Traffic policy matches on VLAN tags may incorrectly match untagged packets by interpreting part of the untagged packet as a VLAN tag. (671474)
- If a large number of subinterfaces is configured on a parent interface which is a breakout port, and a speed change of another interface causes that breakout port to become inactive, the SandFap agent may unexpectedly restart once, flapping all links. After the SandFap agent restarts, the system will recover, the speed change will be completed, and the affected subinterfaces will no longer be configured in the hardware.

As a workaround, remove all subinterfaces from the associated breakout ports before changing the speed of a front-panel port. (691732)

- For a traffic-policy containing matches with police actions applied to a linecard's interfaces, the police action may fail to reprogram on that linecard after a linecard hot swap.

To recover, re-install the traffic policy on the interfaces where the police action failed to program. (714936)

- In an EVPN network, SandMact agent might restart due to a SIGQUIT while processing the MAC addresses polled from the hardware MAC table if the number of MAC addresses is close to the hardware MAC table limit. (718204)
- When a static MAC is configured on an MLAG port-channel, traffic may be dropped while the local port-channel is down.

The workaround is to enable the fast MLAG unicast convergence feature. (723709)

- VXLAN source VTEP pruning with VTEP-to-VTEP bridging does not work for VXLAN ARP packets. (797624)

- VXLAN source VTEP pruning with VTEP-to-VTEP bridging might be broken after some underlay route changes, such as ECMP expansion.

The workaround is to run "agent SandL3Unicast terminate" command, which could be traffic impacting depending on which features are configured. (797824)

- Ingress IPv4 and IPv6 router ACL features do not work on GRE, IP-in-IP, and GUE decapsulated IPv4 and IPv6 traffic, respectively. (799085)
- SandL3Unicast agent may unexpectedly restart under VXLAN tunnel ECMP FEC table overflow condition. (801962)
- DirectFlow cannot enforce a new destination with its actions 'output nexthop' and 'output interface' for traffic destined to the CPU. The flow's counters will increase but the traffic will not be redirected. (809100)
- An interface with a 1000BASE-SX or 1000BASE-LX transceiver inserted configured with "speed auto" connected to a peer that has auto-negotiation disabled may incorrectly link up, causing packet loss.

Workaround is to disable auto-negotiation. (829740)

- L2PF profile applied to some interfaces may not take effect after system reboot or SandFap agent restart. The affected interfaces are not displayed in the output of "show L2-protocol forwarding interface" output. Removing the L2PF config from the affected interfaces and reapplying them is a known workaround (881627)

7500R3, DCS-7130, DCS-7280R3, DCS-7280R3A, DCS-7500R3, DCS-7800R3 and DCS-7800R3A Series

- When a remote MAC is learnt on an Ingress PE with L2 EVPN MPLS configuration, the traffic destined to the remote MAC may get dropped for a brief amount of time. (551995)
- In the presence of virtual MAC with mask configuration, packets that undergo recirculation may not be properly forwarded. (559242)
- Latency based drop feature is not compatible with features using dynamic rate classes and is applicable only to green packets in traffic stream. Following set of features use dynamic rate classes:
 - 1) Configurable static tail-drops/FADT per speed, per speed+traffic-class tail drop thresholds
 - 2) WRED
 - 3) ECN

4) Interface based per TXQ and per drop precedence tail drop thresholds (618130)

- When the mapping from DSCP, CoS, or MPLS EXP to traffic-class is changed from its default by a qos map, then if the device ingress experiences congestion due to a large number of small packets at line-rate, packets with higher traffic-class with the new mapping may be dropped before packets with lower traffic-class. (626382)
- The "Last LOC Detected" and "Last LOC Cleared" timestamps shown in the output of `show cfm continuity-check end-point` command get cleared if the SandOam agent is restarted. (629830)
- "mac address virtual-router" command no longer works on 7280R3, 7500R3, 7800R3 devices. (633307)
- Egress IPv6 ACL logging do not work on packets subject to layer 3 MPLS termination. (640491)
- If the 'update interface hitless strict' configuration is active, when the FEC associated with the routes from redirect actions used by interface traffic-policies change and traffic-policy hardware usage is near or above 50% of available resources, the FEC programmed in TCAM may not be updated, causing traffic matching a redirect rule to either misroute or drop. The most common reasons for FEC changes are SandL3Unicast restart and/or SandFapNi restart.

To prevent this issue, it is recommended that traffic-policy hardware usage does not exceed 50% of the available resources. (642948)

- SandAcl agent will restart continuously if a TCAM based feature is being configured/modified on a FAP that has linerate traffic ingressing through it. For instance, bootup of linecard. Workaround is to reduce the ingress traffic on that FAP. (645398)
- When config-replace of configuration having VPLS commands is performed back-to-back in quick succession, the SandMact agent can restart unexpectedly. (690059)
- IPv6 linklocal packet received on front panel ports or sourced by CPU do not get headend replicated to remote VTEPs when IPv4 or IPv6 multicast routing is enabled on the VXLAN VLAN. (690105)
- Synthetic Loss Measurement does not work if sFlow, sampled flow tracking or mirroring is active.
Disable or remove sFlow, flow tracking or mirroring configuration. (716811)
- When ACL application fails as expected due to TCAM exhaustion and if a new member is added to a Port-channel with ACL already applied on it, the member interface may remain in blocked state even after the ACL is removed.

Need to restart SandAcl agent to recover from this state. (729108)

- Altering MEP attributes, like changing direction, changing profile etc, while CCMs are already running may result in unexpected RDI being reported in the network. A workaround is to restart the SandOam agent on the affected node, i.e. the node sending the RDI after such a churn has been made to the MEP attributes. (766982)
- With EVPN OISM, IPv6 ND link local packets received from the VXLAN tunnel are not flooded to all ports in the VLAN. (772926)
- Running 64 byte packet size RFC2544 generator throughput or frame-loss test causes subsequent RFC2544 Initiator throughput or frame-loss test's of higher packet sizes to fail.
Workaround is to run 128 byte packet size test immediately after running 64 byte packet size test or configure both 64 byte and 128 byte when using multiple packet sizes in RFC2544 generator profile. (773588)
- When removing a port-channel member from a routed port-channel with segment-security configured, traffic may leak through the other port-channel members. (773786)
- sFlow hardware acceleration may be disrupted in a forwarding chip if a linecard is removed from the system.
The workaround is to disable sFlow , restarting the SandFapNi agent for the corresponding linecard, enable sFlow. (781284)
- When an SVI is configured on the VXLAN VLAN, unknown unicast and multicast packets are unnecessarily sent to the CPU after VXLAN decapsulation. (788956)
- CFM Loss Measurement sessions running on MEPs configured on the same interface at different Maintenance Domain Levels may report incorrect losses. (793764)
- In the presence of a large number of VXLAN MAC addresses, using configure replace to change VXLAN dual stack underlay to IPv4 underlay and also changing the TCAM profile at the same time can lead to permanent traffic loss toward some remote hosts. Workaround is to shut and un-shut the SandMact agent to recover from this state. (802920)
Series Affected: DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series
- SSO can cause an egress queue to be stuck (809410)
- During boot-up, line card agent might restart. (811948)
SKUs Affected: DCS-7816-CH

- Configurations to disable MAC learning applied to trunk Port-Channels may stop working on VLANs with ACL or policy-maps applied after LAG membership changes. An MLAG peer-link should never learn MACs. If that peer-link is a Port-Channel, then it is susceptible to this issue, and may start to incorrectly learn MACs.

One workaround is to ensure the MAC learning settings for all LAG members match that of the LAG. This can be done by explicit configuration on the LAG members, even though it's dormant configuration when the member is part of a LAG. (813959)

- When MLAG shared router MAC is configured in the presence of both IPv4 and IPv6 VXLAN underlay, removing either IPv4 or IPv6 underlay might unexpectedly delete the MLAG shared router MAC entry in hardware.

The workaround is to toggle "vxlan virtual-router encapsulation mac-address mlag-system-id" command which will add back the MLAG shared router MAC entry in hardware. (814471)

- Applying a TCAM profile where both `feature traffic-policy port ipv4` and `feature traffic-policy port ipv6` contain the `set-vrf-secondary` action and at least one common packet type will cause the SandTcam agent to restart continuously. To work around this issue, remove all common packet types between these two features. (825460)
- Hardware accelerated sFlow of Vxlan decapped packet miss first 12B from header. (829460)
- Traffic flowing through the repair path of a VRF selection policy that is forwarding via an IP-in-IP decap and encap tunnel may be transiently misforwarded after either a stateful switchover of supervisor or a Layer3 forwarding agent restart. (834783)
- On a switch having an interface configured with traffic-policy unshared policer, unconfiguring the traffic-policy while the SandAegis agent is restarting may cause the SandAegis agent to unexpectedly restart one more time. (836424)
- Executing a config replace between two scaled traffic-policy configurations with unshared policer actions may fail. Additionally, the subsequent config rollback may also fail causing the switch to unapply traffic policies configured on an interface. (838906)
- When modifying interface traffic-policies where we have enough LPM hardware resources for a hitless install, but we do not have enough TCAM resources

for a hitless install, we can get into a state where we are incorrectly matching on prefixes from older policy configuration versions. (847950)

- When a switch is configured as a virtual VTEP by adding a secondary IP address on the loopback interface associated with VXLAN tunnel interface, then either after a reboot or SandAcl agent restart, duplicate packets received over the VXLAN tunnel in certain scenarios might not be dropped. Workaround is to restart the SandL3Unicast agent. (852272)
- Sub-optimal convergence may be seen in certain scenarios, with "hardware next-hop fast-failover" and BGP PIC configuration. This issue is observed only on non-BK R3-series systems using non-default MDB profiles (l3-xxl, l3-xxxl, or balanced-xl). (854811)
- Upon VLAN membership changes, policy-map, ACL, or similar features may not be applied to access ports in the VLAN. For a full list of impacted features please contact support.

Possible workarounds include: reapplying the feature to the impacted VLANs or SVIs, shut/no-shut the impacted access ports. (864384)

- Removing member interface(s) from a Port-Channel where an ingress ACL was applied, followed by ACL removal from the Port-Channel may lead to stale hardware entries. Traffic might be impacted on the member interfaces that are part of Port-Channel at the time of ACL removal.

Workaround to get device out of issue state is reload of device. (866025)

- In a VXLAN routing setup using IPv6 underlay, if the overlay SVI is configured with "ipv6 address virtual" then the virtual VTEP functionality will not work correctly. Specifically, neighbor advertisement responses destined to this SVI will fail to sync across to other L3 VTEPs and hence neighbor resolutions may fail. Also, the ND proxy functionality on AG L3 VTEPs in campus deployments will not work. (868642)
- Traffic-policies with matches on inner fields of GTPv1 packets may match incorrectly on inner L4 protocols (like TCP, UDP). (875226)
- Adding an InternalRecirc interface on a switch with a MLAG configuration will cause the Mlag agent to restart unexpectedly and repeatedly.

Remove all InternalRecirc interface configuration. (876825)

- IS-IS sessions spanning across an LDP pseudowire not using the FlexEncap CLI for connectors will not be established.

One workaround is to switch to the FlexEncap CLI configuration model for LDP

pseudowires. Another workaround is to upgrade to an EOS release where this limitation has been resolved. (882952)

- Moving an interface where an ingress ACL was applied, to a Port-Channel may lead to stale hardware entries. Traffic might be impacted on the member interface being moved to Port-Channel. (893953)
- If only egress traffic-policy is configured in the TCAM profile, but ingress traffic-policy is not, running 'show traffic-policy interface counters' for an egress traffic-policy will cause the ConfigAgent to restart. (908122)
- Under rare circumstances, unknown-unicast or broadcast or L2 multicast or L3 multicast traffic forwarding during multicast related hardware programming in the switch can result in complete traffic blackholing.
A workaround is to reset the FAP(s) on which traffic is blackholed. (915576)

DCS-7280R and DCS-7280R2 Series

- When ACLs are applied on a large number of interfaces (more than 1000 SVI interfaces), the SandHalo agent may restart with a SOCKET_CLOSE_LOCAL syslog message. (253015)
- RACLs are not supported for VXLAN traffic ingressing on a MLAG peer-link when AlgoMatch is enabled. (485677)
- For an ACL applied to a LAG subinterface when all members of the LAG subinterface on a given chip are down, subsequent edits to any ACL on that chip may cause the default configured action to be applied to all interfaces on the chip where an ACL is present.

The default configured action is to either drop or permit if the "hardware access-list update default-result permit" is present.

To recover when this bug:

- Remove the "ip access-group" configuration from LAG subinterfaces that are down
- Remove the LAG members that are down from the LAG that has subinterfaces with an ACL configuration (598255)

- When AlgoMatch is enabled and MPLS is a configured feature in the TCAM profile, mpls-in-ip packets (ip-protocol=137) may bypass ACLs configured on interfaces. (679923)
- When using AlgoMatch ACLs, certain sequences of configuring a routed port with ACL applied that is later configured to be part of a LAG with an ACL applied may cause traffic to be dropped on all ports with ACLs applied:
 1. configure routed port with an ACL
 2. configure routed LAG with no ACL applied

3. add routed port as member of routed LAG
4. configure routed LAG with an ACL

The only way to recover from this is to restart the SandHalo forwarding agent. (800795)

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7260X, DCS-7260X3 and DCS-7300X Series

- When performing reload hitless, link flaps on BASE-T ports might be observed. (152459)
Series Affected: DCS-7050TX Series
- MAC ACLs and PACLs cannot co-exist on the same interface if QoS policy ACLs exist in the switch (186949)
- The StrataVlanTopo agent may unexpectedly restart after many linecard online insertion and removal operations. (193787)
Series Affected: DCS-7300X Series
- Ethernet65, Ethernet66 may experience extra link flap when performing fast boot. (221059)
SKUs Affected: DCS-7260QX-64-F, DCS-7260QX-64-R, DCS-7260QX-64-SSD-F and DCS-7260QX-64-SSD-R
- The Strata fabric or linecard agent may restart due to fabric link flaps leading to traffic outage. (225302)
Series Affected: 7320X and DCS-7300X Series
- DirectFlow flow in table type vfp will not match malformed Arp requests. (233922)
- When traffic is redirected using PBR and when the FIB lookup based on destination IP address indicates ARP miss, the PBR redirected traffic may be rate-policed.

The workaround is to use count action for the ARP-needed class. (238084)

- When the utilization of the MAC address table grows close to the full hardware capacity, the StrataL2 agent may encounter an out of memory (OOM) condition in the system, causing the agent to get restarted. (238115)
Series Affected: 7320X, DCS-7050TX, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7260X, DCS-7260X3 and DCS-7300X Series
- If nexthop resources are exhausted as indicated by syslog
VXLAN_HWHWER_NEXTHOP_RESOURCE_FULLL, forwarding of BUM traffic to that VTEP

will be affected. (246391)

Series Affected: DCS-7060X Series

- During an SSO failover, Strata Fabric agent may not detect hardware, resulting in a fatal error reboot with increased traffic outage. (250129)

Series Affected: DCS-7300X Series

- In an L3 EVPN setup, traffic that is sourced by a VTEP that is a pure VXLAN router (i.e., it performs no VXLAN bridging and has no hosts attached to it) might incur some loss at the receiving VTEP, where the packets are decapsulated and forwarded.

Configure a custom PDP policy with the vxlan-vtep-learn class set to count to work around this issue. (251117)

- BUM traffic may temporarily loop into the MLAG pair when a member interface in the peer-link port channel is brought up. (259213)
- When IPv4 uRPF exceptions are already programmed, configuring IPv6 uRPF exceptions imposes a heavy processing load on the Strata Slice agent to do a transition from IPv4 to IPv4-v6 exception mode in the hardware. Therefore, under a scaled config scenario, the Strata Slice agent may restart repeatedly.

Removal of IPv6 exceptions will help in recovering the system if this condition is hit. (264449)

- IS-IS configured over LAGs may see neighbor adjacency flapping during a hitless restart. (276232)
- A PCI error can break counter collection for an ASIC.

Resetting the ASIC with "platform trident reset" will fix this error. (277519)

- After hitless speed change from 1G to 40G or 100G, link may remain down.

Workaround is to perform hitful restart of forwarding agent. (283175)

SKUs Affected: DCS-7050CX3-32S, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050SX3-48YC12, DCS-7050SX3-48YC12-F, DCS-7260CX3-64, DCS-7260CX3-64-F and DCS-7260CX3-64-R

- Egress ACL on SVI for BUM traffic will not match ACL rules for packets being sent across fabric interfaces. (284543)

SKUs Affected: DCS-7304 and DCS-7308

- Applying a large configuration containing many interface changes may result in Strata forwarding agent restart, which could cause traffic loss during the reconfiguration. (330840)

- When VXLAN is configured and when a route or ARP points to an ECMP group with both local and remote nexthops, then a packet destined to this ECMP group may get looped between the linecards and fabric cards resulting in link congestion and packet drops. The work around is to disable the source port based hashing by removing "ingress-interface" in "ip load-sharing trident fields" command. (368073)

Series Affected: 7320X, DCS-7260CX and DCS-7300X Series

- Irrespective of IGMP snooping being enabled or disabled, IGMP packets will not be flooded to remote VTEPs. (368106)

Series Affected: 7300X3 and DCS-7050X3 Series

- Link partners might see high number of uncorrected codewords when Reed-Solomon FEC is in use.

A workaround is to run 'shutdown' followed by 'no shutdown' on the affected interfaces. (370583)

Series Affected: DCS-7060X2 and DCS-7260X3 Series

- After switching scheduler mode, the StrataCounters agent might be unable to complete a hitless reload shutdown stage, which will cause a hitless reload to be hitful. This can be pre-identified by seeing "SBus ack error in schan command" in the StrataCounters agent log. Workaround is to reset the ASIC before running reload hitless. (388694)

Series Affected: DCS-7050X3, DCS-7060X, DCS-7060X2 and DCS-7260X3 Series

- Having any ACLs with UDF qualifiers as well as any of the following LAG/ECMP hashing features present in the configuration: RRoce hashing, flexible hashing or TTL/HopLimit hashing, will result in incorrect operation of the LAG/ECMP hashing where the switch incorrectly hashes on IP payload. To avoid this problem the conflicting features should not be configured at the same time. (403182)

Series Affected: DCS-7060X, DCS-7060X2, DCS-7260X3 and DCS-7300X Series

- Changing PFC Watchdog configuration while traffic is flowing can result in sub-second traffic outage when PFC Watchdog is in non-disruptive mode and multi-second traffic outage when PFC Watchdog is in disruptive mode. (403790)
- DirectFlow flows with mirroring action may not be reprogrammed after a supervisor switchover or StrataMirror agent restart.

The workaround is to restart the Strata agent, which restores the flows. (406857)

- Restart of StrataL2 may lead to continuous StrataL2 restarts when there is mismatch of VTEP IPs for recovered MAC addresses.

Workaround: reload the system. (408138)

- DirectFlow flows do not take precedence over conflicting router ACLs. (408734)
Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX2, DCS-7050SX3 and DCS-7050X3 Series
- DirectFlow flows can be uninstalled from a linecard on reaching the idle timeout for the flow even though matching traffic continues to match the flow on another linecard. The flow can be restored by re-programming it or disabling the idle timeout.

Series Affected: 7300X3 (411851)

Series Affected: 7300X3, 7320X and DCS-7300X Series

- In an MLAG setup, under certain conditions, packets entering a LAG interface on one MLAG peer destined to the other peer might end up getting dropped on the MLAG peer link.

A workaround for this is to restart the StrataLag agent. (412307)

- Hosts may incorrectly learn the ARP of virtual IP behind the switch MAC when the switch is reloaded or configuration is flapped.
Workaround is to clear the ARP on the host to trigger a new ARP reply with the correct VARP MAC. (418626)
- Configuring more than 16K MPLS pop/swap labels will cause continuous restart of forwarding agent. (419243)
- If the number of next-hop groups configured exceeds maximum supported hardware NHGs, then STRATA-3-MPLS_MAC_DA_PROFILE_TABLE_FULL syslog message is logged and L3 forwarding agent may continuously restart.

The workaround is to scale down the number of next-hop groups configured to stop the agent from continuously restarting. (420234)

- Ethernet49-54 may not link up with peers over long copper cables.
Workaround is to enable autoneg on both ends of the link with 'speed auto <speed>' command. (421121)
- After a speed change, flows with mirroring actions on affected interfaces may not be mirrored as expected. Strata Slice agent can be restarted to restore expected behavior. (422221)

SKUs Affected: DCS-7060SX2-48YC6-F and DCS-7060SX2-48YC6-R

Series Affected: CCS-720XP, DCS-7050SX3 and DCS-7260CX3 Series

- The forwarding agent may continuously restart if the speed is changed from 100G to 50G when a 50G QSFP transceiver is inserted.
This only affects 2-lane 50G transceivers, and does not affect 100G transceivers operating in 50G mode.
A workaround is to configure the speed before inserting the transceiver.

The forwarding agent restarts can be stopped by temporarily configuring the affected interfaces in 25G mode. (424990)

- With scale policy-map configuration, Strata agent may restart continuously with SIGQUIT. (430654)
- During Leaf-SSU or SSO, interfaces with 100GBASE-SRBD transceivers may experience a single link flap. (440451)
- Changing Acl configuration rapidly and multiple times may cause the forwarding agent to repeatedly restart unexpectedly. (447584)
- Slice agent may restart unexpectedly on performing speed change and/or shutting/no shutting on a PFC enabled interface with traffic flowing out of the interface.

The issue will not be seen if PFC is not configured or traffic is not flowing. (462139)

Series Affected: 7368 Series

SKUs Affected: DCS-7060DX4-32-F and DCS-7060PX4-32-F

- With PFC pause stream for lossless priorities on egress interface and data traffic destined to the egress interface for lossy and lossless priorities, few seconds of loss may be seen on lossy priority traffic on toggling PFC globally.

Workaround: Stopping PFC pause frames before toggling PFC globally. (477941)

- Aggregate storm control rate limit is not honored for broadcast and multicast traffic. (490511)

Series Affected: DCS-7060X, DCS-7060X2, DCS-7260X3 and DCS-7300X Series

- If an MBA supplicant is not assigned a dynamic VLAN when it is authenticated on a port, and if the supplicant later moves to a new port which is in unauthorized state, MBA authentication on the new port will not succeed and the port will continue to stay in unauthorized state. (491230)

Series Affected: 7300X3, AS7326, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, DCS-7050SX3, DCS-7060CX, DCS-7060SX2, DCS-7260CX and DCS-7260CX3 Series

SKUs Affected: 7304X3-FM2, 7308X3-FM2, 7320X-32C-LC, 7324X-FM, 7328X-FM, DCS-7010TX-48, DCS-7050CX3-32S, DCS-7304 and DCS-7308

- Programming or updating a Router ACL may result in forwarding plane slice agent restart and can leave the corresponding vlan permanently blocked. (498502)

- MLAG connectivity between the peers might go down after linecard OIR. Flapping MLAG peer-link interfaces will bring the MLAG connectivity back up (500162)

- With "switchport default phone vlan <vlan-id>" configuration, DHCP offer packets tagged with the default phone vlan tag might be sent out untagged on trunk ports (501749)
- It is possible that storm-control configuration will still be applied in hardware after removing the configuration. Workaround is to restart the Strata agent. (504901)
- Enabling per VLAN routing with an L3 interface that only has an IPV6 address will result in no routing of packets received on this L3 interface. Workaround is to flap the interface. (528490)
- A different set of routed interfaces may operate after reload or StrataL3 process restart if the supported scale of routed interfaces is exceeded. (533468)
 SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DP-24S, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-48S, CCS-720DP-48S-2F, CCS-720DP-48S-F, CCS-720DT-24S, CCS-720DT-24S-2F, CCS-720DT-24S-F, CCS-720DT-48S-2F, CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8, CCS-722XPM-48ZY8-F and DCS-7010TX-48
- After a switch reload, L2 unicast traffic with a known destination MAC address may flood the VLAN.

A workaround is to remove the MAC address and re-add it. (539560)

Series Affected: DCS-7050SX3 Series

- Changing the speed of a QSFP or SFP interface from 10G to 1G may cause a Strata forwarding agent restart under some circumstances. This will cause an interruption in traffic on all ports. Workaround is to change from 10G to another speed before changing to 1G. (550269)
 Series Affected: 7300X3, CCS-720XP, DCS-7050SX3 and DCS-7260CX3 Series
- TCAM resource exhaustion can lead to VXLAN forwarding to fail.
 Workaround: Reduce TCAM resource utilization to allow VXLAN entries in TCAM by unconfiguring other features that occupy TCAM. (561231)
- When "show platform trident vxlan multihoming groups non-peering" is executed, StrataL3 agent may restart unexpectedly. (561420)
- Fabric interface on a linecard can sometimes end up in a state where it can discard all packets at egress.
 When this condition happens, the "show platform trident fabric counters queue detail" command will show that the drops increment while there is no increase in packets or bytes transmitted on the problematic fabric interface. To recover from the problem, it is necessary to restart the linecard forwarding plane agent by running "platform trident <linecard> reset". The recovery process can cause traffic disruption across several

ports in the switch. (562860)

SKUs Affected: DCS-7304 and DCS-7308

- Committing security access lists through CVP may restart forwarding plane slice agent. Post restart forwarding plane functionality will continue to work as is.

Workaround is to use "hardware access-list update default-result permit" CLI. (571289)

- After performing SSU from a PTP unsupported release to a PTP supported release, PTP will not be functional.
The workaround is to reset the forwarding switch by issuing the CLI command "platform trident reset" after the SSU and prior to enabling PTP. This will cause all interfaces to flap and a brief traffic outage to occur. (582970)

Series Affected: DCS-7060X4 Series

- When many VLANs and/or LAGs are configured, port channels may flap during SSU. (592390)
- PIM over GRE interface does not support over 1K multicast routes (605368)
- Performing SSO with increased route scale while Segment Security is configured may result in a fatal error reboot and extended traffic outage. (623639)

- After an interface flap, multicast traffic may not forward correctly in data plane.

Workaround is to remove and re-apply multicast configuration. (627318)

Series Affected: CCS-710P and DCS-7010TX Series

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- StrataL3 forwarding agent may restart unexpectedly when many L3 interfaces are removed and then added again. (636201)
- Migrating from static VXLAN to EVPN VXLAN configuration may result in stale remote MAC entries in hardware leading to black holing of traffic destined to those MAC addresses.

Workaround is to clear the MAC address table entries for those VLANs. (644589)

- If StrataL3 is restarted when there are insufficient hardware resources to add a virtual port for a VXLAN remote VTEP as indicated by the STRATA-3-VIRTUALPORT_RESOURCE_FULL syslog, StrataL3 agent may not become warm resulting continuous traffic drops.

Work around is to reduce the number of VXLAN remote VTEPs.

Workaround: Reduce VTEP scale so that all VTEPs have an allocated virtual port before restarting StrataL3 to recover. (646694)

- Adding and removing "no switchport mac address learning" command while MAC address flaps between local and remote can trigger stale MAC entry left in MAC address table even after clearing all MAC addresses. (664960)
- When ECN is configured globally, non-ECT packets could be dropped if the total buffer occupancy exceeds the global threshold, even if there is sufficient memory available in reserved space. (666378)
- When "hardware counter feature vtep encap" CLI is configured, the deletion of VTEPs can trigger an unexpected restart of the StrataL3 agent. (678643)
- Use of "no mac address learning" command with VXLAN configuration may result in mis forwarding of the VXLAN encapsulated traffic received.

Workaround is to remove "no mac address learning" configuration on VXLAN VLANs. (681651)

Series Affected: DCS-7050QX, DCS-7050TX, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7300X Series

- After a Leaf Smart Software Upgrade (Leaf SSU), Strata agent might restart unexpectedly, resulting in increased traffic outage. In some cases, this might also cause an SSU stage timeout without Strata agent restart, resulting in a fatal error reboot. (689713)

SKUs Affected: DCS-7060DX4-32 and DCS-7060PX4-32

- Removing an SFP BASE-T transceiver from Ethernet33 or Ethernet34 may cause the forwarding agent to restart and a brief traffic outage to occur. (715516)

Series Affected: DCS-7050X4 and DCS-7060X4 Series

- During a link toggle or speed change, Strata forwarding plane agent may unexpectedly restart due to "OVERFLOW IRQ from ISEC_PDF_STATUS or ISEC_IDF_STATUS".

Switch is expected to get back to working state after the restart. (717855)

Series Affected: CCS-750 Series

- If a VRF is dissociated from a routed port that also has IPv6 enabled through "ipv6 enable", the StrataL3 agent can restart unexpectedly.

Use "no ipv6 enable" under that routed port before dissociating the VRF from it, followed by "ipv6 enable". (719039)

- If MLAG configuration is applied when there is congestion, MLAG stays in inactive state and StrataVlanTopo might restart. When the congestion is stopped MLAG becomes active (727116)

- StrataLag agent will continuously restart unexpectedly if there are more LAGs configured (including LAGs with no members) than the hardware can support.

The LAG hardware limits can be seen by executing CLI 'show port-channel limits', field 'Max port-channels per group'.

Workaround:

- remove configured LAGs until the number of LAGs is not exceeding the hardware capacity. (732733)

- An interface configured to be a member of a LAG may fail to be programmed in the hardware.

The CLI 'show port-channel detailed' will show the status of the interface as "Waiting for hardware to program port to RX...".

Workaround:

Remove and add the affected member interface from/to the LAG (738580)

- Transitioning to data plane learning from EVPN control plane may result in remote MAC addresses not being learnt.

Workaround is to flap the VXLAN interface. (748461)

- 5GBASE-T ports may always report linkup regardless of the actual link state, when configured with forced 100M speed using CLI command `speed 100full`.

Workaround is to use BASE-T auto-negotiation with capable link partners instead of forcing the speed to 100M. (759421)

SKUs Affected: CCS-720XP-96ZC2

- When the switch is configured in ALPM mode and when the default route is forwarded over a VXLAN tunnel, L3 agent forwarding restart may result in some routes not getting programmed in the hardware. There is no workaround for this issue (764586)

- During config replace, StrataL3 forwarding agent may restart unexpectedly when the interface moves from one VRF to another. (776041)

- MAC initialization failure may cause Strata forwarding plane agent to unexpectedly restart during stateful switchover, and this may cause stateful switchover to fail, and reboot the switch. (776821)

Series Affected: 7300X3 Series

- When a recirc channel is present without VXLAN configured, another port-channel may stop forwarding traffic after a hitless speed change.

Restarting the StrataVlanTopo agent fixes the issue. (780110)

Series Affected: DCS-7060CX Series

SKUs Affected: DCS-7060CX2-32S, DCS-7060SX2-48YC6, DCS-7260CX-64, DCS-7260CX3-64 and DCS-7260QX-64

- IP PACL configured on interface receiving VXLAN traffic encapsulated to VARP VTEP IP may result in drops.
Workaround is to configure port range rule in the ACL or apply the following command: ``platform trident tcam port-acl vfp disabled``. (784384)
- Looping or double delivery of L2 protocol packets received on the peer-link may be observed in an active-active MLAG setup when L2 Protocol Forwarding is configured on the peer-link. (784857)
- When an interface having "switchport trunk private-vlan secondary" configuration is removed from a port-channel we may observe incorrect tagging on packets egressing the interface.
A workaround is to remove and add the "switchport trunk private-vlan secondary" configuration. (786324)
- Groups of four adjacent BASE-T ports share a common hardware resource; for example, Et1-Et4 and Et5-Et8 and so on. Due to an issue with the shared hardware resource, all 4 ports may fail to link up, even if the link partners report a linkup.

On EOS versions prior to 4.29.0F, on fixed systems, the workaround is to run ``platform trident reset``; on modular systems, the workaround is to run ``platform trident <linecard> reset`` on the affected linecard. On EOS versions 4.29.0F and onwards, the workaround is to run ``reset platform trident phy interface INTF pcs`` on one of the affected interfaces. (792495)

SKUs Affected: CCS-720XP-24Y6, CCS-720XP-48Y6, CCS-750X-48THP-LC and CCS-750X-48TP-LC

- Groups of four adjacent BASE-T ports in the range Et1-Et24 share a common hardware resource; for example, Et1-Et4 and Et5-Et8 and so on. Due to an issue with the shared hardware resource, all 4 ports may fail to link up, even if the link partners report a linkup.

On EOS versions prior to 4.29.0F, on fixed systems, the workaround is to run ``platform trident reset``; on modular systems, the workaround is to run ``platform trident <linecard> reset`` on the affected linecard. On EOS versions 4.29.0F and onwards, the workaround is to run ``reset platform trident phy interface INTF pcs`` on one of the affected interfaces. (792506)

SKUs Affected: CCS-720DP-48S, CCS-722XPM-48Y4, DCS-7010TX-48 and DCS-7010TX-48-DC

- If a mirroring session is configured with a VLAN source as well as a session ACL, and then if the ACL on the session is changed, removed, or if the entire session is unconfigured, the StrataMirror and/or Strata-FixedSystem agent may unexpectedly restart. (794048)

- IPv6 Underlay is not supported. Receiving EVPN routes from IPV6 VTEPs may result in continuous restart of L3 forwarding agent. (794373)
Series Affected: 7320X, 7388, DCS-7050QX, DCS-7050QX2, DCS-7050SX, DCS-7050SX2, DCS-7050TX, DCS-7050TX2, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series
- When Reducing number of members in ECMP groups while running traffic with DLB enabled can cause some packets drop.
Workaround is to restart the traffic or disable/enable DLB. (797519)
SKUs Affected: 7358X4-SC, DCS-7050DX4-32S, DCS-7050DX4-32S-F, DCS-7050PX4-32S and DCS-7050PX4-32S-F
- Active supervisor removal may cause the PlxPcie-system to unexpectedly restart and, in some circumstances, may prevent the uplink ports on the supervisor card from being detected.
Workarounds: reinsert the card again or reload the device. (807283)
SKUs Affected: CCS-755-CH and CCS-758-CH
- VLAN translation may stop working on a port when:
 - the port was removed from a LAG
 - that LAG has the same VLAN translation configured as the port
 Workaround:
 - restart StrataL2 agent (809637)
- If there are interfaces with transceiver inserted but repeatedly flapping or unable to link up due to poor signal, high CPU usage of the Strata slice agent may be experienced.

To workaround this, shut down the interfaces or remove the transceivers. (811800)
- BessMgr agent CPU usage may be high when underlay multicast and sFlow are configured. (811844)
- If a MAC flaps between a port-security protect mode interface and another interface rapidly, it may be displayed on the wrong interface in 'show mac address-table' and might not be counted in number of allowed addresses on the port-security interface (814589)
- When the switch is configured with lot of VXLAN VLANs and trunk ports, changing the VXLAN loopback interface's IP or flapping the loopback interface may lead to VXLAN DIP termination entry not programmed in hardware resulting in dropping all the VXLAN encapsulated traffic.

Workaround: Reload of the device or shutdown edge interfaces before flapping VXLAN loopback interface. (819174)
Series Affected: DCS-7050QX, DCS-7050SX, DCS-7050TX, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series

- Pause Flow control might not work after Autoneg resolution.

Workaround :

Reconfigure Pause Flow control.

1. "no flowcontrol send"
2. "flowcontrol send on" (820102)

Series Affected: 7388 Series

SKUs Affected: 7358X4-SC, 7388X5-SC and DCS-7060DX5-64

- If VLAN translation is configured on a port-security protect mode port and the StrataL2 agent is restarted, some MACs learned on that port may be able to move and age.

A workaround is to remove and re-add the port-security configuration for that port. (821390)

- After an invalid MMU queue interface-profile is applied, the StrataMmu agent may keep restarting repeatedly after any of the below:

- a) StrataMmu agent is restarted
- b) Device is reloaded

A workaround is remove the invalid MMU interface-profile (834515)

Series Affected: 7388 Series

SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4-32S-F, DCS-7050PX4-32S, DCS-7050PX4-32S-F, DCS-7060DX5-64, DCS-7060DX5-64-F, DCS-7060DX5-64S, DCS-7060DX5-64S-F, DCS-7060DX5-64S-R, DCS-7060PX5-64E and DCS-7060PX5-64E-F

- L2 traffic to peer static and configured static MAC addresses may be broadcasted in the VLAN after that VLAN is removed and reconfigured. (841385)

Series Affected: CCS-710P, CCS-722XPM and DCS-7010TX Series

SKUs Affected: CCS-720DP-24S, CCS-720DP-24S-F, CCS-720DP-48S, CCS-720DP-48S-2F and CCS-720DP-48S-F

- After an invalid MMU queue interface-profile is applied, the forwarding plane agent may keep restarting repeatedly after any of the below:

- a) Forwarding plane agent is restarted
- b) Device is reloaded

A workaround is remove the invalid MMU interface-profile (843327)

Series Affected: 7368 Series

SKUs Affected: AS7326-56X, AS7726-32X, CCS-710P-12, CCS-710P-16P, CCS-720DF-48Y-2F, CCS-720DF-48Y-F, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-48S-2F, CCS-720DP-48S-F, CCS-720DT-24S-2F, CCS-720DT-24S-F, CCS-720XP-24Y6-F, CCS-720XP-24Y6-R, CCS-720XP-24ZY4-F, CCS-720XP-48Y6-F, CCS-720XP-48Y6-R, CCS-720XP-48ZC2-F, CCS-720XP-96ZC2-F, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8-F, CCS-755-CH, CCS-758-CH, DCS-7010TX-48-DC-F, DCS-7010TX-48-DC-R, DCS-7010TX-48-F, DCS-7010TX-48-R, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050CX3-32S-SSD-F, DCS-7050CX3-32S-SSD-R,

DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7050DX4-32S-F, DCS-7050SX3-48C8-F, DCS-7050SX3-48C8-R, DCS-7050SX3-48YC12-F, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050SX3-96YC8-F, DCS-7060CX-32C-F, DCS-7060CX-32C-R, DCS-7060CX-32S-F, DCS-7060CX-32S-F (HwRev:13.20), DCS-7060CX-32S-R, DCS-7060CX-32S-R (HwRev:13.20), DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060PX4-32-F, DCS-7060SX2-48YC6-F, DCS-7060SX2-48YC6-R, DCS-7260CX-64-F, DCS-7260CX-64-R, DCS-7260CX-64-SSD-F, DCS-7260CX-64-SSD-R, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F, DCS-7260CX3-64E-R, DCS-7260CX3-64LQ-F, DCS-7260CX3-64LQ-R, DCS-7304 and DCS-7308

- If many continuous MAC moves occur in a short period of time, the device may not be accessible via SSH while the MAC moves continue. (845283)
- A chip reset may occur as a result of a fatal parity error reported in the MMU egress buffer block if the system is configured to exceed the hardware replication limit of 1024 replications per packet. (850455)
Series Affected: 7358, 7368, 7388, DCS-7050X4, DCS-7060X4 and DCS-7060X5 Series
- Support for the 64-bit image binary of EOS on this platform will be removed in EOS-4.32.0F. Customers should plan to migrate to a 32-bit image binary of EOS if presently deployed with the 64-bit image binary of EOS. Please refer to Field Notice 78 for additional details. (880630)
SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DF-48Y-2F, CCS-720DF-48Y-F, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-48S-2F, CCS-720DP-48S-F, CCS-720DT-24S-2F, CCS-720DT-24S-F, CCS-720XP-24Y6-F, CCS-720XP-24Y6-R, CCS-720XP-24ZY4-F, CCS-720XP-48Y6-F, CCS-720XP-48Y6-R, CCS-720XP-48ZC2-F, CCS-720XP-96ZC2-F, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8-F, DCS-7010TX-48-DC-F, DCS-7010TX-48-DC-R, DCS-7010TX-48-F and DCS-7010TX-48-R
- Configuring and removing static MAC address entries for MAC addresses that have already been learnt may result in high CPU usage by the StrataL2 agent. (886655)
- Enabling LANZ and PFC simultaneously may lead to CPU resource contention causing a slow down in Strata agents such as StrataL2 and StrataCounters. (893364)
Series Affected: 7368, 7388 and DCS-7060PX4 Series
SKUs Affected: 7368X4-SC, DCS-7060DX5-64, DCS-7060PX5-64E and DCS-7060PX5-64E-F
- "show qos interfaces wred counters queue" command may lead to ConfigAgent agent restart after interface flap . Reload the Device would fix this issue. (905912)
Series Affected: 7368 Series
SKUs Affected: 7358X4-SC, CCS-755-CH, CCS-755-X3-SC, CCS-758-CH, CCS-758-X3-SC, DCS-7050DX4-32S, DCS-7050DX4-32S-F, DCS-7050PX4-32S, DCS-7050PX4-32S-F, DCS-7060DX4-32, DCS-7060DX4-32-F, DCS-7060DX4-32S, DCS-7060DX4-32S-F and DCS-7060DX4-32S-MON

CCS-750 Series

- Switchcard switchover may disrupt PTP time synchronization and cause PTP to go into holdover state temporarily. PTP will recover and exit holdover state when time synchronization is re-acquired without any intervention. (728467)

7300X3, 7320X, DCS-7060X and DCS-7260X Series

- Packet buffer memory corruption can result in packets getting discarded without visibility in counters.
When this condition occurs, the buffer usage reported by the "show platform trident mmu queue status" command exceeds the total number of buffers reported by the "show platform trident mmu buffers" command.
A workaround is to run the "platform trident reset" command, which results in forwarding agent restart causing the links of the system to flap momentarily. (553299)

DCS-7260X3 Series

- Forwarding agent may restart when a 40GBASE-SRBD transceiver is inserted after the hitless reload. To work around, configure "speed forced 40gfull" on the desired interface prior to inserting the transceiver. (417311)
- Transceiver insertion or removal from a port used for recirculation in VXLAN routing may result in continuous traffic loss. Workaround is to flap the port by doing shut/no shut. on the interface. (590389)
- When egress port mirroring to CPU is enabled, ingress packets of size 143 or 144 that undergo VXLAN encapsulation may cause parity errors in the egress data buffer.

As a workaround, disable all egress port mirroring to CPU sessions. (898804)

7368 and DCS-7060X4 Series

- 802.3X pause is not supported (348758)
- If speed of a member interface of a port channel belonging to a VLAN is changed, then it can cause packet drops on that interface. Workaround is to unconfigure and reconfigure that port channel or kill L3 agent. (380019)
Series Affected: 7368 Series
SKUs Affected: 7368-16C, 7368-16S, 7368-4D and 7368-4P
- MPLS traffic will be dropped if the nexthop MAC address ages out or not known.

The workaround is to set the MAC address aging-time to a larger value or simply turn off MAC aging if possible. (389425)

- Speed change on a port-channel member can cause a flap in BFD and BGP sessions for the entire port-channel (520638)
- Changing the speed of an interface with priority flow control enabled may result in a forwarding agent restart. (860732)

DCS-7060X5 Series

- Interfaces with copper (CR) transceivers may require auto-negotiation to reliably link up at all rates except 10G and 40G. Auto-negotiation can be enabled using the "speed auto <speed>" interface configuration command. (736210)
SKUs Affected: DCS-7060DX5-64 and DCS-7060PX5-64E
- The forwarding agent may restart when attempting to run "show interfaces phy detail". (756995)
- FEC advertisement may be incorrect for 25G and 50G-2 auto-negotiation. The workaround is to perform shutdown followed by no shutdown on affected interfaces. (781470)
SKUs Affected: DCS-7060DX5-64-F and DCS-7060PX5-64E-F

7388 and DCS-7060X5 Series

- StrataMmu agent has high CPU usage when PFC watchdog is configured with a low polling interval, such as 0.010. (636394)
- With VXLAN and CPU TX mirroring configured, an uncorrectable parity error, not associated with any memory or register, may corrupt CPU bound control plane traffic.

As a workaround, CPU TX mirroring can be deconfigured, and then reset the forwarding ASIC via "platform trident reset" CLI command. (770294)

CCS-710P, CCS-720DP and CCS-720DT Series

- Scaled IP Locking configuration may result in an unexpected restart of Strata forwarding agent. (678926)
SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DP-24S, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DT-24S, CCS-720DT-24S-2F and CCS-720DT-24S-F
- If more than 2 VFP groups requiring counters are configured, Strata slice will continuously restart.
Workaround is to unconfigure features requiring VFP until below the limit. (884210)
Series Affected: CCS-710P, CCS-720DP and CCS-720DT Series

CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series

- Egress mirroring does not encrypt mirrored packets when MACsec is enabled on the destination port. (769793)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- DirectFlow is not supported on MACsec enabled interfaces (861705)
Series Affected: CCS-722XPM Series
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- ARP and IPv6 neighbours may be in unresolved state over VXLAN tunnels when MACsec is enabled on core facing interfaces. The workaround is to disable MACsec on core facing interfaces. (897240)
Series Affected: CCS-722XPM Series

CCS-720DF and CCS-720XP Series

- Setup with only Storm-control and scaled IP Locking configuration may cause slice agent to restart continuously. (599935)
- A single event upset (SEU) in a TCAM after a SSU upgrade may not be detected. This may cause traffic to be misclassified. A forwarding agent restart is required to forward traffic normally. (621616)

CCS-750 Series

- When report flooding is disabled in IGMP Snooping configuration, report packets may still be flooded on port-channel member interfaces.

The workaround is to disable and re-enable IGMP snooping. (534531)
- The StrataCes Agent may unexpectedly restart when a linecard is hot swapped. (568170)
- After a hitless restart of the Strata-Switchcard agent, configuring global ECN fails with "ECN Profiles Exhausted message".

To work around the issue, restart "Strata-Switchcard" agents using the following commands:

```
platform trident Strata-Switchcard1 reset
platform trident Strata-Switchcard2 reset
```

The operation is hitful and temporarily disrupts traffic across all the interfaces of the switch. (618006)

- Linecard slice agents may restart after supervisor switchover. The system recovers normally. There is no workaround. (623381)

- SSO with L3 VXLAN may lead to extended outage greater than 1 second. (637602)
- Inserting secondary switch card may result in unexpected restart of StrataL3 forwarding agent. (639631)
- When packets are flooded on a VLAN and the ingress interface is a downlink, the ingress interface will have lower egress bandwidth for other traffic flows even though the flooded traffic does not egress the ingress interface. (649424)
- Hotswapping the active supervisor, and then issuing "redundancy manual switchover" will cause stateful switchover to fail. (718863)
Series Affected: CCS-750 Series
- A PCIe error during hardware removal may result in a kernel panic. (734166)
- Performing SSO may result in extended traffic loss. (748533)
Series Affected: CCS-750 Series
- Feature agent StrataL2 would continuously restart when the number of unique multicast replication list goes higher than 16k. Workaround is to reduce the multicast scale (784949)
- In some rare cases, after a stateful switchover, the switch could start discarding all packets. To recover from the problem state, "platform trident Switchcard1/0 reset" and "platform trident Switchcard2/0 reset" must be executed." (787513)
- When the core VLAN of a VXLAN VLAN has a LAG and needs to flood edge-to-core traffic to all of its ports, the addition of another non-LAG port to that core VLAN results in traffic not egressing the LAG. The removal of the LAG and the addition of another port to the core VLAN results in traffic not egressing this new port.

Restart the slice agent to recover from this condition. (828814)

- Known unicast traffic may not flood to all ports of a core VLAN when the underlay MAC of the nexthop to the remote VTEP hasn't been learnt.

Flap the VXLAN interface or restart the StrataL3 agent to recover. (838183)

- after a hitful reset on Trident3X4 platform, vxlan encap in the case of v6 vxlan underlay would fail. (855762)
- When a linecard is removed and re-inserted, the Strata-SwitchcardCes agent may restart, disrupting traffic across all ports on the switch. (862911)
- If the same egress ACL is configured across a large number of ports, the Strata switchcard agent may continuously restart.
Workaround is to configure the egress ACL across a fewer number of ports.

Alternatively a duplicate egress ACL can be created and half the ports can be configured with the original egress ACL, and the other half can be configured with the duplicate egress ACL. (880861)

- If active supervisor is removed from the chassis with traffic flowing through the supervisor's uplink interfaces, traffic may not converge to standby supervisor's uplink interfaces for 100 seconds.

Workaround is to remove the active supervisor's uplink interfaces, waiting 5 seconds, and then removing the active supervisor from the slot. (892286)
Series Affected: CCS-750 Series

DCS-7050X3 Series

- Changing interface speed can cause all subsequent link up and link down notifications to be delayed by approximately one second on all ports. The workaround is to restart the forwarding agent after the speed change using the CLI command "platform trident reset". This causes all ports to flap and a brief traffic outage to occur. (810073)

7300X3, AS7326, AS7726 and DCS-7050X3 Series

- On interfaces with 40GBASE-SRBD transceivers operating at their default speed and "no transceiver qsfp default-mode 4x10G" configured, SSU may cause more than 100 seconds outage.

The workaround is to configure such interfaces with 40GBASE-SRBD transceivers to "speed forced 40g" (761971)

- Momentary MAC Rx local faults and remote faults may not be reflected in MAC Rx fault counters and may not cause the interface state to flap. (880612)
SKUs Affected: 7300X3-32C-LC, 7300X3-48TC4-LC and 7300X3-48YC4-LC
- SSU from a release < EOS-4.24.0F may fail with a fatal error if interfaces on a single QSFP port are configured with speeds 50G and 10G, 25G or default. Such interfaces are identifiable by the speed-misconfigured errdisabled state. The workaround is to configure matching speeds on all interfaces in a QSFP port. (891375)

7358 and DCS-7050X4 Series

- When more than 500 access list rules are configured and applied to range of interfaces, slice agent may encounter SIGQUIT and restart when shut and no shut on the range of interfaces.

Workaround is to reduce number of access list rules. (679292)
Series Affected: 7368 Series

- StrataL3Unicast agent restart may result in the flapping of Layer3 sub interfaces. (732310)
- After parity error occurs in LTS_TCAM memories, traffic forwarding may not work as expected.
Workaround is to restart the Strata-FixedSystem agent after the parity error. (756038)

7388 and DCS-7060X5 Series

- When StrataL2 agent restarts, some static MAC addresses pointing to port-channel may not get programmed in hardware table resulting in flooding.
Workaround is to removing and add the static MAC address configuration. (787583)
- With the scale of VXLAN VLANs and flood VTEPs, churn in flood VTEP list may result in exhaustion of replication group used for flooding.
Workaround is to flap the VXLAN interface. (824256)
- Changing VXLAN source-interface under "interface vxlan" may not take effect resulting in sending the VXLAN encapsulated packets with incorrect source IP address. Workaround is to flap the VXLAN interface. (824839)
- CPU monitoring sessions may cause parity errors in the egress data buffer when VXLAN traffic is traversing the switch.

As a workaround, disable all CPU monitoring sessions. (858716)

Series Affected: 7388 Series

SKUs Affected: 7388-16CD, 7388-16CD2, 7388-8D, 7388-SUP, 7388-SUP-D, 7388X5-SC, DCS-7060DX5-64-F and DCS-7060PX5-64E-F

- CPU bound traffic may get dropped in the presence of multiple broadcast / multicast / unknown unicast streams with high traffic rate. (862954)
- StrataL3 agent may restart unexpectedly when underlay reachability for a remote VXLAN VTEP changes. (863256)
- Some VXLAN flows may get encapsulated with incorrect Destination VTEP IP Address after underlay reachability/route changes.

Workaround: Restart StrataL3 agent or reload the device. (913104)

Transceiver Series

- When performing an accelerated switch upgrade with the command reload fast-boot, some transceivers may experience a link interruption. This can occur on dual-speed QSFP100 optical modules that support both 100G and 40G, specifically when the module is configured to operate at 2x50G. Examples include the 100GBASE-PSM4 and the 100GBASE-CWDM4. (343684)

- 100GBASE-SR4 modules may not link up or experience errors when inserted into a QSFP-DD slot or an OSFP slot using an OSFP to QSFP Adapter (OQA) on some platforms.

The workaround is to override the transceiver tuning with the command 'transceiver electrical lane 1-4 rx-output-amplitude module-default rx-output-pre-emphasis module-default tx-input-equalization module-default' (631109)

SKUs Affected: 100GBASE-SR4

- After ASU2 upgrade, some QSFP power class 8 transceivers (including 100GBASE-ZR4) may stop transmitting.

Operational functionality of the transceiver can be restored by issuing interface configuration command 'transceiver diag simulate removed' followed by 'no transceiver diag simulate removed' (755695)

- SFP-10G-MRA-T 10GBASE-T ports configured with 'speed auto 100full' or 'speed 100full' (100M port speed) may experience temporary traffic loss during the Leaf Smart Software Upgrade (SSU) boot up sequence.

Workaround is to configure 10GBASE-T port with 1G/10G speed before SSU (866682)

- QSFP-100G-CWDM4 transceivers with serials beginning with XYL installed in ports 41-48 of a 7800R3A-36DM2-LC may experience uncorrectable FEC errors and link flaps. The config-if mode command 'transceiver electrical lane 1-4 tx-input-equalization 4' on the /1 of ports containing the transceivers is an effective workaround. (873386)

SKUs Affected: 100GBASE-CWDM4, 100GBASE-XCWDM4 and 7800R3A-36DM2-LC

- 10G SFP optical transceivers which require 1.5W or more may fail to initialize. (873619)

SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DF-48Y, CCS-720DP-24S, CCS-720DP-48S and CCS-720DT-24S

- Inserting a third-party 400GBASE-ZR or 400GBASE-ZRP transceiver that does not implement CDB commands correctly may cause XcvrAgent to restart repeatedly.

Removing the transceiver or issuing interface config CLI command 'transceiver diag simulate removed' will allow the XcvrAgent to start correctly. (921923)

SKUs Affected: OSFP-400G-ZR, OSFP-400G-ZRP, QDD-400G-ZR and QDD-400G-ZRP

7300X3, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX and DCS-7050X3 Series

- Tagged packets with any VLAN priority received through MLAG peer links are always mapped to the same egress queue, and not complying with "show qos map". (407056)
- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (416694)
- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (417223)
- Some NAT rules may not work correctly if the StrataL2 forwarding agent happens to restart. Packets may not get translated as expected
The impacted rules need to be un-configured and then re-configured. (483821)
Series Affected: DCS-7300X Series
- Having any ACLs with UDF qualifiers as well as any of the following LAG/ECMP hashing features present in the configuration: RRoce hashing or TTL/HopLimit hashing, will result in incorrect operation of the LAG/ECMP hashing where the switch incorrectly hashes on IP payload. To avoid this problem the conflicting features should not be configured at the same time. (497016)
Series Affected: DCS-7050CX3, DCS-7050SX3 and DCS-7300X Series
- If NAT configuration is modified or if a NAT-enabled interface is shutdown and enabled while traffic is flowing, then some of the NAT connections might not be setup properly leading to packet loss or untranslated traffic.

Workaround is to be quiesce traffic before modifying the NAT configuration and then resume traffic. (497045)

- Disabling static NAT access-list sharing with "no hardware nat static access-list resource sharing" may in some cases leak hardware VFP resources.

Workaround is to reload the switch. (539951)

- Some multicast NAT translations may stop working after the restart of Strata forwarding plane agent.

Workaround is to remove and re-add the affected rules. (544827)

- With GRE tunnels or IPinIP tunnels or Nexthop-group tunnels configuration along with subinterface in the tunnel core interface, tunnel encapsulated packet will not be tagged with dot1q tag as configured in the subinterfaces, instead packet will go out tagged with internal vlanId assigned for the

subinterface. There is no workaround for this issue other than removing subinterface configuration in the core interface. (548327)

- When VXLAN encapsulated packet is received with an inner payload tag, the switch may decapsulate and send the packet with inner tag stripped when the egress chip is not same as the ingress chip. There is no workaround to this issue. (548492)

SKUs Affected: 7320X-32C-LC, CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F

- Some NAT translations may stop working after rebooting a switch.

Workaround is to restart the Strata Slice Agent. If the problem persists, perform a cold reboot of the switch. (563526)

- When VXLAN IPv6 underlay is configured and when core interface is a port channel on a SVI, continuous port channel interface flaps may result in persistent loss of VXLAN traffic to remote VTEPs that are reachable through that port channel. (567455)
- VXLAN encapsulated packets received with inner payload having more than one dot1q tag will result in outer dot1q tag being stripped when egressing out on a port. (567587)

Series Affected: 7300X3 and CCS-750 Series

SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F

- Replacing new scaled configuration with old scaled config may cause StrataL3 agent to restart. Feature will function normally after agent restart. (591050)
- When a linecard with LAG members configured statically is removed, subsequently moving or adding the LAG to a different VLAN may cause BUM traffic on that VLAN to be blackholed.

To work around this issue, remove the removed LAG members from the LAG. (603388)

SKUs Affected: 7300X3-32C-LC and 7300X3-48YC4-LC

- When StrataLag and StrataL3 forwarding agents restart, Egress VLAN translation (configured using "switchport vlan mapping") configured on port channel may not work.

The workaround is to shut and un-shut the impacted Port Channel. (603906)

Series Affected: CCS-710P, CCS-720DP, CCS-720DT, CCS-720XP, CCS-750 and DCS-7010TX Series

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- On switch reboot/reload with NAT configured, the ports come up and then NAT configuration is applied. If traffic is flowing before the NAT configuration is

fully programmed, then some packets might go out untranslated or dropped. Workaround is to stop traffic till NAT configuration is fully applied. (606202)

- If dynamic NAT configuration is modified or if the admin shutdown/no shutdown performed on an interface configured with dynamic NAT while NAT traffic is flowing, NAT connections may be left untranslated.

Workaround is to manually flush all the entries using the command "conntrack -F" in the corresponding network namespace bash mode. If configured on a non-default VRF (vrfName, for example), run the command "ip netns exec ns-vrfName conntrack -F". (615735)

- A single event upset (SEU) in the IFP_LOGICAL_TABLE_SELECT TCAM or IFP_TCAM_WIDE may not be detected. This may cause traffic to be misclassified.

A forwarding agent restart is required to forward traffic normally. (626452)

- Configuring 'switchport source-interface tx [multicast]' while the Strata-FixedSystem/Strata-Linecard agents are shutdown will result in an unexpected StrataL3 agent restart.

Workaround: Ensure the Strata-FixedSystem/Strata-Linecard agents are running before configuring 'switchport source-interface tx [multicast]' (635275)

- When IGMP Snooping is enabled, unknown multicast traffic may have a brief traffic outage on all interfaces when a speed change or transceiver gets inserted or removed on a QSFP port. (638649)
- A single-event upset (SEU) in the IFP_TCAM can lead to forwarding agent restart, after which the switch forwards traffic normally. (641105)

- With EVPN multicast solution configured, first few underlay multicast <S,G> entry may not get programmed; leading to any traffic arriving with the underlay <S,G> to be punted to CPU.

Workaround is to create an extra VLAN/VRF to underlay multicast group mapping before configuring the rest of the mappings. (643866)

- During config replace with NAT configuration, TCAM churn may result in forwarding agent restart. (676117)
- With IPV6 VXLAN Underlay configuration, when routing is toggled VLAN encapsulated packets sent on routed ports may be sent incorrectly with a VLAN tag resulting in packets getting dropped.

Workaround is to flap the core underlay interface. (679980)

- SSO switchover may result in remote ARP entries getting refreshed resulting in brief traffic loss for traffic routed to VXLAN remote hosts. (724315)

SKUs Affected: CCS-755-CH, CCS-758-CH, DCS-7304 and DCS-7308

- Under heavy lookup scenario, if the host CPU performs a write to one of hardware table responsible for MPLS/VXLAN tunneling feature, while a tunneled packet is undergoing a second lookup in the same table, it can result in packet drops or mis-forwarding for the flow.

Workaround is to restart the slice agent after stopping the traffic. (753304)

Series Affected: 7300X3, DCS-7050CX3, DCS-7050SX3 and DCS-7050TX3 Series

- With a large route scale, executing the CLI command "show platform trident l3 software alpm-table routes" may result in the restart of multiple forwarding agents. There is no workaround for this issue. (762247)
- After linecard hotswap or slice agent restart, static or peer-learned MAC addresses may go missing in hardware, which can result in unexpected flooding. Further, such missing macs may incorrectly be displayed in the output for "show mac-address-table".
A workaround for this issue is to remove and readd the MAC address. (762363)
- On multichip systems, reload may cause VLAN traffic-policy rules to not be reprogrammed across all chips. Workaround would be to reconfigure the VLAN traffic-policy rule. (824297)
SKUs Affected: 7300X3-32C-LC, 7300X3-48TC4-LC, 7300X3-48YC4-LC, DCS-7304 and DCS-7308
- On devices running EVPN VXLAN multicast OISM, Overlay multicast traffic with TTL1 arriving on VXLAN enabled VLAN in a VRF with `evpn multicast` enabled is not VXLAN encapsulated when there is a receiver in the same VLAN on remote VTEP.

Use TTL>1 or enable `redistribute igmp` under router BGP for the VLAN on the receiver VTEP. (833210)

- Executing "platform trident * reset" continuously may result in unexpected restart of StrataL3 forwarding agent. (854688)
- If StrataNat agent is killed or shutdown while NAT traffic is flowing, it may cause the StrataNat agent to continuously restart. System needs to be rebooted to recover from it.

Workaround is to stop NAT traffic before shutting down the StrataNat agent. (862801)

Series Affected: CCS-720DP, CCS-720XP, DCS-7050CX3 and DCS-7050SX3 Series

- Traffic sent from the CPU to a port that has tx mirror-to-cpu enabled (without recirc-channel) may get looped around and punted back to the CPU excessively. The workaround is to use a recirculation channel for tx mirror-to-cpu. (865752)

- When NAT translation is applied on an ingress packet with TCP checksum 0xFFFF (instead of 0), the TCP checksum is not recalculated after modifying the IP/TCP headers. The translated packet is sent out with TCP checksum 0xFFFF. (869547)
- Sometimes the acls programmed for dynamic NAT rules (in the TCAM) may not be unprogrammed after the corresponding rules are un-configured.

Restarting the forwarding plane agent should force the original rules to be unprogrammed. (881869)

- Global routing disable with MSS-G enabled can result in switch local IPs being unreachable.

Workaround: Enable routing globally. (890058)

7358 and DCS-7050X4 Series

- Traffic loss may be seen during StrataL2 restart (745553)
- MAC Moves may happen for very small duration when forwarding agent restarts on any of the peers of a MLAG. In case of VXLAN BGP EVPN, this may result in the peer learned MACs being locally learned and being advertised back to remote VTEPs. (749358)
- When a scaled QoS policy-map config with policers is applied in ingress direction, forwarding agent(s) may unexpectedly restart. (751937)
- CPU bound traffic may get dropped in the presence of multiple broadcast / multicast / unknown unicast streams with high traffic rate. (759085)
- Ipv4 microBfd packet sent to non-routed port with interface destination mac will be treated like regular L2 traffic. (769736)
- When BFD is configured but the session is down, BUM traffic can cause loops in a LAG. (778369)
- Simultaneously restarting multiple forwarding-plane agents might result in a small memory leak in AsicResourceMgr agent. (784711)
- If QoS Policy-maps are using some ACLs with rules having many L4 port ranges, in some cases, it is possible that rules are not programmed correctly in hardware. There is no error reported while programming the rules and it will have unexpected behavior in terms of traffic flow. (786011)
 SKUs Affected: 7358X4-SC, DCS-7050DX4-32S, DCS-7050DX4-32S-F, DCS-7050PX4-32S and DCS-7050PX4-32S-F
- BGP 'update wait-install' command is not supported on this platform. If this command is present in the startup-config, BGP routes are not advertised to peers.

To workaround this issue, remove 'update wait-install' command from startup-config and replace the config using 'config replace flash:startup-config' command or reboot the switch. (796360)

- Simultaneously restarting multiple forwarding-plane agents might result in a small memory leak in AsicResourceMgr agent. (807391)
- When IGMP Snooping is enabled some of the reserved multicast packets may not get flooded to all members of the VLAN if there is a multicast MAC entry to which the reserved IP matches. (830872)
- Scaling down static MAC address configuration while traffic with the same source MAC address as one of the removed MAC addresses is running might result in the MAC address not being learnt.
A workaround is to execute the "clear mac-address-table dynamic" command to prompt relearning of such MAC addresses. (852039)
- Scaling down static MAC address configuration while traffic with the same source MAC address as one of the removed MAC addresses is running might result in the newly learnt dynamic MAC address still being considered a static MAC address.
A workaround is to execute the "clear mac-address-table dynamic" command to prompt relearning of such MAC addresses. (852203)
- Multicast packets with an incoming TTL of 1 may get routed and sent out with a TTL of 0 instead of getting dropped. (877936)
- 2 bits parity error on some tables might not get corrected, causing traffic to be dropped unexpectedly.

Workaround is to restart slice agent . (884919)

- TCAM entries for the IPV4/v6 counters may not be freed even after deconfiguring the feature (891312)
- On interface flaps or peer reload, partial OIF replication for VXLAN bridging and multicast traffic may occur. The workaround is to restart the StrataL3Unicast agent. (906156)

Limitations and Restrictions in 4.29.6.1M

Accelerated System Update

- During a hitless reload upgrading from a release before 4.21.3, more than 3 LACP packets could be transmitted within a one second interval. This should not cause problems. (338048)
SKUs Affected: DCS-7260CX3-64, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F and DCS-7260CX3-64E-R

- Aborting 'reload fast-boot' or 'reload fast-boot now' command with Ctrl-C is unsupported. It may lead to the incorrect reload cause in 'show reload cause' in the subsequent reload. (344574)
- SSU of VRRP Master and Backup routers at the same time is unsupported and may cause extended traffic loss. Successfully SSU one router then SSU the other router. (345657)
- On systems with 4GB of /mnt/flash space, there may be insufficient space to store both an old swi image and a new swi image to perform an SSU upgrade. The workaround is to delete the old swi image from /mnt/flash prior to performing the upgrade. (385364)
- Performing SSU to boot into an image with this limitation may require that the currently running image is deleted to ensure that there is sufficient space available to perform the upgrade. (525870)

BGP EVPN L2/L3 VXLAN/MPLS

- On a BGP router, EVPN creates dynamic VLANs for L3 VNIs in receiving EVPN RT2 or RT5 advertisements even when the router does not have any L3 VRF or import route-target configurations. (294328)
- On a BGP EVPN device, if a VLAN-aware bundle is configured to include VLANs from multiple IP VRFs and if there are any IP addresses reused across multiple VRFs, then only one of the corresponding MAC/VLAN ARP bindings is distributed over EVPN. As a workaround, configure VLAN-aware bundles such that all VLANs in the bundle are in the same IP VRF. (514980)
- VRRP+VXLAN is not supported in deployment with EVPN VXLAN. (549525)
- In an EVPN/VXLAN environment with VLAN-Based MAC-VRF, type 2 and type 3 routes are incorrectly imported when the route-targets match even if the VNIs are different.

The workaround is to use different route-targets for the VLAN-Based MAC-VRF's to prevent the routes from being incorrectly imported. (631589)

- For an EVPN gateway to advertise received IP-PREFIX routes with next-hop-self, the corresponding IP-VRF must be active, which requires an interface to be configured in the VRF or the configuration of Router ID. BGP IP-VRF status is shown via "show bgp instance vrf VRF". (703768)
- IP ARP proxy is not compatible with EVPN VXLAN enabled VLANs.

Disable IP ARP proxy on VLANs with EVPN VXLAN enabled. (710344)

- In a EVPN DCI deployment, Symmetric Integrated Routing and Bridging is not supported in the default VRF. (733143)

- In networks using EVPN Active/Active multihoming with the modulus designated forwarder election algorithm, EOS versions prior to 4.29.0 consider all of the VLANs configured in a vlan-aware-bundle MACVRF when calculating the DF. EOS versions starting with 4.29.0 only consider VLANs configured in both the MACVRF and on the device itself. This can lead to disagreements on the elected DF if the range of VLANs configured in a bundle MACVRF is wider than the range of VLANs actually configured on the device, and the devices in the ethernet segment are running disparate EOS releases on either side of version 4.29.0. If different devices in the same ES elect different DFs, traffic loss can ensue.

As a workaround, either make sure all devices participating in an EVPN active/active multihoming ethernet segment are running the same version of EOS, or alternatively, before upgrading any device, ensure that the lowest VLAN configured in a vlan-aware-bundle is also configured on the device itself. This will ensure that the VLAN considered as the lowest VLAN in the MACVRF is the same pre and post-upgrade. (868750)

Containerized EOS

- ISIS/OSPF/OSPF3 may not function properly in cEOS if the kernel interfaces are prefixed with "eth" (397410)

vEOS Router / CloudEOS

- The /mnt/flash/cloud_ha_config.json based configuration for the CloudHa agent is not supported from this release. If you are upgrading from an older image (< 4.20.5) please reconfigure HA using EOS CLI. Please see information on converting json file based config to CLI commands in the vEOS Router Configuration Guide. (264660)
- While vEOS instances in an AWS cloud can be created using either a Bring Your Own License (BYOL) or Pay As You Go (PAYG) format, there currently is no show command in vEOS to easily help the user or support determine which mode this instance uses. (272649)
- Tx/Rx ring parameters of an "Elastic Network Adapter (ENA)" interface cannot be set through config.json in vEOS. (276382)
- vEOS supports up a maximum of 8 physical cores. With hyperthreading enabled it is 16 vCPUs (278286)
- In dpdk based vEOS (MODE=sfe in /mnt/flash/veos-config), even a packet matches a data plane ACL rule with log action, the log message might not show up if there is a similar rule with more exact ttl/dscp value match before that rule. (293641)

- In dpdk based vEOS(MODE=sfe in /mnt/flash/veos-config), SR-IOV mode doesn't work with Intel x520/82599 on ESXi. As a workaround use mixed mode instead of the SRIOV only mode in ESXi (296718)
- Hot detaching network interface instances is unsupported. Reboot the instance after detaching a network interface. (306132)
- CloudEOS does not support more than 10k /30 routes. The SFE agent will restart if more than 10k /30 routes are added to the system. The workaround is to reduce the number of /30 prefixes. (540204)
- NIC i40en driver version 2.2.4.0 is required on ESXi 7.0 host for SRIOV virtual function trust mode on/off functionality. (680174)
- NAT interface configuration mode is not supported in SFE. This needs to be kept in mind when migrating to SFE mode from Sfa.

The NAT interface configuration will need to be changed to profile based configuration in SFE mode. (682248)

CVX

- MapReduce Tracer is only recommended for deployments with 128 or fewer TaskTrackers.

If a switch has more than 128 directly connected TaskTracker nodes with MapReduce Tracer deployed, then the MapReduceTracer Agent can increase CPU utilization significantly. It is recommended to keep the directly connected TaskTracker count to below 128 when MapReduce Tracer feature is deployed. (80403)

- When a new SDN controller with 'manager ip' command is configured in HSC, HSC retains the configurations (e.g., logical switches) received from the old controller.

The workaround is to disable and re-enable the HSC service. (132171)

- Switches and CVX instances participating in a CVX setup must either all run 4.15.4F or later, or all run images from before 4.15.4F. (146664)
- If a CVX service agent connects to a client switch running a newer software version, the service agent may fail to connect to the client.

Upgrade the CVX software to a version greater or equal to versions running on all client switches. (219403)

- Intercept hosts on trunk ports with native vlan are only supported on 7050X2 platforms (257580)

- The Macro-Segmentation Service (MSS), requires MLAG fast redirection be enabled on the service MLAG TOR pair, when configured to work with a L2 transparent FW. This configuration minimizes traffic loss when any individual interface in the MLAG port channel goes down. (268291)
- The MssClient agent, can restart unexpectedly, when the MacroSegmentation Service, deployed with a L2-transparent firewall and is configured to intercept more than 10,000 end-points on a single switch (a number that's far greater than the total number of intercepts possible on the switch). (283221)
- In a Macro-Segmentation Service (MSS-FW) deployment with Palo Alto Networks Panorama, security zones configured on the Panorama are not reported in the API MSS uses.

As a workaround, configure the zones directly on the firewalls. (363371)

- The MacroSegmentation Service (MSS) allows users to configure multiple tags on CVX that can be used in firewall policy definition, to identify policies MSS should work on. It is however not possible to delete one of the multiple tags configured.

A workaround would be to remove all tags and reconfigure the subset required. (372320)

- In Macro-Segmentation Service (MSS) for L3 Firewalls , Virtual SVI IP addresses cannot be treated as intercepted hosts and thereby cannot be added to any redirect policy. However, these IPs can still be part of offload policies, even though implicit redirect will not work for them. (372504)
- If the Macro-Segmentation Service (MSS), working with a L3 firewall, is configured to intercept traffic from a host generating multicast traffic, the traffic can be black-holed.

To work around this issue, configure the following DirectFlow rule on all TOR switches MSS is programming rules on and assign it the highest priority:

```
flow bypassMulticast
  priority 65535
  table trident ifp
  match ethertype ip
  match destination ip 224.0.0.0 240.0.0.0
  ! (375156)
```

- MLAG devices in an environment with the Macro-Segmentation service (MSS) configured (with L3 firewalls), should have a high-priority DirectFlow rule configured to ignore traffic over the MLAG peer link. This can be configured using the following flow definition on each switch (assuming Po1 is the MLAG peer link):
flow bypassPeerLink

```

priority 65535
table trident ifp
match input interface Pol
match ethertype ip (375185)

```

- The Macro-Segmentation service (MSS), running with L3 firewalls, requires the following flows to be configured on the service VTEP devices connected to the firewall manually in order to prevent return traffic from the firewall from being subject to MSS rules again. The flows required are:

```

flow bypassFwIntf3
  priority 65535
  table trident ifp
  match input interface Pol103    >>>>>>> lag interface connected to FW
  match ethertype ip
!
flow bypassFwIntf4
  priority 65535
  table trident ifp
  match input interface Et31      >>>>>>> phy interface connected to FW
  match ethertype ip
! (375189)

```

- The MssPolicyMonitor agent might unexpectedly start running and eventually exit when the CVX functionality is disabled on the master CVX instance. This does not have any side-effect on the operation of the Macro-Segmentation Service (MSS). (378077)
- When the Macro-Segmentation Service (MSS) is enabled and works with a L3 firewall, there can be a momentary traffic outage on reloading a MLAG peer if the CVX to VTEP (MLAG switch) connection is not re-established before the MLAG reload delay expires at the reloaded peer. (418560)
- If MSS policy enforcement rules configuration is changed from verbatim to group verbatim, the SandAcl agent can restart unexpectedly (433592)
- When Macro-Segmentation Service (MSS) is enabled on the CVX in an environment where no switch is configured as CVX client, multiple "excessive warmup delay" syslog messages for MssPolicyMonitor agent might be noticed. (497132)
- In a Macro-Segmentation Service (MSS) deployment with a PaloAlto firewall, enforcement policy configured with application using dynamic port assignment is not supported. (515749)
- Macro-Segmentation Service (MSS) running in consistent enforcement mode may emit "CVX-3-MSS_HOST_UNPROGRAM_ERROR" syslog messages for the hosts attached to only one MLAG peer when that peer is reloaded.

These messages can be ignored. "CVX-3-MSS_HOST_RECOVER" messages will be

observed as soon as these hosts are learned again at the reloaded MLAG peer afterwards. (524705)

- For Macro-Segmentation Service (MSS) deployment running in consistent enforcement mode, "CVX-3-MSS_HOST_UNPROGRAM_ERROR" or "CVX-3-MSS_RULE_UNPROGRAM_ERROR" syslog messages may be seen when a switch is reloaded (planned reload).

Use "clear cvx connections client <switch-id>" command to clear the switch from CVX registration. (524717)

- In a Macro-Segmentation Service (MSS) deployment in a multi-VRF environment with EVPN as control plane, when an MLAG peer is reloaded, CVX-3-MSS_HOST_PROGRAM_ERROR message is logged for the intercept hosts followed by CVX-3-MSS_HOST_RECOVER message. (525851)
- In a Macro-Segmentation Service (MSS) deployment with a Palo-Alto Firewall, when using an aggregator interface in a virtual system, all its member interface must be in the same virtual system. (528511)
- In a Macro-Segmentation Service (MSS) deployment with a Palo-Alto firewall, a policy with multiple source or destination zones will be enforced over only one source zone and/or one destination zone.

A workaround is to split the policy into multiple policies with single source and/or destination zone. An alternative is to add the source and destination subnets to the policy configuration. (567376)

- In a Macro-Segmentation Service deployment with Layer 3 firewall (MSS-FW), when a policy matching a subnet with a mask less than /16 is tagged as a group-based redirect or a regular offload policy, the service may become non-operational. However, any subnet mask length is supported for all variants of verbatim policies. (716283)

General

- CLI does not allow access to files whose names contain spaces. (539)
- 'speed forced 1000full' is not supported on interfaces with the "Management" prefix such as ma1. Use speed autonegotiation instead. (1506)
- Ports configured as mirroring destinations will be shown as active even when they are shutdown. (65221)
- When VmTracer is used in a network that includes Cisco devices, ESX hosts connected to the Arista switch will see CDP frames from the Arista switch as well as from the other Cisco devices. VmTracer will display the VM info when the ESX host reports Arista's CDP info, but then will delete it when the ESX host reports other CDP info.

The workaround is to drop all inbound CDP on all ports using MAC ACLs. (101756)

- EOS SDK-based applications must be run via the 'daemon' command configuration. They cannot be run directly from bash. (158431)
- EOS extension scripts and agents which do not use EOS-SDK will need to be modified to work in 4.16.6M and beyond. (158467)
- EOS swix extensions containing RPMs which were based out of or procured from Fedora distributions prior to Fedora 18 may fail to install due to dependency issues. Suggested workaround is to recreate the swix files with the corresponding RPMs from Fedora 18 distribution. (162325)
- Packets with size greater than MTU of egress interface will not honor directflow flows and may not be forwarded (180927)
- On the DCS-7150 series, IEEE 1588 PTP transparent clock is not supported on 40G interfaces (182897)
- SSO is hitful with subinterfaces. Protocols running on subinterfaces may experience session flaps during SSO. (188070)
- ACL drop counters cannot be cleared. (202905)
- when ACLs are configured with stats over port-channels, these stats are kept at the physical port level not at the port-channel level. (203131)
- With international and HW-REV-01 images and the introduction of "ip access-group <access-list name>" other configuration commands beginning with 'ip' that are configurable in the global config mode are not accessible in "router bgp" mode.

To access those commands exit the "router bgp" config mode to switch to the global config mode. (219390)

- Changes made to running-config after a config session is created might be reverted when the config session is committed which contains changes in similar or related areas. Use 'show session-config diff' inside the config session before committing to verify what changes will be made. (226850)
- The uptime in the output of "show module" may not match the uptime as shown by "show version" or "show uptime". The uptime in "show module" represents the time the module has been up since the last change in the "Status" column in the output of "show module", which differs from the meaning of uptime in "show version" and "show uptime". (248230)
- To avoid interoperability problems when using strong SNMPv3 encryption algorithms, follow the following minimums:

- When using AES192 for privacy, use a minimum of SHA224 for authentication.
- When using AES256 for privacy, use a minimum of SHA256 for authentication. (268290)
- Mirroring to an MLAG port channel is not supported. (276973)
- Copy commands with sftp: or scp: URL do not work via eAPI by passing password through the "input" parameter. SSH keys have to be setup to run the commands without a password. (283716)
- EOS support alphanumeric VLAN names, however ODL gives an error when reading the leaf /interfaces/interface/routed-vlan/config/vlan if the name contains alphabetic characters.

To workaround use numeric VLAN names only in EOS. (361244)

- The Macro-Segmentation Service (MSS) fails to skip processing a tagged firewall policy if one of the interfaces in the zone are not configured with a valid IP for that zone.

To workaround this problem, configure a valid IP address on the interface. (367835)

- Enabling next-hop sharing (using "ip hardware fib next-hop sharing" command) could result in slower convergence in the case of link failures. (369305)
- A reload with "The system rebooted due to a watchdog on the lower card" as the reload cause may occur unexpectedly on hardware built prior to 2020. (372163)

SKUs Affected: DCS-7260CX-64-F, DCS-7260CX-64-R, DCS-7260CX-64-SSD-F, DCS-7260CX-64-SSD-R, DCS-7260QX-64-F, DCS-7260QX-64-R, DCS-7260QX-64-SSD-F, DCS-7260QX-64-SSD-R, DCS-7280CR-48-F, DCS-7280CR2-60-F, DCS-7280CR2A-60-F, DCS-7280CR2K-60-F, DCS-7280QR-C72-F, DCS-7280QR-C72-M-F, DCS-7280QR-C72-M-R and DCS-7280QR-C72-R

- For static tunnel interfaces, duplicate or conflicting GRE tunnels are tunnels that have the same source, destination, tunnel mode (GRE) and key configurations as an existing tunnel. A duplicate tunnel is always in the down state. When there are duplicate static tunnel interfaces and a config-replace is executed, it is non-deterministic which one of the conflicting tunnels will be in the up state. (373135)
- When an interface profile is used to configure an interface's access VLAN, the VLAN is not created automatically. The VLAN must be manually created with the vlan command. (376621)
- When using interface profiles, configuring interfaces may take a few minutes, depending on the number of interfaces being configured. It may be faster to

apply an empty interface profile to the interfaces, and then update the profile. (382120)

- The CLI commands "locator-led [multifan | powersupply | chassis]" will not cause multifan, powersupply and chassis LEDs on the front panel of the system to flash, as they did on previous generations of fixed systems. The beacon LED on the back panel of the chassis is purple by default. The CLI commands "locator-led chassis" causes it glow red instead of expected blue. (395892)

SKUs Affected: DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3K-32P4-F, DCS-7280PR3-24-F and DCS-7280PR3-24-M-F

- The configuration lock feature does not prevent changes to the configuration via SNMP. (402556)
- If an EOS CLI command is used to create a reference to a VRF that does not exist, a subsequent OpenConfig YANG operation may cause the VRF to be created in EOS. (411908)
- Large amounts of monitoring traffic processed at the control plane via aggressive sflow sampling or mirroring to CPU could affect other control plane functions due to CPU bandwidth contention. (415992)
- IP Locking may send out additional queries when STP changes state (even possible in portfast config which is common for the IPLocking use-case). (417806)
- When hardware or software flow tracking is enabled and tracking flow is a VXLAN flow, the exported IPFIX packet may contain incorrect inner VLAN ID. (421437)

Series Affected: CCS-720XP Series

- When "logging format hostname ipv4" is configured, log messages may not display the correct IP address if the IP address is updated. (422666)
- Egress MAC ACL logging is not supported (433169)
- In the 'show interfaces interactions' CLI, the list of speeds an interface is limited to does not account for the limitations that result from speed-group configuration. (434363)
- SSH may fail when sFlow and Management SSH have the same QoS DSCP value configured, and sample rate for sFlow is set to dangerous. (434763)
- Per-port denial of service protection (PDP) support has been removed. (440074)

- Mirror-On-Drop does not monitor IP packets with an invalid IP version. (443839)
- When we first configure a port-channel lag_type, we must set values for both aggregation/config/lag-type and ethernet/config/aggregate-id (i.e. member configuration) in the same transaction. The lag-type is populated in EOS once the port-channel has one member port and remains set. (445421)
- When eAPI is configured with invalid SSL profile, nginx will stop running and all http/s requests will be refused (474422)
- Currently EOS does not support parsing and separating multiple ACL rules in the same NAS-Filter-Rule AVP using zero byte delimiter. (475210)
- EOS-2GB.swi is no longer available. (475680)
- Installed licenses without corresponding licensing features configured might not be visible from the CLI (504793)
- CVE-2019-16905: The version of openssh shipped in EOS does not enable support for XMSS, so EOS is not affected by CVE-2019-16905. (511376)
- CVE-2020-1968: EOS is susceptible to CVE-2020-1968. The workaround is to only use ephemeral (DHE) ciphersuites. (519024)
- If Octa, OpenConfig or TerminAttr are started under low memory conditions, they can cause a spike in CPU usage before unexpectedly restarting. (534841)
- Configuring NTP to serve clients in more than 500 VRFs fails with a syslog message "ntpd: Too many sockets in use". (543116)
- When the "tunnel destination" address is configured as a broadcast address for static tunnel interfaces then there could be packets drops during decapsulation at the tunnel endpoint. (544032)
- For a given interface, multiple simultaneous MAC address flush requests on different VLANs may result in clearing all MAC addresses on that interface. This is an optimization to support a higher scale of VLANs and interfaces. To disable this optimization, use the "mac address-table flushing interface vlan aggregation disabled" CLI configuration. (544712)
- The "dot1x re-authenticate" command is not supported for failed supplicants. Workaround is to flap the interface. (567028)
- Removing "vlan" key-field from custom TCAM profile for ingress security ACL feature causes ACL application on SVIs to fail.

To workaround this issue, add "vlan" key-field to custom TCAM profile for ingress security ACL features. (570680)

SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F

- In SSO redundancy mode, the switch may take longer (more than 20 minutes) to be declared switchover ready. This is to accommodate for all agents to become ready for switchover. (572615)
- When unsupported transceiver support is unlocked using a key, the ports that are in disabled state before unlocking will continue to be disabled.

The workaround is to remove such transceivers and re-insert them after unlocking the unsupported transceiver support. (574146)

- The 'logging level all' CLI expands to individual commands in the configuration, preventing dynamic updates of the logging list as part of EOS upgrades. (579047)
- TACACS+ authentication and authorization through SSH does not work with Pulse Secure server because EOS sends different port names in authentication ("ssh") and authorization (actual vty) requests which are rejected by the server. (583678)
- Sampled flow tracking exports IPFIX flow records with egress interface as 'discard' for flows on pseudowire patched interfaces. (603916)
- A CPU internal error (BIT30 SyncFlood BIT11) may cause an unexpected reboot (623392)

SKUs Affected: DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7050SX3-48C8-F, DCS-7050SX3-48C8-R, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050SX3-96YC8-F, DCS-7050SX3-96YC8-R, DCS-7050TX3-48C8-F, DCS-7050TX3-48C8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-36S-F, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-36S-F, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F, DCS-7280PR3K-24-F, DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R, DCS-7280SR3K-48YC8-F and DCS-7280SR3K-48YC8-R
- IP Locking does not support supervisor stateful switchover (SSO). (631563)
- Dot1x agent incorrectly accepts larger than advertised values for idle timeout and session timeout RADIUS response attributes (639283)
- /mnt/flash cannot be used as the log archiver destination on systems formatted as VFAT. To check if a product is formatted as /mnt/flash, from bash, run "mount | grep flash", and look for "type vfat" or "type ext4". If you see "type vfat", then that product's /mnt/flash storage is formatted as VFAT. (642415)

- locator-led fantray & locator-led powersupply CLI commands do not operate successfully on Modular Systems and may report "This LED cannot be used as locator-led target". (652512)
Series Affected: 7368, 7388, CCS-750, DCS-7300X, DCS-7300X3, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 Series
- "mac security profile" is not supported on L2 subinterfaces. (680967)
SKUs Affected: 7800R3-48CQM2-LC, 7800R3K-36DM-LC, DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- When Path MTU discovery is enabled under static tunnel interfaces using "tunnel path-mtu-discovery", the discovered MTU is only used for software forwarding in the slow path. When MTU is configured under static tunnel interfaces using "mtu" config, the configured MTU value is used for hardware forwarding. (681101)
- OpenConfig does not support BFD configuration for VLAN interfaces.
BFD configuration under VLANs must be removed for any gNMI set operation to succeed. (682920)
- If RFC2544 test is started on a sub-interface and the parent interface has "traffic-loopback source system device mac" configured, the test packets will get into infinite loop.

Unconfigure traffic-loopback on the parent interface using "no traffic-loopback source system device mac" to stop the traffic loop. (683129)

- If the AVA switch sensor ("monitor security awake") is enabled along with the TerminAttr (CloudVision), the TerminAttr agent will also process IPFIX packets sent to the CPU, and the "Traffic flows" feature on the CloudVision portal is automatically enabled and can cause additional CPU usage.
If "Traffic flows" feature is not required on CloudVision, disable TerminAttr IPFIX processing by adding "-ipfix=false" argument to "exec /usr/bin/TerminAttr" under "daemon TerminAttr" configuration. (688791)
Series Affected: CCS-720XP Series
- EOS images are optimized during installation on selected platforms. The optimized EOS image saved on non-volatile storage may be incompatible with other platforms. An attempt to install an incompatible SWI will result in the following warning: "Next image does not support required optimization". Only install full EOS images or compatible optimized EOS images. (691810)
- From EOS release 4.29.2F, the 64-bit image of EOS supports "monitor security awake" or AwakeSensor.swix/NDRSensor.swix on platforms with at

least 8GB memory. EOS releases prior to 4.29.2F, the 64-bit image of EOS does not support AwakeSensor.swix on any of the platforms. Use the 32-bit EOS version instead. (719835)

- System will not boot if a cable connected to an ethernet switch is connected to the console interface of the system. All system LEDs may remain red and the console cable will subsequently be unresponsive, until the cable connected to an ethernet switch is removed from the console interface of the system. (724181)

SKUs Affected: DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7050SX3-48C8-F, DCS-7050SX3-48C8-R, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050SX3-96YC8-F, DCS-7050SX3-96YC8-R, DCS-7050TX3-48C8-F, DCS-7050TX3-48C8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-36S-F, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-36S-F, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F, DCS-7280PR3K-24-F, DCS-7280SR3-40YC6-F, DCS-7280SR3-40YC6-R, DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R, DCS-7280SR3E-40YC6-F, DCS-7280SR3E-40YC6-R, DCS-7280SR3K-48YC8-F, DCS-7280SR3K-48YC8-R, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R

- Inbound packets destined for local agents on the switch (e.g., LACP, STP) will not be decoded properly when the "tcpdump" command listens on the "any" interface.

The workaround is to direct the "tcpdump" command to listen on the interface on which the packets are being received. (730292)

- RAIL is not supported on Trident3X1 and Trident3X2 chipset based products (736685)
- PTP traffic cannot be captured using a monitor session on PTP enabled ports for 750X devices configured as a PTP Boundary Clock (738554)
- Smbus card agent may contain assertion errors when restarting due to card removal. There is no effect on service. (740093)

Series Affected: 7368 Series

- Removing one or more rule(s) with sequence numbers which causes a hole in ACL sequence number, i.e., removing rules from middle of an ACL and, re-adding ACL rules (with auto-generated sequence numbers) can cause sequence numbers to be re-used.

This series of ACL operations will cause an "Duplicate sequence number" error to be emitted in CLI.

To workaround this issue, if a rule in middle of ACL needs to be removed then,

1. Copy the rules which appear after the rule of interest,
2. Remove all rules from the rule of interest till the end of ACL,
3. Re-add rules copied in step 1. (743473)

- `* network-instances/network-instance/policy-forwarding/policies/
policy[policy-id=POLICY-ID]/rules/rule[sequence-id=SEQUENCE-ID]/ipv4/config/
source-address`
`* network-instances/network-instance/policy-forwarding/policies/
policy[policy-id=POLICY-ID]/rules/rule[sequence-id=SEQUENCE-ID]/ipv4/state/
source-address`
configures/displays at most 1 source-address prefix per rule, despite EOS CLI supporting >1 source-address prefix per rule.

As a workaround, either use OpenConfig exclusively for configuration/query/telemetry on IPv4 source-address prefixes, OR do not use OpenConfig for configuration/query/telemetry on IPv4 source-address prefixes if >1 IPv4 source address prefixes per rule have already been configured via EOS CLI. (747080)

- `* system/logging/remote-servers/remote-server/selectors/selector/facility`
`* system/logging/remote-servers/remote-server/selectors/selector/config/
facility`
do not filter syslogs sent to the remote-server based on RFC5424 facility. Instead, these two paths configure the RFC5424 facility for all syslogs, regardless of original facility, sent to the remote-server (filtered by having severity >= configured severity). (747096)
- JSON does not specify the precision of a Number in the specification. eAPI can generate numbers up to 6-bit, which can be truncated by some JSON clients with lower limit (e.g., some implementations have a max value of $2^{53}-1$). The only workaround is to use a different JSON deserializer (769194)
- TACACS+ request has a limit of 8150 bytes in body size. If a CLI command is too long, authorization might get rejected. (770296)
- Workaround for the restrictions on the list
`* /system/logging/remote-servers/remote-server[host]/selectors/
selector[facility severity]`
is to update/replace/merge all `/system/logging/remote-servers/remote-server[host]/selectors/selector[facility severity]` list items at once, such that the end configuration is that all `/system/logging/remote-servers/remote-server[host]/selectors/selector[facility severity]` list have (a) exactly 1 list item with identical facility and severity key values, or (b) no list items. (774251)
- When a fifth non-unique TTL value is configured for Static Interface Tunnels on Hardware Forwarded Platforms using the "tunnel ttl" config command, the

Static Interface Tunnel pipeline programming may not be completed due to hardware limitations and the traffic may not flow through the Tunnel. Only 4 unique TTL value are supported across all the Static Interface Tunnels on Hardware Forwarded Platforms. If TTL value is not configured for Static Interface Tunnels using the "tunnel ttl" config command, then it implies that the TTL value is zero. This zero TTL value is also a unique TTL value. For Static Interface Tunnels to be operational, configure 4 unique tunnel TTL values across all the Static Interface Tunnels on a Hardware Forwarded Platform (782720)

- EOS with version 4.29.0 or later generate a different user id from the same username. As a result, files on flash may show random user id instead of the expected user name. For example, a file created by the "admin" user under an older version may show that it is now owned by user "106904808" after upgrading to a newer version. The workaround is to run the following command from bash:

```
find /mnt/flash -uid 106904808 -exec sudo chown admin {} \; (798231)
```

- Across all flow tracking features, when the number of configured exporters for a given tracker exceeds the maximum supported for the feature on the platform, the active exporters may differ after a reload or shutdown/no shutdown of the flow tracking feature.
Reduce the number of configured exporters to maximum supported for the flow tracking feature on the platform. (806831)
- A switch configured with both MACsec L2 protocol forwarding and configurable EAPOL DMAC will send the EAPOL packets with the configured DMAC to CPU instead of transparently forwarding it in the dataplane. (824887)
- When PTP boundary clock is configured, incoming PTP packets on a PTP-enabled access switchport that contain a VLAN other than the one configured on the access port will be sent to CPU. (837200)
Series Affected: CCS-750 Series
- In a MLAG VXLAN deployment with anycast gateways, ARP synchronization between the MLAG peers may result in packet drops in the CoPP "copp-system-selfip" queue and adversely affect any application traffic over that queue. (881798)
- Arista Networks Product Bulletin for EOS Release 4.29.1F incorrectly included DCS-7280DR3A-36-F, DCS-7280DR3AM-36-F, DCS-7280DR3AK-36-F in the New Platform Capabilities section. DCS-7280DR3A-36-F, DCS-7280DR3AM-36-F, DCS-7280DR3AK-36-F were introduced in 4.30.0F. The Arista Networks Product Bulletin for EOS Release 4.29.1F was revised on January 15, 2024 to correct the documentation. (902473)
Series Affected: DCS-7280R3 Series

- If a VRF selection policy and a match rule is created using CLI first, and the rule is updated via gNMI, the resulting VRF selection policy configuration may have unexpected additional rule.
Workaround is to use only either CLI or gNMI instead of both. (910953)

IPSEC

- We do not support NAT on Isec autoVPN in this release. (564431)
- When `ethernet untagged allowed` is configured under `ip security` config, TCP MSS clamping may not apply to packets encapsulated over a VXLANsec tunnel. (685484)

Series Affected: DCS-7280CR3 Series

Layer 1

- The MACSec outPktsEncrypted and outOctetsEncrypted counters are not supported with the MACSEC Proxy feature (281715)
- On Vxlan MACsec proxy, disable learning is not supported on proxy patch vlan. (284348)
SKUs Affected: 7500R2M-36CQ-LC, DCS-7280CR2M-30-F, DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R
- MACsec proxy for front panel ports requires replay protection to be disabled on the MACsec profile, otherwise some packets will be dropped due to sequence error. (339533)
SKUs Affected: 7500R2M-36CQ-LC, DCS-7280CR2M-30-F, DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R
- Clause 37 (BASE-X) auto-negotiation is not supported on SFP25 ports. (492271)
SKUs Affected: DCS-7010TX-48-F and DCS-7010TX-48-R
- Ports configured at PAM4 speeds (50G-1, 100G-2, 200G-4, 400G-8) do not support link-debounce intervals shorter than 2 seconds. (500674)
Series Affected: 7368 and DCS-7060X4 Series
SKUs Affected: DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32P4-F and DCS-7280CR3MK-32P4-R
- InPktsNoSCI counter increments for a MACsec frame with AN that is not in use by the receiver. InPktsSAnotInUse should have incremented instead. (532869)

SKUs Affected: 7500R2AM-36CQ-LC, 7500R2M-36CQ-LC, 7800R3-48CQM-LC, DCS-7280CR2M-30, DCS-7280CR2M-30-F, DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R

- Interfaces with BCM82391 PHY do not support auto-negotiation. (537530)

SKUs Affected: 7500R-8CFPX-LC, 7500RM-36CQ-LC, DCS-7050CX3M-32S, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R

- QSFP to SFP adapters are not supported on front panel interfaces ethernet25/1-36/1. To work around this, use interfaces ethernet1/1-24/1 instead. (625046)

SKUs Affected: DCS-7280CR3-36S-F and DCS-7280CR3K-36S-F

- MACsec frames are dropped when the sci setting in the "mac security profile" does not match the peer's sci setting. (664596)

SKUs Affected: 7388-8D, 7800R3-48CQM2-LC, 7800R3K-36DM-LC, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

- MACsec frames with SL less than 5-bytes are silently dropped on ingress (714497)

SKUs Affected: 7800R3-48CQM2-LC, 7800R3-48CQMS-LC, 7800R3K-36DM-LC, DCS-7170B-64C-F, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

- MACsec frames with SL less than 24-bytes are silently dropped on ingress (714597)

SKUs Affected: 7500R-8CFPX-LC, 7500RM-36CQ-LC, DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R

- Cable diagnostics test(TDR) on a BASE-T port may report inconclusive result if link partner is configured at forced 100M speed. (716208)

Series Affected: CCS-710P, CCS-720XP, CCS-750 and DCS-7050TX3 Series

- The MACsec inPktsNoTag counter includes the non-standard MACsec inPktsCtrl counts. inPktsCtrl (non-standard) counts packets that are parsed as control packets and hence forwarded to the uncontrolled port. For instance, EAPOL (MKPDU) and possibly LLDP or Pause/PFC if bypassed. (735450)

SKUs Affected: 7388-8D, 7800R3-48CQM2-LC, 7800R3-48CQMS-LC, 7800R3K-36DM-LC, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R,

DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R,
DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

- The MACsec inPktsNoTag counter includes the non-standard MACsec inPktsCtrl counts. inPktsCtrl (non-standard) counts packets that are parsed as control packets and hence forwarded to the uncontrolled port. For instance, EAPOL (MKPDU) and possibly LLDP or Pause/PFC if bypassed. (735455)
SKUs Affected: 7500RM-36CQ-LC, DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- 50G-2 autonegotiation is not supported on interfaces with the CRT50216 PHY. (755659)
SKUs Affected: 7358-16C, 7368-16S, 7368-4D, 7368-4P, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3K-48CQ-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4A, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4A and DCS-7280CR3K-96
- 100G-4 autonegotiation with FEC disabled is not supported on interfaces with the CRT50216 PHY. (796531)
SKUs Affected: 7368-16C, 7368-16S, 7368-4D, 7368-4P, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3K-48CQ-LC, DCS-7060DX5-64S-F, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R and DCS-7280CR3K-96-F

Layer 2

- LACP PDU fast rate ("lacp rate fast" on an Arista switch) should not be configured on any port in an MLAG pair, or any port connected to an MLAG pair. (13950)
- Static ARP inspection may not work on modular platforms after SSO switchover. The workaround is to disable and re-enable static ARP inspection after switchover. (244477)
- The command "lacp rate fast" is not supported on modular systems with stateful supervisor switchover (SSO) enabled. Neighboring devices should not have "lacp rate fast" configured on ports connected to the said system. (248121)
- In a MACsec profile, when a primary key's CKN is changed, then operating primary continues to be principal actor even when a fallback key is configured till the time the new primary establishes a successful MKA session. There is zero traffic disruption. (275678)

- Port Security Protect does not support trunk ports. If a port is secured to source MAC <A> in VLAN <X>, traffic from source MAC <A> in other VLANs of the trunk port will also be allowed. (344773)
- While performing SSO, a device's peers might not receive LLDP packets within default TTL (120 seconds) and those peer devices may not maintain LLDP status information for the device undergoing SSO. After SSO is complete, the peers will again have LLDP state available.

To avoid losing LLDP state during SSO, consider increasing LLDP hold time at the remote. (347459)

- An MLAG switch cannot be used as an EVPN multi-homing provider edge. When a switch has MLAG enabled, MAC addresses from ports with an ethernet segment configured are advertised through BGP as if they were single-homing, and local ethernet segments are not advertised through BGP. (397867)
- Even with fallback to unprotected traffic configured, traffic loss in the order of milliseconds may be seen during initialization (491232)
- Upon programming a large number (~250) of VTEPs in a flood set at once, the L2Rib agent may restart. Workaround is to program them in batches, waiting a few seconds between each batch. (544051)
- 802.1X is supported on LAG member ports using EAPOL authentication. (544556)
- The SyncE wait-to-restore (WTR) timer is armed when administratively shutting / unshutting an Ethernet interface. This is inconsistent with disabling / enabling SyncE on the interface in which case the WTR timer is not armed.

To workaround the WTR timer, issue 'no sync-e' followed by 'sync-e' for the given interface(s). (648153)

- When an EEC has redundant links traceable to the same reference (e.g. ITU G. 781 section 5.13.2 "'bundles' of ports"), and the quality level of the reference transitions from a higher quality to a lower quality (e.g. QL-PRC to QL-SSU-A), the EEC may briefly enter a holdover condition for a few seconds. (687326)
- Synchronous Ethernet is not supported for 10G BASE-T transceivers. (689493)
- Dot1x session move is not supported if the supplicant changes authentication methods during the move. (705466)
- There is no Accelerated Software Upgrade(ASU) support for 802.1X dynamic ACL assignment feature. (725100)

- Loop protection is not intended to be used along with Spanning Tree Protocol (STP) (745522)
- A clock source with a lower quality level than the internal oscillator is incorrectly shown as standby in "show sync-e selection" when it is unusable. (761975)
- LAG interface counters will have constant counts discrepancy compared to the aggregate interface counts of its member interfaces when there is a link flap event on LAG interface. (785073)
- If a RADIUS server's hostname is long and crosses a certain character limit, the output of 'show dot1x radius' will be misaligned. (820586)

Layer 3

- ECMP is not supported for routes learned via RIP. (34773)
- DHCP relay is not supported when running virtual machines on EOS (101754)
- CLI does not detect TCAM exhaustion and automatically revert the change when a PBR policy is applied to an L3 interface in the "shutdown" state, since the actual hardware programming happens when the interface comes up. (122651)
- OSPFv3 support for multiple address families (RFC 5838) in EOS introduces support for IPv4 routing with OSPFv3. OSPFv3 for IPv4 AF in EOS requires configuring neighboring OSPFv3 IPv4 routers on the same link with primary IPv4 addresses in the same subnet. Adjacency is established in OSPFv3 IPv4 AF even if the neighbor's primary IPv4 address is in a different subnet but routes through the neighbor on a different subnet will however not be installed. This limitation may be removed in a future release. (140234)
- When BFD is unconfigured with Fhrp as the only client and/or BFD is administratively shut, both VRRP routers will become Master. This transition is temporary and Master-Backup relationship will be restored. (159647)
- BFD sessions established between Arista switches and Cisco peers over IPv6 link-local addresses may flap when BFD echo mode is enabled. (159803)
- vEOS does not support BFD echo sessions. (160770)
- If "bfd per-link rfc-7130" is configured on port-channel(s) with primary and secondary IPv4 addresses configured in the same subnet, and the "bfd neighbor" is configured with peer switch port-channel's secondary IP address and there is a BFD session using primary IP address, then the port-channel

might flap continuously.

Configure secondary IP address on the affected port-channel(s) to different subnet from their primary IP address on both switches. (182622)

- BFD echo functionality is not available on L2 Port-Channels with BFD RFC7130 mode enabled. (190418)
- If a port channel is configured with "bfd echo" and "bfd per-link rfc-7130" and is used for OSPFv3 or PIM, BFD sessions will not come up with echo functionality. BFD with SSO for OSPFv3 or PIM is not supported. (190997)
- An LACP port channel configured with BFD per-link and echo may experience BFD session flap on existing member links if new member link(s) are added to the port channel while the BFD peer device undergoes Stateful Switchover (SSO). (255042)
- Eos does not support programming IPv6 ECMP paths to kernel (258387)
- The DSCP value of BGP control plane packets set by the command 'bgp transport qos dscp <value>' will be overwritten when the platform 'qos rewrite dscp' is enabled and the outgoing interface is a SVI based on traffic-class to dscp map present on the platform. (269764)
- In the output of the "show ipv6 dhcp relay counters" or "show ip dhcp relay counters" command, interface-specific DHCP relay counters might not sum up to "All Req" and "All Resp" counters. (278786)
- Tunnel interfaces flap when configuring ttl, tos and path mtu discovery on them. (281464)
- IPv6 addresses are not supported as tunnel source and destination. (283912)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289217)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289224)
- When BFD is configured with per-link mode rfc-7130 for an IPv6 link-local address on a Port-Channel and the Port-Channel transitions to the Down state, the Port-channel may not transition back to the UP state.

The workaround for this situation is to disable and then re-enable BFD per-link mode rfc-7130 on the Port-Channel. (343814)

- On Macro-Segmentation Service (MSS) deployment with L3 firewalls, a flow entry that fail to get programmed due to lack of TCAM space, is not programmed when TCAM space becomes available at a later time. The workaround is to "shut" and "no shut" on the direct flow config. (372232)
Series Affected: DCS-7050TX, DCS-7050X3, DCS-7060X and DCS-7060X2 Series
- The update wait-install feature may not work as expected if there are routes which fail programming due to hardware resource exhaustion. The routes which failed programming may still get advertised out to peers. (372765)
- DHCP packets will not be processed by EOS DHCP Server when DHCP Relay is relaying packet from one SVI where DHCP Relay is configured to another SVI where DHCP Server is configured. (389017)
- When a sampled packet egresses a trunk port and the path is ECMP, sFlow will arbitrarily pick one of the ECMP next hops for the next hop IP of the router extension. (415747)
- DHCP clients will fail to obtain IPv6 leases when "ipv6 dhcp relay destination" is set to IPv6 multicast address. (426796)
- Destination (TEP) or include hop entry for a RSVP tunnel can be TE router ID of a router or IP address on any interface that has an IGP neighbor. It cannot be any other interface address, like a loopback interface address which is not TE router ID or a stub interface address. (453019)
- Nexthop group tunnel counters does not work, when tunnel counters is enabled using 'hardware counter feature mpls tunnel' command. (474078)
- VXLAN features are not compatible the 'ip hardware fib next-hop sharing' configuration. Configuring VXLAN routes may result in identical FECs not being shared and potential misforwarding after changes to FECs are made. (477471)
- The config 'ip hardware fib next-hop multi-path maximum' does not restrict the size of installed next-hop group ECMP FECs (498484)
- CLI configuration of a static route with nexthop pointing to self ip address is not rejected. (514894)
- NAT dynamic profiles are not supported on multi-chip and modular system. (525213)
Series Affected: 7300X3, CCS-720XP and DCS-7300X Series
- BFD over Vxlan does not work for IPv6 underlay. (526897)

- Rsvp autobandwidth feature does not work for single hop tunnels. The workaround is to configure Rsvp explicit-null on the egress router for single hop tunnels. (571177)
- VARP VTEP IP configuration is not supported on AP aggregation VTEPs

It is recommended that either EVPN VXLAN symmetric or asymmetric IRB be used for distributed VXLAN routing in campus environments. (571517)

- IS-IS Extended administrative group (EAG) sub-TLV is not supported. If a system in the network advertises EAG, Arista router will ignore the sub-TLV. (628807)
- The `show ip hardware fib ecmp resilience` CLI may sometimes incorrectly report an unordered resilient ECMP as ordered. (685318)
- If a preferred IPv6 source address for an IPv6 kernel route is obtained from a local Port-Channel interface that is used in a "software forwarding hardware offload route local interface INTERFACE" command, and the interface has a IPv6 BFD neighbor configured, and the Port-Channel interface is subsequently shut down, the preferred IPv6 source address of the kernel route is not updated as expected.

The workaround is to remove the "bfd ipv6 neighbor" command from the Port-Channel interface and re-apply the command. (698716)

- If a "software forwarding hardware offload route rcf FUNCTION" command is configured under "router kernel/address-family" sub-mode and FUNCTION returns True for a connected route, then kernel traffic to that prefix is forwarded in hardware as expected. But if the connected route later goes away for some reason (e.g., the interface is reconfigured) and is replaced in the FIB by a route for the same prefix but learned from a routing protocol peer, the kernel route is programmed for software forwarding instead of hardware forwarding.

Restarting the KernelFib agent should cause the affected kernel route to be reprogrammed such that forwarding of kernel traffic to the prefix is offloaded to the hardware. (701746)

- Resilient ECMP may log erroneous "RECMF_CAPACITY_EXCEEDED" syslog during interface flap. (722339)
- CLI command show ip route may incorrectly display an unprogrammed route as programmed. (731198)
- BFD running over L3 interfaces can reduce the time it takes to shrink an ECMP when one of the ECMP member links fail. If all the ECMP paths use the

same L3 interface (such as when using SVI interfaces) then the time to shrink the ECMP may take longer. (746361)

- The `show ip hardware fib ecmp resilience` command may display `VRF <name> not found` when given a non-default VRF. (764757)
- IPFIX and sFlow BGP nexthop value cannot be accurately determined for ECMP flows going over the same egress interface that are encapsulated with implicit-null labels. There is no workaround for this issue (816528)
- When "ip virtual-router mac-address mlag-peer" is enabled in the MLAG environment, BFD is not supported on the MLAG peer interface. (818124)
- gNMI replace on /routing-policy/defined-sets/bgp-defined-sets/community-sets/community-set/config/community-member from one representation of standard community number to another might result in redundant CLIs even though the two representations are effectively the same (826724)
- "MIXED" prefix-sets are not supported in OpenConfig in policy constructs. The workaround is to use separate prefix-sets for IPv4 and IPv6 prefixes. (858818)
- ECMP path shrink or fallback to the backup path may be slow when BFD session state goes down if there are multiple next hops with different address families within the ECMP group. (867829)

Multi-agent

- If a BGP peer is configured for BFD fallover, the BFD session is not cleaned up even after the BGP peer is unconfigured. (217199)
- Traffic destined to an L3EVPN route whose underlay has ECMP BGP-LU routes, may experience temporary traffic loss when the underlay goes from n way to n-1 way or lower. The duration of loss can vary depending on the scale of BGP-LU routes or L3EVPN routes. (251692)
- When the multi-agent routing protocol model is configured, the maximum number of next hops for an OSPF ECMP route is 128. (274888)
- With very a very low multi-hop BFD timer configured, a BFD session may flap when a path (belonging to an ECMP set), over which the BFD session is established, goes down. (287571)
- Route-map specified as match-map for dynamic prefix-list is evaluated only against BGP routes. (345444)
- The IpRib agent runs out of addressable memory when a large number of routes are leaked to a large number of VRFs. (346575)

- Routes which are imported into a target VRF using a "leak routes source-vrf" policy cannot be matched on the basis of the source-protocol from the source-VRF as part of the "rib fib policy" (403739)
- In multi-agent mode, extended Sflow BGP data is not supported for BGP LU routes. (411762)
- When using 64-bit versions of EOS, Bgp, BgpCliHelper and Bmp agents may report a virtual memory size up to 50 GB larger than actual memory used. The actual memory usage is better represented by the sum of "RES" and "SHR" from "show processes top". (427364)
- When BGP router has advertised a large scale of routes (e.g. several million) and it reloads, there may be traffic drops while it is coming back up. This happens when the reloading router comes back up before the directly connected peering routers (which are receiving routes via an intermediate Route Reflector) get a chance to process millions of withdrawals and delete all the routes from FIB. (458127)
- An AS_PATH prepend operation in an RCF function applied towards an IBGP peer is ignored. (474006)
- IP ACLs are not supported for control plane route filtering for VRF leaking (536329)
- Match dynamic prefix-list clauses in route maps are only supported for BGP inbound and outbound route maps. (564552)
- CVP will not monitor multiple BGP peers that are configured using the same peering IP address. (588954)
- BMP, SNMP and EosSdk will not report multiple BGP peers configured with the same BGP peering address. (602657)
- When default-originate is configured, default route will not be advertised to a neighbor if a prefix-list configured directly for this neighbor denies the default route. (648587)
- When multiple single-hop EBGp sessions are configured using point-to-point physical interface IP addresses between the same pair of switches, some of the BGP neighborship may get established using the other peer's IP address. When any BGP neighborship comes up using such wrong IP address from cross-link, the BGP neighbor configured on the other link will have connection establishment failure with Notification error: 'connection collision resolution'.

To workaround this issue, either use 'update-source' option for the neighbor or explicitly enable 'no neighbor <addr> route-to-peer' for all the single-hop EBGp peers using point-to-point physical interface IP address. (686390)

- BGP session tracker may not detect quick BGP peer flaps and may not error disable interfaces (687998)
- For the RT membership AFI/SAFI, RT membership routes advertised by a given peer are also reflected back to the same peer. i.e. there is no echo suppression of the RT membership NLRI. (717330)
- BGP TCP-AO is not supported on cEOS. (762218)
- With "next-hop resolution v4-mapped-v6 translation", IPv4-mapped IPv6 next hop address of IPv6 Unicast NLRIs is translated to IPv4 address for next-hop resolution. However, such paths are not advertised to outbound peers if the next hop is unchanged.
A workaround is to explicitly configure outbound route map or RCF for the peer that sets the next hop to be path's next hop. (790764)
- BGP-LU routes directly redistributed into IS-IS using command 'tunnel source-protocol bgp ipv4 labeled-unicast' under 'router isis' mode, uses MED as metric instead of AIGP.

To workaround this issue, IS-IS metric value can be derived from AIGP metric using RCF policy. Note that the 'aigp.metric' attribute value, used by RCF at redistribution point of application, is same as AIGP metric path attribute value received from the peer and doesn't account for the next-hop's IGP metric cost. (791558)

Rib

- PBR recursive resolution for nexthops is not supported for packets routed in software (like packets with IP options) (82077)
- No support for IPv6 labeled Unicast capability when BGP neighbor is established using IPv4 transport. (208402)
- Match dynamic prefix-list clauses in route maps are only supported for BGP inbound and outbound route maps. (261865)
- Dynamic prefix-list match clauses are not supported in BGP neighbor inbound route maps. (269663)
- The BGP configuration commands, "bgp next-hop-unchanged" and "neighbor next-hop-unchanged", are not supported in ribd mode (single-agent mode).

In ribd mode, this functionality can only be achieved via route map configuration. (275007)

- Ribd agent may restart while running with Internet scale routes and with a large number of (50+) Rib-Out groups. Such a high scale is not supported with single-agent mode.

To workaround this issue, either reduce the number of Rib-Out groups by checking the output of 'show bgp update-groups' or upgrade to multi-agent mode which can support very high scale. (558044)

- Static routes with administrative distance of 255 are installed in the RIB and eligible for redistribution into other protocols with "redistribution config". The routes are also installed in the FIB for programming in hardware.

To deny FIB programming of routes, the RIB route control configuration policy in the form of "rib ipv4|ipv6 fib policy" can be used.

To prevent RIB redistribution of routes, the routes can be filtered out as part of the "redistribution" policy into the target protocol. (663662)

- There may be a behavior difference in OSPF between single agent and multi agent mode when both OSPF and BGP are configured in a non-default VRF. In single agent mode, a route that would come from an external type 5 LSA might be dropped where in multi agent mode, the route would not be dropped.

The workaround to prevent loops is to follow RFC4577 (use of the DN-bit) (774087)

MLAG

- Spanning-tree mode backup does not work in an MLAG configuration. (15151)
- The port-channel configured on each MLAG switch with the same MLAG ID must also have the same channel-group ID. (22890)
- Configuring a device connected to an MLAG with a round-robin LAG distribution algorithm is not supported if the device is going to participate in IGMP. (28370)
- On a scaled MLAG setup, Lag+LacpAgent agent may restart unexpectedly if MLAG is reinitialized due to configuration changes. (116125)
- On an MLAG system configured with dual primary detection and with default control plane ACL enforced, MLAG will not form if the management interfaces of the MLAG peers are connected through extra hops (e.g., on different subnets), due to the default control plane ACL dropping MLAG control packets with TTL < 255.

The work around is to apply a control plane ACL which permits MLAG control packets with the correct TTL. (271308)

- TCAM profile switch in a MLAG switch may result in some packet duplication for broadcast or multicast packets (341353)

- On an MLAG device VRRP sessions configured on VLAN interfaces may briefly transition to the master state after device reboot even though the VLAN is not associated with any active L2 ports. The VRRP session will become Stopped once the MLAG reload delay timer has expired. Beyond unexpected log messages there is no impact of device operation. (799723)

MPLS

- The MPLS agent may continuously restart on a router with high BGP LU Tunnel scale. (377939)
- BGP LU Tx MPLS routes corresponding to BGP LU advertisements with multiple labels are not supported. (395029)
- When RSVP is configured with split-tunnels, the `show traffic-engineering rsvp tunnel` command may display a last sampled bandwidth value for the tunnel that is not exactly equal to the sum of the last sampled bandwidth value of all sub-tunnels. This can occur if the command is used while the values are actively changing as it is possible that some values in the output have not yet been updated while others have. (625866)

Multicast

- If a single IGMP snooping port has multiple multicast hosts attached and performs proxy reporting or report suppression from these multiple hosts downstream, immediate-leave must be disabled for the VLAN.

Use "no ip igmp snooping VLAN <vlanId> immediate-leave". Otherwise, some desired multicast traffic might be lost. (21367)

- After ASU reload, if the first PIM hello packet received from the neighbor gets lost and If the neighbor has a high hello query-interval, it will result in traffic loss. Workaround is to use default PIM hello query-interval and increase the PIM hello query-count. (341447)
- With scale greater than 5000 multicast IPV6 routes, show platform sfe mroute ipv6 <vrf> commands might not display the right number of routes when performing multiple addition/deletion operations across VRFs. This is applicable only if user level multicast forwarding agent is configured to be used instead of kernel. (378841)
- Bessd might unexpectedly restart on config replace with scaled setup over 30 VRFs and large number of PIM enabled interfaces. Work around is to reduce the scale. (398185)
- Pim "non-dr install-oifs" feature becomes disabled if the non-dr needs to route multicast traffic for any reason to the interface on which the feature was originally configured. This is how the feature is expected to work. The workaround is to setup the network in a way that the DR always becomes the

assert-winner. However, if the non-DR becomes assert winner for some routes on an interface and feature becomes disabled, flap the pim sparse mode configuration on non-DR on that interface to attempt to get out assert winner state. (401672)

- Multicast routes are not created when a source is transmitting only fragmented multicast packets.
Switch to using SFE for multicast software forwarding. (605326)
- A device with 8GB RAM might be able to support a maximum of 20 VRFs or lesser depending on the memory requirements of other features configured on the device. A workaround is to reduce the VRF scale or use a switch with higher RAM for such a scale requirement. (707709)
- Multicast scale is limited on devices running with 4GB memory. Running `software-forwarding sfe` on these platforms can result in BessMgr agent crashing. (786687)

NAT

- SNAT will not be performed if the original source IP is not fully resolved, and DNAT will not be performed if the translated destination IP address is not fully resolved. For routed ports, a static or dynamic ARP entry must be present for the relevant IP address. For SVIs, both an ARP entry and a MAC table entry for the MAC address associated with the IP address must be present. (756211)
Series Affected: DCS-7130 Series
- Changing the forwarding profile between nat-vxlan and nat may cause the Nat agent to unexpectedly restart. (792534)
Series Affected: DCS-7170 Series
- When an interface configured with static NAT translations is taken out of the shutdown state, untranslated traffic will flow briefly. (805141)
Series Affected: DCS-7130 Series
- If static source NAT is configured on the OUTSIDE interface of an INSIDE/OUTSIDE pair of interfaces, configuring static destination NAT on the INSIDE interface will cause the static source NAT translation to fail. (821248)
SKUs Affected: DCS-7130-16G3-F (HwRev:C1), DCS-7130-16G3-R (HwRev:C1), DCS-7130-48G3-F (HwRev:B1), DCS-7130-48G3-F (HwRev:B2), DCS-7130-48G3-R (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48L-F (HwRev:A2), DCS-7130-48L-F (HwRev:A3), DCS-7130-48L-R (HwRev:A2), DCS-7130-48L-R (HwRev:A3), DCS-7130-48LB-F (HwRev:A3), DCS-7130-48LB-R (HwRev:A2) and DCS-7130-48LB-R (HwRev:A3)

SwitchApp

- IGMP snooping filters are not supported when SwitchApp over forwards IGMP reports. Reports are over forwarded if the "igmp flooding" CLI is enabled through the FPGA instance CLI, or if the FPGA profile does not support correct forwarding of IGMP reports. (601917)
- Multicast boundaries do not filter unicast PIM encapsulated packets sent to the RP. This may result in unicast PIM traffic continuously reaching CPU in certain network topologies. (718224)
- When running the SwitchApp application as a PTP grandmaster, it will distribute its internal free running counter instead of system time, which corresponds to the time the SwitchApp has been started since 1 January 1970. (736982)
- When using SwitchApp with routing and MLAG, if link failure causes more than 64 ARP entries to be needed on the peer link then traffic may be dropped.
Design the network with care to ensure that this limitation does not occur for the relevant failure scenarios. (765647)
- MakoSharedResources agent may restart hitfully if MakoL3Unicast agent performs a hitful restart due to crashing during a hitless restart. The system will recover to normal operation without intervention. (800466)
- SwitchApp platforms are limited to 64 nexthops per L2 port. This includes the inter MLAG peer link. As a result MLAG systems where more than 64 hosts are reachable via MLAG ports may hit this limitation during failure scenarios where traffic for more than 64 hosts need to cross the peer link resulting in traffic loss. It is therefore recommended to design configurations to ensure that this limitation is not exceeded in the event of failures. (811102)
- When the hardware resources become full, new NAT rules will not be programmed in hardware, and the system will log:

```
`%NAT-3-HW_RESOURCE_FULL: Hardware resources are insufficient to program all NAT rules`
```

Those rules will remain pending, even if hardware resources subsequently become available.

The workaround is to remove any pending rules (`show ip nat translation pending`) and re-add them to use the available hardware resources. (848105)

DCS-7170 Series

- If an ingress packet contains more than one 802.1Q tag and the egress port is a trunk port, the egressing packet may get corrupted. (180436)
Series Affected: DCS-7170 Series
- For frames that are routed and VXLAN encapsulated, underlay MTU is not used for MTU enforcement. Therefore the packets will be routed, encapsulated and transmitted without fragmentation. (254767)
- Changing port speed from 10G/25G/50G to 40G/100G will cause the forwarding agent (BfnSliceAgent) to restart. This can result in a brief traffic outage in the order of 100ms on all ports. The link state of other ports will not change. (256519)
- 10/25/50 Gbps interfaces only support MTU up to 7Kbytes. (258823)
- After applying or removing shaping configuration on an L2 sub-interface, the L2 sub-interface must be flapped for the configuration to take effect. (281312)
Series Affected: DCS-7170 Series
- After setting CoPP rate = 0, a few packets (<10) might be let into the CPU, but everything will be dropped afterwards. (350676)
- The number of L2 subinterfaces which can be supported on a parent interface is 63.
On a given parent interface, there cannot be more than one L2 subinterface with the same forwarding vlan. (357946)
- Full mroute scale may not be achievable under certain scenarios due to hash collision for the multicast route table. (380545)
- When egress NAT and MLAG are both used, 'show ip mroute' may show NAT'ed addresses in addition to the receiving mroute addresses. (394571)
- firewall profile on DCS-7170 platform does not support IP multicast (401134)
- Accelerated Software Upgrade (ASU) is not supported in combination with the packet-filter profile. The hitless reload will time out and fall back to a normal reload. (473985)
Series Affected: DCS-7170 Series
- On the DCS-7170 series, ASU is not supported on the firewall profile. (475112)
Series Affected: DCS-7170 Series
- When sFlow and monitor session are both enabled in the system, the mirrored traffic may be incorrectly throttled to the sFlow copp policy rate. (604224)
- The 32-bit image of EOS does not support the 7170B series switches. Use the 64-bit version instead (646589)

- IP Multicast packets that are software forwarded are not subjected to egress static NAT. (771707)

DCS-7170 Series

- On the 7170 series of switches, when IPv6 is used in VXLAN underlay, the UDP checksum will be 0 for all the packets going out with an IPv6 VXLAN encapsulation. (297660)

CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM and CCS-750 Series

- Configuring Authentication Failure VLAN to same value as 'Internal VLAN' or 'AAA-server assigned successful VLANs', may lead to un-desirable behaviour (376571)
- When hardware flow tracking is enabled, tcpControlBits information element in IPFIX flow records for VXLAN encapsulated TCP flows is incorrect. (416079)
- On receiving COA-REQ for changing VLAN for EAPOL supplicant, the NAS deliberately fails authorization for the first time in the new vlan. (449890)
- The per-interface configuration to ignore reauthorization timeouts from the AAA server ("dot1x timeout reauth-timeout-ignore") is not designed to work if manual reauthorization is requested from the command line. The reason for this is that manual reauthentication is designed to force the supplicant to contact the AAA server to retrieve the latest credentials. (453044)
- CoA request for VLAN change for EAPOL supplicant returns CoA-NAK. The VLAN does get changed when device is reauthenticated and the AAA server sends the new VLAN in an Access-Accept response. (459726)
- Hitless FPGA upgrades are not supported due to hardware limitations. Therefore FPGA images will not be upgraded during SSU. (474978)
 SKUs Affected: CCS-720XP-24Y6-F, CCS-720XP-24Y6-R, CCS-720XP-24ZY4-F, CCS-720XP-48Y6-F, CCS-720XP-48Y6-R and CCS-720XP-48ZC2-F
- Negation operator (!) and 'assigned' keyword are not supported in Radius NAS-Filter-rule(attribute-id 92) Attribute. (475207)
- Ip options qualifiers are not supported in the NAS-FILTER-RULE attribute if present in the Access-Accept or COA packet from the AAA server. If Ip options are present in NAS-FILTER-RULE, it may cause supplicant to fail authorization. (480659)
- WOL feature is currently only available for phone trunk port, but not for regular trunk port. (481468)

- Fallback to MBA authentication feature may not work as expected with Caching Auth feature. (482333)
- Timeout value configured for MBA fallback feature can block MBA Auth for upto 60 seconds more than the configured value since first non-EAPOL packet is received. It is recommended to configure the hold period to be lower than the timeout value for MBA fallback. (483995)
- MBA for passive device can not be triggered by ARP/ICMP if IP locking is enabled (500439)
- If a switchport loses power, the credentials of the attached phone and PC will remain cached when authentication caching is enabled. However, when power is restored, the EAPOL enabled PC may not be reachable within the default reauthorization limit of 3 attempts and the supplicant will be removed as a result.
The solution in this situation is to increase the reauthorization limit using the per-interface dot1x command: dot1x reauthorization request limit 10 (500488)
- Limitation - 802.1X Web Authentication Implicit ACL feature doesn't work when IP Locking is configured on the interface. The workaround is to use 802.1X ACL based Web Authentication. (557986)
- PFC and pause generation are not supported when MACsec is enabled. (577310)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- An interface that acts as a Dot1x supplicant cannot be paired with a third party device that acts as a Dot1x authenticator. (643209)
Series Affected: CCS-750 Series
- Mirrored packets sent to a MACSec enabled destination interface are sent out unencrypted. (649877)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- Command "ptp free-running source clock system" is not supported. Although it may complete without errors, it will not take effect. (708316)
Series Affected: CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series
- PTP will not synchronize over interfaces configured with half-duplex link speeds. (740727)
- Phone sending LLDP packets before getting authenticated (by sending non-LLDP packets like DHCP) may not get classified as phones.

The workaround is to configure "dot1x protocol lldp bypass" (742933)

- "VTEP to VTEP bridging" is not supported. (791265)
SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F

CCS-750 and DCS-7060X4 Series

- When re-enabling a disabled power supply using `power supply enable module PowerSupply<n>`, the power supply may not re-enable.

The power supply can be recovered by removing and re-applying the input power to the power supply. (857791)

SKUs Affected: PWR-511-AC and PWR-511-DC

DCS-7170, DCS-7280R and DCS-7280R2 Series

- On devices with Algomatch enabled, removing an ACL on an interface might fail because the remaining ACLs on the system may not fit in the hardware tables. (192811)
- Ingress ACLs on IPv6 packets do not work when IPv6 uRPF is configured and AlgoMatch is enabled. (512394)

DCS-7130 Series

- This platform only supports 64-bit versions of EOS. (615075)
- Several link flaps may be seen as the links are coming up. Debounce time CLI do not apply to Ethernet Interfaces on dropbear products. Link monitoring, Debounce and EOAM rely on PCS link status and not enabled on these since these interfaces use network PMA to set link status. The failure this reported here is a test issue. (633947)
- In the event of power loss, the switch will not report the correct reload cause. If power is interrupted briefly, the switch will report the reload as being caused by a watchdog. If the power loss lasts more than a few seconds, the switch will report the reload as cause unknown. (700291)
Series Affected: DCS-7130 Series
- CL28 Autoneg is not supported. This means auto-negotiation on copper cables at speeds 10G or lower is not supported. (719438)
Series Affected: DCS-7130 Series
- When link loss is experienced on a MetaMux input interface, MetaMux may erroneously increment the error count for that MetaMux input interface. (723415)
Series Affected: DCS-7130 Series

- SwitchApp's published routing scale limits are approximations and simplifications so it is possible for a system to run out of internal resources for certain prefixes even though the overall scale is within the published limits. (764479)
- Use of unidirectional L1 configuration for a front panel Ethernet interface populated with a BASE-T type transceiver is not supported. Setting unidirectional configuration to such an interface may produce unexpected results, including traffic loss. Workaround is to use unidirectional configuration with non BASE-T type transceivers. (860391)
Series Affected: DCS-7130 Series

7358, 7368, 7388, CCS-750, DCS-7300X, DCS-7500R, DCS-7800 and DCS-7800R3 Series

- When a system shuts down due to a power overload, EOS does not log a 'power overload' reload cause (371160)
- Mixed supervisor types are not supported on dual-supervisor modular systems. (450053)
- With scaled VRF, SVI and Arp entries, typically beyond claimed support, the Arp agent may experience a SIGQUIT during SSO switchover, on the new active supervisor. There is no workaround for this issue and switchover is ultimately successful. (495528)
- IP reachability between the ManagementActive interface and physical management interfaces may be impacted if both sets of interfaces are configured in separate VRFs. (885866)
- If Management0 interface redundancy is enabled with ipv6 neighbor monitoring, the backup interface can become active on supervisor switchover instead of the primary interface (893164)

CCS-710P, CCS-720DP, CCS-720XP, CCS-722XPM and CCS-750 Series

- Polycom SoundStation IP 7000 phone might stop sending traffic after a switch power cycle if the phone stays on PoE power without an L1 link for some time.

Workaround is to toggle PoE power on the affected interfaces with "poe disabled" followed by "no poe disabled". (502688)

7500R3, DCS-7020, DCS-7130, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280R3A, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800, DCS-7800R3 and DCS-7800R3A Series

- The L3 MTU on interfaces is not enforced on SVIs. (11659)
Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- Disabling IPv6 routing on an L3 interface is not supported when IPv4 routing is also enabled on that L3 interface. (43093)
- On DCS-7500 series, stateful switchover is not hitless. (54305)
- show interfaces counters rates' for port channels may show a slightly inaccurate data rate (including above 100%). (67367)
- A VXLAN encapsulated packet after de-encapsulation is not forwarded back to the port on which the original packet was received. Similarly, a native ethernet packet after VXLAN encapsulation is not forwarded back to the original receive port towards the remote VTEP. (75075)
- Due to hardware limitation, if an egress IP ACL is applied on an interface which is part of a VLAN with an egress router ACL, the filtering behavior for egress traffic of that interface may be undefined. (82094)
- For packets that hit rules across multiple ACLs, the hit counts are accounted for only in a single ACL. PACL counters take precedence over RACLs and MAC ACLs, and RACL counters take precedence over MAC ACLs. (85440)
- When an egress IPv6 RACL is applied to a VLAN, IPv6 packets may not be subjected to MTU checks. (88778)
- An egress ACL on a destination port of a monitor session only takes effect for Rx mirrored packets which are IP forwarded. The egress ACL will not work for bridged packets or Tx mirrored routed packets. (89915)
- When a shaper is configured the output of the shaper has a tolerance of +30% with 100-byte frames. This reduces to +5% as the frame size is increased to 600 bytes. The shaper variance reduces when the port happens to be the endpoint of a tunnel. (93546)
- The egress per-queue byte counter is not exact for certain packet types. For routed packets the egress per-queue byte counter may be offset by +/-3 bytes. For any packets originating from the CPU, the per-queue byte counter may be offset by -12 bytes. (97660)
- Packets that attempt to do a GRE key lookup where the GRE key lookup misses will have egress IPv6 ACLs applied with the ingress VLAN (102455)

- GRE key forwarding packets that are recirculated will have no QoS policy applied to them. (102458)
- If a double tagged frame is received on an untrusted or DSCP-trusted interface, the CoS is retained if the outgoing frame is tagged, unless the frame is routed (131614)
- On the DCS-7280E and DCS-7500E series, configuring QinQ or SelectiveQinQ along with configurable TPID on the same interface is not supported (131757)
- IPv6 Egress ACL deny logging is not supported on subinterfaces. (146487)
- Disabling IPv4 routing on an L3 interface is not supported when IPv6 routing is enabled on that L3 interface. When "no ip routing ipv6 interfaces" is configured to enable the forwarding of IPv4 packets over IPv6 nexthops over interfaces that do not have IPv4 address, but only a IPv6 address, it will also allow routing IPv4 traffic over these interfaces. (151356)
- Egress IPv4/IPv6 RACL does not work for VXLAN encap-route case. (154551)
- Congestion on an egress interface used by a GRE tunnel that is a mirroring destination can cause drops on forwarding ASICs where the mirrored traffic ingresses the switch. (158001)
- On the DCS-7500E and DCS-7280 series, traffic outage is expected during transition from a single member LAG to two member LAG and vice versa. Use "platform sand lag hardware-only" to avoid this outage. (160439)
- When MLAG peer mac routing feature is enabled on both MLAG peers, the IPv6 neighbor may not be resolved when source ipv6 address in solicited neighbor is link local address. (165067)
- When the 'vxlan-routing' hardware TCAM profile is configured, VXLAN traffic flows in overlay may not work alongside IPv6 traffic flows in underlay. The workaround is to use the CLI command 'no platform sand ipv6 host-route exact-match'. (190017)
- L2-ECMP over VXLAN is not supported for multi-homed hosts under EVPN. (193475)
- In Tap Aggregation mode, Tap Aggregation features may not work for 802.1BR/VN tagged packets when 802.1BR/VN tag stripping is not configured. (198798)
- Links are not compatible with systems running EOS 4.18.1F

The workaround is to upgrade both sides of the link. (199152)

SKUs Affected: 7500R-8CFPX-LC

- Traceroute to a IP route where the IP route is getting forwarded through a MPLS nexthop-group may incorrectly display the first hop as being

unreachable. The actual first hop will appear in the second position.
(209273)

- When ingress packet truncation is enabled in Tap Aggregation mode, truncated packets are not counted as packets received on the ingress interface. (215346)
- Frames less than 64 bytes are dropped (as expected), but not counted.

Command "hardware phy cmx42550 <intf > diag read64 mac line 0x138" can be used to dump the counter for the purpose of debugging". (218429)

SKUs Affected: 7500R2M-36CQ-LC, 7800R3-48CQM-LC, DCS-7280CR2M-30-F, DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R

- Native VLAN setting on Tap interfaces does not take effect. (227770)
- While in Tap Aggregation mode, ingress VLAN membership filtering does not work when packets are destined to a tool interface configured with egress packet truncation. (227841)
- While in Tap Aggregation mode, interfaces configured for egress truncation will continue to consume hardware resources even when they are down. (228384)
- LLDP functionality is not supported on interfaces with ingress truncation enabled. (229983)
- Sflow agent CPU utilization is around 10% even when hardware acceleration is enabled and the agent is not processing any flow samples. (251674)
- With MTU increased beyond 9100B, there can be some multicast packets of size greater than 9100B that are dropped as a port gets close to linerate. (261446)
- Changing PMF profile on one linecard may cause a brief disruption of PMF-based features on other linecards which should not be affected by the same change. (270849)
- A QoS policy, with a class map which has either set-tc or set-dscp on the ingress port and which has DSCP rewrite map on the egress port, results in indeterministic behavior with respect to DSCP remarking of the outgoing packet. (273320)
- If a QOS and PDP Policy both having police action are applied on an interface and the traffic hits both the policies then QOS Policer Counter is messed up. The "show policy-map type qos pmap-name" will give the PDP Policer Counter, not the QOS Policer Counter. (275994)

- Rate values for subinterfaces shown by "show interface counter rates" may deviate up to 15% from corresponding rate on parent interface. (278700)
- With TCP MSS Clamping feature configured, TCP SYN packets might get dropped under high CPU usage conditions if the incoming TCP SYN packet rate is high. (279370)
- Packets larger than 10196 bytes will be discarded by the forwarding ASICs. (280459)
- PFC on DSCP trusted ports is not supported. (281214)
- Routing is not supported on the parent interface of a L2 sub-interface. (287016)
- Loop protect feature is not supported when AlgoMatch HW is enabled. (290706)
- When a port becomes a LAG member, there may be a very brief moment where a few packets may be duplicated on the LAG. (295260)
- Fragmented IP packets will not match on flow-spec rules that include a match on TCP/UDP port. (297309)
- When all the L2 subinterfaces belonging to a parent interface are shutdown, packets are bridged and mac addresses are learnt on parent interface (300202)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- VXLAN routing is not supported on a modular system on which one of the linecards is configured for TAP aggregation. (306180)
- Show interfaces counters ip only show IPv4/Ipv6/MPLS counters per physical interface and not per port-channel (306209)
- When a BGP peer session is cleared, traffic which hits existing flow-spec rules may increment the wrong counter for a brief period of time. (339523)
- If non shaped sub interfaces are configured under same interface having shaped sub interfaces, traffic will egress out of the non shaped sub interfaces in a round robin fashion. It is recommended to not have shaped and unshaped sub interfaces under the same parent. (341851)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Hardware accelerated sflow with profile sflowaccel-v2 doesn't work with 7500E linecards. (342130)
- Stale telemetry flow-spec statistics may be seen until the next "show flow-spec" command is issued. (345195)
- The sFlow sample output interface of IPv6 packet subject to an egress ACL, is set to unknown output interface. (354746)

- Hardware accelerated Sflow does not support collectors reachable over arbitrary tunnel types (354749)
- IPv4 packets with options in the IP header will not be subjected to Unicast Reverse Path Forwarding (uRPF) checks, and hence will not be dropped even if the source IP is unknown. (355542)
- unicast RPF is not enforced on IP packets when the destination IP needs ARP resolution (356347)
- Software forwarding is not supported for GRE or IP-in-IP tunnel terminated packets using decap-groups. (358369)
- The egress DSCP value in the Outer IP header for packets that undergo headend replication after routing is derived from the Traffic Class (and not copied from the ingress DSCP value). (366189)
- While in Tap Aggregation mode, if the identity tag configuration of a Tap port is changed, traffic received on this port may be sent to an invalid multicast destination momentarily. This will cause the DchUnreachables counter to increment. (372757)
- TX mirroring shows VXLAN encapsulated packets that are actually dropped by split horizon. (375795)
- If more than maximum supported Ipsec tunnels are configured, it is random as to which ones are assigned flow entries within on-chip encrypt/decrypt engine. The set of Ipsec tunnel interfaces assigned flow entries may not be the same as the set of tunnel interfaces which have IKE connections established. As a result, some tunnel interfaces with IKE connection established may not get link up. (376431)
- Configuring more than 250 BFD sessions on a physical interface may result in many of those sessions flapping because of hardware drops on the interface. (389387)
- Drop Precedence based tail drop thresholds can be ignored when multiple queues are congested and buffering resources are low. (395693)
- Queue build-up doesn't show up for 64 byte packets in case of congestion via following commands:-
 - 1) platform fap diag diag cosq non_empty_queues
 - 2) show interfaces queue length

Following command should be used instead:-
show queue-monitor length (398937)
- PIM sparse mode feature which allows installation of outgoing interfaces on the non-DR for a faster failover does not work on L3 sub-interfaces. This

feature relies on an egress ACL to stop multicast traffic from going out of the interface. Multicast egress ACLs do not work on L3 sub-interfaces due to a platform limitation. If the feature is configured, there will be duplication of traffic. The workaround is not use the feature on such interfaces. (405696)

- When a policy-map with match on inner VLAN is programmed on an interface, if single tagged packet traffic is sent to that interface, the behaviour of that packet is indeterministic as inner VLAN tag is calculated using an offset and vlan-tag-format qualifier doesn't differentiate between single and double tagged packets. (415587)
- When MPLS label is pushed on the traffic ingressing CoS trust port then EXP bits are not derived from TC based on packet's COS. (417108)
- Byte counts are not supported for IPSec tunnels. (419031)
- SflowAccel supports up to 200 VRFs. (423892)
- PAUSE frames with correct DMAC=01:80:C2:00:00:01, Ethertype=0x8808, opcode=0x0001 are never forwarded and are always consumed by the hardware. (426047)
- Header truncation for mirrored packets can not be used when the destination of a monitor session is a GRE tunnel. (428547)
- A port transmitting packets smaller than 100B that were tunnel terminated on the switch might be limited to ~90% linerate (429797)
- When the VXLAN routes are programmed using 2 hierarchical FEC levels in hardware, and the outgoing interface of a VXLAN underlay route changes from one forwarding chip to another, then it can negatively impact VXLAN overlay route convergence. (440233)
- Egress MAC ACLs are matched on the incoming packet's ethernet header and not on the rewritten ethernet headers. (458724)
- Dynamic port-channels coconfigured using LACP are not supported as destinations for monitor sessions (463186)
- When the feature "flow" in a TCAM profile is configured with only the "drop" action, the profile is rejected. The workaround is to add any other action to feature "flow". (486849)
- When a LAG is used as a destination for a monitor session, a given set of packets mirrored from a Rx source will hash differently compared to the same set of packets captured from a Tx source, due to hardware limitations. Consequently, symmetric hashing will not work as expected with a monitor

session which captures both Rx and Tx traffic from a single port, and using a LAG as a destination.

A workaround is to only capture traffic on the ingress (Rx direction).

Another workaround is to use two different monitor sessions, and destinations, for the two directions of the target traffic, and implement symmetric hashing on a second stage. (486994)

Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series

- When GUE is enabled with outer IP hashing, inner IP fields are not included in the load balance key of MplsOverGre transit packets in tc-counters TCAM profile. (491706)
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- When "ip hardware fib optimize" is enabled for a non-nibble aligned prefix-length, local-remote MPLS Pseudowires must have Control Word enabled for MPLS decap forwarding to work correctly. If Control Word is not enabled, then ETHoMPLSoETH traffic that ingress with the first nibble of the inner DMAC starting with "0x4" will be speculative parsed as IPv4oMPLSoETH and incorrectly dropped. (496624)
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- For Packets with GRE tunnel encap header, MTU enforced on the hardware is 3 bytes more than configured value in Cli. (498470)
Series Affected: DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- All asymmetric IRB traffic flows forwarded over EVPN VXLAN toward to a single multi-homed destination are not load balanced across the VTEPs that form the L2 ECMP group. However, note that the traffic destined to different hosts are statically load balanced across the VTEPs. (501343)
- VRF label termination into non-default VRF is not supported for multi-label MPLS packets. (504763)
- When BGP neighbors advertise Flowspec rules with police actions where the number of rules exceed what can be programmed in TCAM, policers will be allocated in hardware for all requested rules even though best-effort programming will only program the subset of rules that fit in TCAM. As a result, hardware policer resources will be unnecessarily consumed for rules which are not installed in hardware. (506688)
- SVI egress counters are not incremented for IPV4 or IPV6 multicast packets. (508471)

- Egress counter features configured via "hardware counter feature ..." do not include packet or byte counts for packets sent from the local CPU. (525380)
- Counter information flows from the forwarding chips to the tables that support TerminAttr in a very bursty manner. When averaged over a several minute window the rates will be accurate, but for shorter intervals calculated rates may oscillate significantly above and below the nominal rate.

For better accuracy, please use a longer interval to generate rate information. (540116)

- Hitless restart is not supported with Macsec proxy (540419)
- `show forwarding destination` may give incorrect results when the traffic would hit an egress IPv4 ACL using bridged IPv4 traffic or hit an egress IPv4 ACL using one of the following qualifiers: TCP established, fragment, and L4 port range. Note that TCP and UDP egress IPv4 ACL rules use the fragment qualifier. (554480)
 Series Affected: DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- Hitless restart of Layer 3 forwarding agent is not supported unless "ip load-sharing prefer local" configuration is removed. Note that this configuration is added by default on this system. (555772)
 SKUs Affected: DCS-7280QR-C36
- Route churn under hardware FIB exhaustion may cause BFD sessions to flap. (556089)
- On PTP enabled switches, CRC errors may occur in internal time synchronization. This can cause PTP time to be unstable. PTP time stability will recover on its own after a few seconds. (557370)
 SKUs Affected: DCS-7280CR2-60-F
- For products with paired QSFP100 (odd numbered) and QSFP+ (even numbered) ports, the even numbered port may experience link flaps when a transceiver with a different media type from the previous installed transceiver is inserted in the odd numbered port. (565969)
 SKUs Affected: DCS-7280QR-C72-F, DCS-7280QR-C72-M-F, DCS-7280QR-C72-M-R, DCS-7280QR-C72-R, DCS-7280QRA-C36S-F and DCS-7280QRA-C36S-R
- VRRP IPv6 using VRRP IPv4 MAC prefix may have unexpected packet forwarding behavior when the VRRP IPv6 and IPv4 have different states and one of the states is Master. (567504)
- Routes resolved over LAG sub-interfaces that have shape rate configured may forward traffic to CPU after SandL3Unicast agent restarts. The workaround is to flap the parent LAG by using the shut and no shut command. (570311)

- Different streams being rate-limited by the same group policer might exhibit bias towards some streams and fair policing across all the streams is not guaranteed. (578041)

Series Affected: DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series

- When all members of a port channel have "forwarding transmit disabled", traffic will not be forwarded to other VLAN members. (587648)
- While in Tap Aggregation mode, traffic steering service policy rules that match on fields from the MPLS header may result in false matches on packets without an MPLS header.

Remove the MPLS key-fields from feature 'tapagg mac' in the active TCAM profile. (592184)

- While in Tap Aggregation mode, traffic steering service policy rules that match on TCP flags may result in false matches when using the tap-aggregation-user-tcp-flags TCAM profile, or another profile based on this one.

Remove L4 key-fields from the `tapagg mac` feature in the active TCAM profile. (592191)

- While in Tap Aggregation mode, traffic steering service policies with VLAN matching may match untagged traffic erroneously.

Add the vlan-tag-format key-field to features that filter on VID in the active TCAM profile. (597603)

- VLAN matching in ACLs attached to monitor sessions may result in false matches on untagged traffic. (598186)
- For IPsec tunnel mode on Jericho2, the output of show ip security connection detail incorrectly shows the replay window as 16384, the correct replay window is 128 and fixed to that value. (609180)

- The SandAcl agent may restart repeatedly when more mirroring ACLs are applied than can be configured on any given FAP in the system.

A workaround is to reduce the number of mirroring ACLs applied; or apply smaller subset of mirroring ACLs to different source ports on different FAPs. (628026)

- Security counters are not updated correctly when using Vxlansec with SecTag2 encapsulation format. (630807)

- This platform only supports 64-bit versions of EOS. (646592)
SKUs Affected: 7816R3-FM, DCS-7816-CH, DCS-7816-SUP and DCS-7816-SUP1S
- When 'show forwarding destination' is used in an L2 forwarding scenario, the source MAC address of the packet can be learnt in the MAC address table. This can result in MAC moves if the MAC address was already learnt.

A workaround is to clear the MAC address after running the 'show forwarding destination' command. (650968)

- Packets originating from or forwarding via the CPU through an SVI are not counted by the egress SVI counter feature ("hardware counter feature vlan-interface out"). (688993)
- Counter sampling for ingress sFlow on sub interfaces is not supported. (718833)
- In an EVPN VXLAN active/active multihoming setup, configuring L2 subinterfaces as non-DF ports is not supported and may lead to unexpected SandMcast agent restarts. (739986)
- IPv4 and IPv6 strict mode Unicast Reverse Path Forwarding (uRPF) does not work if the route matching source IP address resolves via an ECMP nexthop. Such packets will be dropped. (757793)
- With scale QoS scale policy-map config with ACL rules, configure-replace might CLI time-out (766811)
- Configure replace with a scaled QoS policy-map config with multiple ACL rules may timeout. (766984)
- For tunnel destinations, the decision to fragment is based on unencapsulated packet size, while MTU enforcement is based on the encapsulated packet size. This may cause unencapsulated packets near the MTU packet size to be dropped instead of fragmented. (768410)
- At high combined scale of VLANs and pseudowires in VPLS or L2 EVPN, exhaustion of "managedTt inLif" resources may occur. This may prevent hardware programming of any configured patch panel patches, resulting in "Unprogrammed" status for one or more patches. It may also result in VPLS packets drops in the decapsulation direction. This situation can be detected by "show hardware capability".

The workaround is to reduce the scale of these or other features which use the managedTt inLif resource and then perform a shut/no shut on the affected patch panel or subinterface. (781993)

- "ip hardware fib next-hop sharing" is not supported with VXLAN. (795334)

- Traffic-policy 'redirect next-hop' and 'redirect next-hop group' actions on L2 interfaces are not supported. (821920)
- IP nexthop group entries are categorized as MPLS tunnels in the 'show hardware capacity' and 'show platform sand l3 summary' CLI commands. (823863)
- The VLAN interface counter feature is not compatible with dynamic VLANs used in EVPN VXLAN prefix routing setups. When the VLAN interface counter feature is enabled, a "Counters unavailable" message will be shown for each of EVPN VXLAN dynamic VLAN in the "show interface counters" CLI command output. (840994)
- The virtual-cable feature is unsupported, but the configuration CLI is not blocked. Configuring this unsupported feature will result in the Sand agent repeatedly restarting.

The workaround is to remove the virtual-cable feature from the running config. This will allow the Sand agent to recover. (893360)

- If 'switchport vlan tag validation' is configured, packets with more than 2 VLAN tags will be dropped. (898006)

DCS-7280R and DCS-7500R Series

- When a linecard is inserted in a chassis there can be a temporary marginal drop in fabric throughput for other cards. (205880)
- On the system with DCS-7500R linecards, if DCS-7500R2 linecards are inserted, routes may consume more resources in the hardware routing table. (215393)
- Jericho chip buffering DRAM test at agent startup does not catch bad parts while it should fail to force RMA of bad parts. (218971)
- IPv6 multicast routes are stored in two separate TCAM resources. The group IP is stored in one TCAM while the source IP is stored in another. Since one TCAM contains the source IP it will contain an entry for every multicast route. However, the group IP TCAM will reuse entries for duplicate group IPs. For example, if you have routes (S1,G1) and (S2,G1), one TCAM entry is required for the group IP and two TCAM entries are required for the source IP. This means it is possible that more TCAM resources are allocated to the source IP lookup than the group IP lookup. The scaling achieved during the initial distribution of multicast routes will be preserved. However, if the distribution later changes, optimal scaling for the new distribution may not be reached. (257364)
- The CPU traffic policy feature does not support filtering on IPv6 extension headers. (364996)

- Ingress ACLs and PBR are not supported for MPLS tunnel terminated traffic using VRF label. (366695)
- VRF terminated traffic is not sampled with sFlow. (369010)
- Drop-precedence thresholds are not supported on mirror session ports. (388932)
- We do not support incoming multicast LDP packets with a destination multicast MAC address.
RFC 5332 doesn't mandate the use of sending labeled multicast traffic in a multicast ethernet frame. Multicast MAC destination address could be used in upstream label assignment. (393406)
Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- Egress Ipv4 ACLs and Egress MAC ACLs cannot be configured together (443671)
- Egress Ipv6 ACLs and Egress MAC ACLs cannot be configured together. (443673)
- With `counters per-entry` for egress mac acl, we can uniquely count max of 3967 rules per fap. (460819)
- MTU is not enforced for MPLS packets forwarded by a MPLS swap route, which have a length one byte more than MTU. (464742)
- SandTcam crashes while changing the system TCAM profile to a profile that has feature 'tcp-mss-ceiling ip' on switches that has some linecards using Algomatch as the access-list mechanism and some other linecards using TCAM as access-list mechanism.

Workaround : Configure 'hardware access-list mechanism tcam' (474595)

- The two rate three color (trTCM) QoS policy-map counters can show incorrect values for packets marked green or yellow if the ingress traffic is broadcast or multicast traffic. (479267)
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7500R and DCS-7500R2 Series
- On TCAM profile change, sub-interfaces that have a L2 Protocol Forwarding profile applied or those sub-interfaces for which a sub-interface on the same physical interface has a L2 Protocol Forwarding profile applied, will briefly forward all traffic that is otherwise supposed to be trapped to CPU by default. (525564)
- If the CPU is included in monitor sessions with multiple destinations, it will only work if the device mirror0 is selected. If multiple mirroring to CPU sessions with multiple destinations are configured, all packets will be sent to mirror0. (525919)

- With RFC5549 ACLs configurations, the Egress IPv6 ACL deny logging on port channel subinterfaces stops working on performing agent SandL3Unicast restart. (549181)
- MPLS explicit NULL termination isn't supported with flex-route configuration containing non-nibble aligned prefix length. (576801)
- If MPLS tunnels are allocated in uncountable egress banks due to transiently high scale, they can be left in an uncountable bank which will prevent traffic counters from being aggregated for these tunnels. As a workaround, a hitfull restart of the forwarding plane agent will reallocate the tunnels in countable egress banks 1 & 2. (635152)

DCS-7280R2 and DCS-7500R2 Series

- On platforms with Jericho+ switches and with large number of pseudowire patches (2500 or more) configured, disabling hierarchical FECs (HFECs) using CLI command "ip hardware fib hierarchical next-hop disabled" may leave some patches unprogrammed and not operational.

Workaround is to save config with HFEC disabled to startup config and reload the switch. (569868)

- During normal operation, the internal time synchronization in a switch may log the message 'TIMESYNC_ERROR ... (slave signal CRC error)'. The internal time synchronization will correct this error without disrupting the normal operation of the switch. (818143)

SKUs Affected: DCS-7280CR2-60

- With "no platform sand ipv6 host-route exact-match" config, outer IPv6 IP-IP/GRE/GUE encapsulated transit packets will not have PBR and other similar policy-map features applied.

One workaround is to re-enable the "platform sand ipv6 host-route exact-match" config. (885787)

7500R3, DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series

- "Drop" MAC entries only drop packets whose destination MAC address matches the "drop" MAC entry. Packets whose source MAC address matches the "drop" MAC entry are not dropped by the "drop" MAC entry. (341123)
- Strict Priority scheduling for 400G interfaces don't work properly for 64 byte packets when traffic rate is above PPS limit of the chip. (388517)
- IP forwarded traffic that is mirrored to a GRE tunnel destination will have the PCP bits of their VLAN tag set to 0. (400909)

- Application of low rate shaping on multiple interfaces with line-rate traffic into the box will create a backlog in DRAM. Due to this majority of traffic will be served to egress from slower DRAM instead of SRAM. In the congestion state, even after the removal of shaping with line-rate traffic ingress from multiple ports, switch won't be able to deliver full throughput. The recovery from DRAM to SRAM will happen over a period of time once the congestion is removed. (414481)
- Set dcsp will not apply to mpls routed packets that egress with mpls labels (418791)
- Enabling Tx mirroring may cause the counter VOQ_SRAM_ENQ_RJCT_PKTS to increment, with reason QnumNotValid, as shown in the output from "show platform fap counters pipeline". There is no impact to mirrored or forwarded traffic. (459862)
- Unidirectional links are not supported at 40G speed on QSFP100 interfaces using the CRT50216 PHY. (467672)
 SKUs Affected: 7368-16C, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3-48CQM-LC, DCS-7280CR3-32D4, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-96, DCS-7280CR3-96-F, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-96 and DCS-7280CR3K-96-F
- SFlow samples of packets which are dropped by an ACL will have incorrect output interface indication. (470650)
- Certain speed groups are limited to one active serdes rate configuration. These speed groups are those that show only one serdes rate in the Active column of the output of "show hardware speed-group status". (486659)
- The sFlow flow samples generated for VXLAN decapsulated packets, have the first 6 header bytes of the sampled packet set to zeros. (502957)
- The "phy link serdes mapping direct" CLI command requires the 25g serdes rate to be configured, unlike what is displayed in "show hardware speed-group configuration". (523850)
 SKUs Affected: 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3K-48CQ-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32P4, DCS-7280CR3K-96, DCS-7280CR3MK-32D4 and DCS-7280CR3MK-32P4
- Encapsulating PTP packets over VXLAN tunnel is not supported. (526461)

- When Hardware Accelerated Sflow is configured the uptime in sFlow datagram is updated every 1/4 of milli sec, rather than every milli sec as prescribed in sFlow standard. (531769)
- Sequence number in sFlow datagram doesn't reset to one when hardware accelerated sFlow is disabled.
The sequence number in sFlow datagram gets reset to 1 for those sFlow sub-agent when the line-card is reset/ restarted on which the sFlow sub-agent is residing. (531983)
- With loose mode uRPF configuration, the switch doesn't drop packets with an uRPF lookup (using source IP) matching a drop route. (531998)
- When SandCounter agent is down, datagram counters are not accounted. Hence in show sflow output, the "Number of Datagrams" entry displays lesser value than expected. (536213)
- If all front-panel ports are populated, and traffic load is high, an improperly programmed current limit could be exceeded generating a spurious power-off. (536331)
SKUs Affected: DCS-7280CR3-96 and DCS-7280CR3K-96
- Symmetric hashing is always on for IPv6 addresses, when both source and destination addresses are included in the applied load-balance profile. (557714)
- In transparent two-step mode, PTP packets egressing the same port with the same sequenceId may cause transient jumps in correctionField.

The workaround is to use one-step transparent clock mode (564682)

- In Tap Aggregation mode, when both header removal and truncation are configured on a tap port, a packet will not egress the switch if it has a size less than 64 or has an invalid ethertype, after header stripping and truncation actions are performed. (733144)
- MAC access-lists are not applied on L2 protocol packets even when L2 Protocol Forwarding profile is applied. (735328)
- When using the Uniform-4x100G-2-R L1 module profile, the "/3" and "/4" interfaces will not support single lane speed modes (e.g. 50G-1). (786990)
SKUs Affected: 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3A-36P-LC, 7800R3A-36PM-LC, 7800R3AK-36DM-LC, 7800R3AK-36DM2-LC, 7800R3AK-36PM-LC, DCS-7280DR3A-54-F, DCS-7280DR3AK-54-F and DCS-7280DR3AM-54-F

7500R3, DCS-7130, DCS-7280CR3, DCS-7500R3 and DCS-7800R3 Series

- When operating under fabric egress replication mode (`platform sand multicast replication default fabric-egress`), multicast traffic egressing

(or dropped in the egress pipeline of) Ethernet ports is incorrectly counted as unicast traffic in the "Egress Queue Counters" section of the ``show interfaces counters queue`` command. In other words, all traffic will be counted as unicast. (488383)

DCS-7280R3A and DCS-7800R3A Series

- Packets mirrored from a TX source to a GRE tunnel destination will have incorrect metadata in the "key" field of the GRE header. (763403)

DCS-7020 Series

- The interface bin counter for both in and out packets with sizes 1523-Max may be inaccurate. Packets with sizes between 9217-9236 are not counted by the switch for interfaces Ethernet 1-48. (188121)
- Mirroring and sFlow can not be supported on the same source interface at the same time due to a hardware limitation. (209060)
- AES128/SHA-1 and AES256/SHA-256 are the only supported encryption/authentication algorithm pairs. (342005)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Using an Ipsec tunnel interface for policy-based routing nexthop is not supported. (378861)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Ipsec tunnel interfaces are not supported as part of nexthop-group. (378890)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Dynamic NAT, AKA Port Address Translation, is not supported for IPSec tunnels. (384624)

DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

- L4 ops and fragments rule cannot be programmed together. (94197)
- While in Tap aggregation mode enabling either VN or BR tag stripping will disable VXLAN header stripping functionality. Traffic steering on the inner IP header of VXLAN packets will also be disabled.

Removing the VN/BR tag stripping configuration will restore VXLAN stripping and traffic steering functionality. (175829)

- An IPv6oMPLS packet terminated by L3 EVPN MPLS configuration will not be forwarded. The workaround is to disable IPv6 exact match host routes using 'no platform sand ipv6 host-route exact-match' and disable IPv6 route optimizations for prefix length 128. (201414)

Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series

- Vxlan decapped packets will not match rules that match on 'tcp established' in egress acls (226582)
- VXLAN and GRE tunnels cannot be configured at the same time. (248239)
- If a PTP packet is mirrored to a GRE destination which has time stamping enabled, the timestamps on such packets are not reliable. (280362)
- If a packet is mirrored to a GRE tunnel, but the forward copy is trapped to the CPU, then the ingress VLAN ID carried in the GRE key will be set to 0 when the mirrored copy is encapsulated in the GRE envelop. (434871)
- In the event that two rate three color (trTCM) policers are configured for an incoming broadcast or multicast traffic, drop precedence (dp) classification may not happen correctly for those packet that are subject to special Copp class policers. A portion of the ingress traffic is marked with as yellow. (481688)

Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7500R and DCS-7500R2 Series

- Egress MAC ACLs do not work on interfaces that are not front-panel interfaces, even though they can be configured. (483583)
- Mirroring to GRE destinations will not work if VxLAN encapsulation is enabled (548429)
- MPLS label popped packets will not match rules that match on 'tcp established' in egress IP acls (565015)
- PTP v1 packet forwarding is incompatible with both Tap Aggregation BR/VN tags stripping and GUE packet decapsulation.

The workaround is to disable these features before enabling `ptp forward-v1` configuration. (572447)

- BGP Flowspec rules will not work on SVIs if the TCAM profile does not include a "port qualifier size N", N > 1 configuration for the flow-spec features. (588565)
- L4 destination port matching in egress IP ACLs does not work on untagged bridged traffic. (609461)
- Hitless traffic-policy update is not supported on Sand Gen3 systems. (622066)
- When the non-head-fragment field qualifier is not present for the traffic-policy feature on Sand Gen3 chips, L4 qualifiers will not match on any fragment packets, including head fragments.

As a workaround, add an explicit rule for handling fragmented packets.
(628773)

- L2 EVPN MPLS won't work in presence of TCAM feature "VLAN Editing Extended-64" in its TCAM profile. (648996)
- In Tap Aggregation mode, egress IP ACLs are not supported on tool interfaces that have egress truncation configured. (668716)
- The following two configurations are incompatible:
 - IPv6 route optimization with "ipv6 hardware fib optimise prefix-length 128 []" for /128 prefix-length.
 - Using "no platform sand ipv6 host-route exact-match" to disable IPv6 host routes in exact match.

The workaround is to remove either of the two commands from the configuration. (671966)

- Enabling IPoGUE outer UDP hashing may break IP-in-IP hashing. (676230)
- On adding a member to a port-channel sub-interface that has a named QoS-map configured, the ingress packets get classified based on the parent interface configuration for a short duration of time. (680689)
- Egress Port ACL does not work for traffic that is bridged after VXLAN decapsulation. (688880)
- TCAM feature 'acl port ipv6 egress' does not add support for egress IPv6 PACLs when added to a profile. (699451)
- In an EVPN IRB setup, enabling both VXLAN and MPLS is not supported. Traffic loss is expected in this configuration. (703701)
- BFD/BGP/OSPF/IS-IS protocol packets received over VXLAN take generic CoPP queues to CPU instead of the protocol specific ones. (719542)
- L2 EVPN MPLS configuration will not work when the same trunk port is used to carry traffic towards both the Provider Edge (PE) and the Customer Edge (CE). (727163)
- When IP PIM sparse mode is enabled on the VLAN interface, egress port ACLs applied to switch ports that are a member of that VLAN match on random fields for bridged multicast traffic. (728267)
- Egress IPv4 ACLs do not apply to unknown multicast traffic (728269)
- QoS policy-maps are not supported on dot1ad, unmatched or untagged subinterfaces. (732743)
- QoS policy-maps are not supported on dot1ad, unmatched or untagged subinterfaces. (733077)

- IS-IS over VXLAN is not supported when the CLI command "vxlan decap floodset unknown-multicast unknown-unicast exclude cpu" is configured. (749279)
- The "Uniform-4x10G-1" L1 module profile is incompatible with the "boot port numbering octal interfaces 4" CLI configuration.

Users should remove the "boot port numbering octal interfaces 4" command from their configuration and apply the Uniform-4x100G-2-R" L1 module profile on card's where 4 interface octal port numbering schemes are desired. (779375)

SKUs Affected: 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3A-36P-LC and 7800R3A-36PM-LC

- Changing the profile for hardware accelerated sflow via "sflow hardware acceleration profile" will not take effect until the switch is restarted. (851412)
- The IPv6 decap-groups feature is unsupported when IPv6 host routes are programmed in exact match tables. There is a CLI warning when configuring IPv6 decap-groups if this feature conflict happens, but a system log message is not generated.

The workaround is to disable the IPv6 host routes in exact match table feature, which is enabled by default, with "no platform sand ipv6 host-route exact-match". (860223)

- Enabling a GRE destination on a monitor session while traffic is running on the monitor source may cause the DCH unreachable discard counter to increment by a small amount, due to some superfluous traffic being mirrored into the fabric without a valid destination. (882484)
- QoS features configured using tx-queue mode of interface are not supported on ports having 2 priority level.
"show platform fap mapping interface <interfaceName>" shows number of priorities supported by the port under QPairs column

List of unsupported QoS features on the tx-queue of the 2 level priority port include:

- 1) Shaping/Scheduling
- 2) WRED
- 3) ECN
- 4) PFC
- 5) Latency based drops. (893100)

7500R3, DCS-7130, DCS-7280R3, DCS-7280R3A, DCS-7500R3, DCS-7800R3 and DCS-7800R3A Series

- Priority tagged packets (802.1Q tagging with the VLAN set to 0) being reflected by an Ethernet reflector interface configured in the egress direction with MAC address swap action will have the priority tag stripped from the packet before they are transmitted. (354190)
- Only port-channel interfaces are supported as peer links. Single member port-channel interfaces are supported. (409958)
- If a packet that is decapsulated with a decap group is mirrored to a GRE tunnel, the packet's ingress VLAN will be set to 0 in the upper 2 bytes of the GRE key. (430986)
- Egress ACLs does not match packets with UDP ports 4754 and 4755. (455922)
- Traffic matching PBR rules will be dropped if it violates the MTU of the PBR destination port (471548)
- A learnt MAC address won't age out in case there is only continuously running unicast RPF drop traffic from the source MAC. The workaround is to clear the dynamically learnt MAC from the MAC address table. Once the MAC address is cleared, the MAC address won't be learnt if all traffic with the source MAC address is dropped by unicast RPF. (536215)
- If there are multiple mirror to GRE monitor sessions sharing the same GRE tunnel parameters (i.e. source and destination IP address, DSCP values, TTL, protocol, all matching), and if any of these sessions has rate limiting enabled, then all these monitor sessions may either have rate limiting functionality disabled, or share the rate limit configured in one of these sessions, following a switch reload.

A workaround is to remove these monitor sessions and reconfigure them again. (548973)

- During EVPN MPLS encapsulation, the VLAN priority bits of payload Dot1Q header will be marked using the traffic-class to CoS map configuration in the global QoS map. (550987)
- QoS policy-map is not supported on Pseudowire decapsulated traffic flows. (574600)
- When there is congestion in Tc6/Tc7 across multiple queues. Tc6/Tc7 Queues can hog up entire DRAM. This can lead to DRAM rejects of lower TC traffic of unrelated stream thus causing some discards. In the current case we had 2 to 3% lower instead of 100% throughput of data traffic (591523)

- The switch does not learn the MAC address of routed packets ingressing on VLAN interfaces when using IPv4 route optimization with an internet profile configuration or with a single prefix length compression configuration.

The workaround is to set the ARP aging timer to be smaller than the MAC address

aging timer, so that ARP packets refresh MAC addresses. (592117)

- Renamed TCAM Qualifiers l4ops_ip_pacl, l4ops_ip_qos, l4ops_ip6_pacl, and l4ops_ip_racl into l4ops_24b, l4ops_7b, l4ops_3b and l4ops_18b. This can cause customer scripts using `show platform fap pmf` command to fail if they are using any of the above renamed TCAM qualifiers. (595868)
- On Gen4 Faps , SSO for TCAM based ACL rules is hit full. There is a brief window where ACLs are not applied to the flows. (606756)
- On Gen4 Faps , SSO is hit full for TCAM based ACL rules. (606764)
- When outer IP hashing is enabled, GUE decap followed by MPLS NULL label termination is not supported. (607977)
- Configuration of "encapsulation dot1q vlan" is only supported on sub-interfaces. (608383)
- Software forwarding of MPLS-over-UDP terminated packets is not supported. (613102)
- Added support for L4 port match on VXLAN IPv6 underlay transit packets (forwarded on either outer L2 or L3 header). (614691)
- If a packet matches both a Tap Aggregation steering rule, as well as either a security ACL or traffic-policy rule to drop the packet, the steering policy counter will still increment. (620429)
- SandTm restart is hitful for named TC to DSCP QoS maps. (632481)
- SandOam may restart with 200+ DM sessions running at 10 milliseconds. The workaround is to use higher tx-interval like 100 milliseconds, 1 second, 10 seconds, 1 minute, 10 minutes. (642106)
- When the transmit interval of CCMs is increased, some of the remote MEPs are shown as unreachable even though CCMs are being received from the remote MEPs. As a side-effect, this also causes the RDI bit in the generated CCMs to be set.
There are a couple of workarounds:
 1. Wait for 5-10 seconds before triggering a CCM interval change on the peer device. This can potentially avoid this issue completely.
 2. If the issue is still seen, restart the SandOam agent. (643283)

- SandL3Unicast on the peer DUT causes a few RMEPs to be marked as unreachable on the local DUT. A SandOam restart on the local DUT recovers from the problematic state. (646125)
- IP UDFs in ACLs may not work correctly on IPv6 packets with a routing extension header when matching on bits beyond the IPv6 header. (650286)
- VXLAN is not supported with greater than 256-way ECMP in the underlay. (652266)
- In case of Delay Measurement, the device does not work as responder to DMM transmitted by any other device which configures the TLVs in the DMM. The device is incapable of copying the received DMM's TLVs to the DMR and if the other device expects the configured TLVs in the DMR then the device can reject the such DMRs. (654226)
- While not in Tap Aggregation mode, packets mirrored to multiple destination ports may be actually forwarded out of the mirroring destinations, or may be malformed.

Do not use mirroring to multiple destinations outside of Tap Aggregation mode. (670046)

- GUE tunnel decapsulation does not support UDP destination port 2152. (675837)
- In RFC2544 throughput test, oversubscribing the ingress interface causes hardware test engine lockup and test traffic won't be generated. Restart the SandFapNi agent. (684101)
- A pseudowire with Cos-To-Tc or Dscp-To-Tc QoS map applied will classify all of the decap traffic to default traffic-class of 1 in the presence of entropy or the null label. (694370)
- While RFC2544 test is in progress, received and sent packet counts may not match. (694683)
- If egress MAC/IPv4 ACL is applied to the front panel port and egress IPv4/IPv6 ACL is applied to SVI with ACL sharing enabled for the SVIs, only the ACL applied to the SVI will match the routed traffic. (695784)
- When using IPv4 route optimization using "ip hardware fib optimize .." command, the number of VRFs supported on the system is reduced to 30 non-default VRFs. (720042)
 SKUs Affected: 7800R3-36D-LC, 7800R3-36P-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3-48CQM-LC, 7800R3-48CQM2-LC, 7800R3-48CQMS-LC, 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3AK-36DM-LC, 7800R3AK-36DM2-LC, 7800R3K-36DM-LC, 7800R3K-48CQ-LC, 7800R3K-72Y-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-36S, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4A, DCS-7280CR3K-32P4,

DCS-7280CR3K-32P4A, DCS-7280CR3K-96, DCS-7280CR3MK-32D4,
DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4S,
DCS-7280DR3-24, DCS-7280DR3K-24, DCS-7280PR3-24, DCS-7280PR3K-24,
DCS-7280SR3-48YC8, DCS-7280SR3K-48YC8, DCS-7280SR3K-48YC8A and
DCS-7280SR3M-48YC8

- The MPLS no-php mode is not supported with L3VPN and VPWS services. (720213)
- The packet counting logic for CFM Loss Measurement adheres to Version 1 of the protocol although the LMM is sent with version 0. This can result in an incorrect loss being reported. (720859)
- OAM doesn't generate SLR's in response to SLM's if Test ID in the SLM is more than 1048575 (0xFFFFF), instead it generates LMR's. (726371)
- The TLVs present in the Loss Measurement Message (LMM) are not copied to the Loss Measurement Reply (LMR) that is generated in response. This can cause the LMM initiator device to invalidate the LMR. (727007)
- The "Flag" field in the Loss Measurement Reply that is generated in response to an ITU-T.Y1731 version 1 Loss Measurement Message is set to 0, indicating an on-demand session. (727056)
- CFM Loss Measurement will report incorrect loss if the loss between consecutive LMM/LMRs is greater than 32K packets. (729327)
- When using IPv4 route optimization using "ip hardware fib optimize .." command, the number of VRFs supported on the system is reduced to 14 non-default VRFs. (731721)
 SKUs Affected: DCS-7280SR3-40YC6 and DCS-7280SR3E-40YC6
- Traffic throughput on an internal recirculation interface is capped at 100Gbps. (738474)
- Three level hierarchical IP nexthop groups are not supported. (751053)
- When policy configuration changes are made with segment-security, some counts may be missed or misclassified under "show segment-security counters" for up to 30 seconds. (769949)
- Post SandOam restart RFC2544 generator test is initiated for subinterfaces that are operationally down. (772320)
- VXLAN bridging of packets that are received on a 802.1Q tunnel port with more than one 802.1Q tag does not work (782097)
- Traffic sent towards control plane for unresolved ARP or next-hop which is marked with software drop-precedence 1 gets dropped. (792663)

- CFM PDUs received on the Up-MEP interfaces that are to trapped/dropped get counted in the following platform counters: "Tx Total Err On Sop Pkt Counter", "Tx Total Pkt With Err Counter", "Total Byte With Err Counter", "Tx Total Err On Sop Byte Counter", "TxTotalPktWithErrOnEopCounter" and "TxBytesWithErrOnEopCounter". (796338)
- When ETH/MPLS/IP/GTP packets are MPLS terminated, hashing on GTP TEID will not occur.

A workaround to hash on GTP TEID is to use the custom UDP payload hashing feature. (813466)

- Hardware resources are exhausted when hardware accelerated sFlow is enabled with scaled subinterfaces, predominantly with port-channel subinterfaces. Possible workarounds are disabling sFlow hardware acceleration or disabling subinterface ifindex stamping by running "no sflow sample input subinterface". (813602)

Series Affected: 7500R3, DCS-7280R3 and DCS-7800R3 Series

- After deleting the operational TCAM profile and then reloading the switch, no TCAM banks would be granted.

The workaround is to re-configure the TCAM profile. (836367)

- When the device is configured with 'vxlan qos dscp ecn propagation decapsulation disabled', the inner DSCP for VXLAN decapsulated packets is retained even if the ingress port is in cos or untrusted mode and global DSCP rewrite is enabled. (836505)
- Mirror copies, or sFlow samples, of decapsulated VXLAN traffic will have bytes missing from the underlay headers. (843636)
- In Tap Aggregation mode, if an incoming packet has 3 or more VLAN tags and 2 tags are selected to be removed via the Tap Aggregation DOT1Q tag removal feature, the packet may be malformed at egress with an additional VLAN tag.

Using the DOT1Q remove feature to remove a single tag is supported. Removal of additional tags can be implemented by recycling the packet via a tap-tool and traffic loopback feature. (850175)

- The switch does not learn the MAC address of routed packets ingressing on VLAN interfaces when using IPv4 route optimization involving two prefix lengths.

The workaround is to set the ARP aging timer to be smaller than the MAC address aging timer, so that ARP packets refresh MAC addresses. (868827)

- The following two configurations do not work together:
 1. IPv4 optimize configuration of internet profile or any other compression configuration - "ip hardware fib optimize prefixes profile internet" or "ip hardware fib optimize prefix-length 21 compress <>".
 2. PBR policy matching on nexthop-group.


```
policy-map type pbr PMAP1
  10 match ip any any nexthop-group <>
```

The work around is to do either of the below:

1. Remove the optimize configuration
 2. Use an optimize configuration without compression
 3. Use urpf-internet profile. (893189)
- DSCP remarking does not use the custom CoS-To-Tc map. The egress packet DSCP value is derived using the global CoS-To-Tc map (on ingress) and global TC-To-Dscp map (on egress) (902952)
 - If routes are learnt via BGP, RFC5549 traffic egressing from L3 sub interfaces and SVIs is not subject to egress IPv4/IPv6 ACLs.

Egress IPv4/IPv6 ACLS on RF5549 traffic will be functional after configuring "forwarding-type" key field in the "acl vlan ipv6 egress" TCAM feature in the operational TCAM profile. Some unused fields from "acl vlan ipv6 egress" feature of the operational TCAM profile needs to be removed to make space for this "forwarding-type" key. (903500)

- CFM protocols do not work if egress sFlow is active on the interface on which CFM MEP with direction Up is defined.
Disable sFlow on CFM UP MEP interface. (905833)
- RFC2544 Initiator benchmarking tests with 64 byte packet size is only able to test max throughput of 256Gbps because of hardware limitation.
400Gbps RFC2544 Initiator throughput testing is supported for packet sizes greater than 100bytes. (915701)
- A traffic-policy match which includes 'encapsulation gtpv1 ipv4', but no inner match criteria will match on both IPv4 and IPv6 traffic. This same behavior is also seen for 'encapsulation gtpv1 ipv6'.
Workaround: add qualifiers inner-src-ip-label, and inner-dst-ip-label, to feature "traffic-policy port ipv4/6" in the TCAM profile (917902)

DCS-7280R and DCS-7280R2 Series

- CBF does not work in AlgoMatch mode on some custom TCAM profiles (603206)

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7260X, DCS-7260X3 and DCS-7300X Series

- Untagged packets received on a port with PFC enabled may be assigned the wrong priority. In case of congestion, PFC frames may not be sent or may be malformed.

The issue is observed when a port uses a default CoS value that is mapped to a traffic class with a different value. For example, if the default CoS is 0 and CoS 0 is mapped to traffic class 1. The workaround is to change the default CoS value on the port or the global QoS mapping such that the default CoS and the traffic class match. (23922)

- System does not stop transmitting traffic for the exact amount of time quanta received in PAUSE or PFC frame. This can lead to packet loss during congestion if the peer switch does not have enough buffers. The workaround is to use XON/XOFF instead of time based flow control, or to increase the time to account for one MTU sized packet. (27190)
- QoS shaping and guaranteed bandwidth will only work in store-and-forward mode. They are not supported in cut-through mode. This includes 'shape rate X' and the 'bandwidth guaranteed X' commands. (56790)
- Accelerated Software Upgrade (ASU) from 4.13.x (x < 6) to 4.13.6 and above requires a special upgrade procedure due to a software image format change. (90097)
- 10 Gbps packet replication on all SFP+ ports at the same time may lead to packet discard. (91149)
- Egress L2 MTU violations in cut-through mode may result in unexpected discards of other frames. (95577)
- VXLAN encapsulated packets cannot be TX mirrored at the egress properly. The mirrored packets will have an incorrect MAC header. (97272)
- When IGMP snooping is enabled a single unknown multicast floodset is shared across all VLANs with IGMP snooping enabled. This can cause unexpected flooding to a trunk port that has no multicast routers on a given VLAN, but has multicast router attached to the same trunk port on another VLAN. (105188)
- When "hardware access-list resource sharing vlan in" is configured, Link Local Multicast traffic will not be intercepted by the ACL attached to the VLAN interface. (105504)

- There may be a momentary traffic disruption every time the nexthop specified in a PBR policy resolves to a new destination: affected packets may be forwarded by normal L3 lookup, or be dropped.

The system will recover automatically; the PBR policy will return to route packets as per the new nexthop resolution momentarily. (105670)

- When any PBR policy is applied, packets that violate the egress MTU and are destined to flood the VLAN will end up being flooded instead of being sent to the CPU. (105680)
- Subinterfaces cannot be used to send or receive VXLAN encapsulated packets. (105767)
- When "hardware access-list resource sharing vlan in" is configured, changing the ACL on a VLAN interface:
 1. From an ACL that is not shared with other VLAN interfaces to one that is shared with other VLAN interfaces can cause deny traffic to get permitted until the new ACL takes effect.
 2. To another ACL that doesn't fit in the TCAM can cause deny traffic to get permitted until the system reverts back to the previously applied ACL. (106646)
- If IP-in-IP decap groups are configured, "qos trust dscp" configuration is not supported on traffic matching the decap groups. (107579)
- Vxlan and MPLS features may not be configured at the same time. (109330)
- When "hardware access-list resource sharing vlan in" is configured and TCAM is close to being full, restart of the forwarding agent or a reboot of the box, can cause some ACLs to not get programmed properly. (109876)
- The Accelerated Server Upgrade feature is not supported with the "hardware access-list resource sharing vlan in" configuration. (110114)
- IGMP snooping is not supported on Vxlan enabled VLANs. The multicast packets will always be flooded locally and to remote VTEPs. (110479)
- A front-panel or fabric port may get stuck in an errored state and drop traffic egressing the port. When this happens, the forwarding agent log will contain "port <internal port #> start error detected". This condition will persist until it is manually cleared. A forwarding agent restart is required to clear the error and forward traffic normally. (112051)
- MPLS is not supported on subinterfaces. (113110)
- Packets being VXLAN encapsulated are constrained to a single underlay physical nexthop per physical egress port.

Topologies involving SVIs over trunk ports, or non point-to-point routed ports towards the core will not work optimally, the encapsulated packets will be sent to the wrong next-hop for all but one of the VTEPs.

Topologies involving virtual VTEPs connected via a downstream L2 switch will not work, as the receiving vswitch will not have the capability to route the packet to the right VTEP. (113722)

- Port ACLs are not supported on VXLAN terminated traffic. (113726)
- If the configured RACLs do not fit in the TCAM, then disabling the RACL sharing feature might cause removing a RACL from an individual VLAN interface to fail.

To recover from this state, delete/remove the access-list and then reconfigure the access-list. (114028)

- L2 flooding of BUM packets on Vxlan vlans uses L3 replication tables, which are also used by IP multicast routing. The replication tables have an entry for each vlan, physical port combination (lag or non lag). We will run out of this table resources if we need more than 64k entries. (114517)
- If the switch is in cut-through mode, L2 MTU violations will not be counted if the packet headers are changed while forwarding. (116760)
- QoS policies are not supported for L3 Unicast flooded packets (routed packets with nexthop MAC not learned) (123883)
- When a PBR is attached to any interface, Layer-3 packets which are getting flooded in the egress VLAN will not be treated with Egress Router ACL if one is configured on that egress VLAN. (133144)
- VXLAN routing is not supported on systems with BCM56750 fabric chips. Such switches can be identified using the CLI command "show platform trident l3 summary", which will show bcm56750 for "Fabric chip model". (134625)
- Mirroring of incoming multicast traffic to a GRE tunnel may result in data traffic loss on the egress interface in an oversubscription scenario. (139591)
- Egress IP/IPv6 router ACL is not supported for VXLAN encapsulated traffic. (148642)
Series Affected: DCS-7050X3 and DCS-7300X Series
- IP ACLs configured on SVIs to match routed IP multicast traffic may also match bridged IP multicast traffic in that SVI. (152523)

- Up to 4 mirroring sessions are supported.

Each direction mirrored on a source port counts towards that limit of 4. For example, if a port is configured with rx and tx mirroring (the default), this counts as 2 consumed sessions. (158490)

- 1) Multicast traffic matching reserved IPv4 multicast address range 224.0.0.X is not counted by L3 interface counters.
- 2) L3 interface counters not supported for IPv6 multicast traffic.
- 3) Unicast traffic to CPU is counted by L3 interface counters only if routing is enabled (i.e. ip routing for IPv4 traffic and ipv6 unicast-routing for IPv6 traffic) (160930)
- Mirroring of Vxlan encapsulated packets in the egress direction is not supported. (167189)
- If a trunk port is configured as a member of both the underlay and overlay VLANs, the following limitations apply:
 - If packets received on the trunk port are flooded in the overlay VLAN, the VxLAN encapsulated copy is not forwarded on the ingress trunk port.
 - Vxlan encapsulated packets received on the trunk port, if flooded in the overlay VLAN after decapsulation, are not forwarded on the ingress trunk port. (185321)

Series Affected: DCS-7050CX3, DCS-7050SX3 and DCS-7300X Series
- If ingress ACL is applied on an interface, traffic sent to CPU from that interface is falsely counted as InAclDrop in "show int counters ingress acl drop". (203465)
- Clause 37 auto-negotiation on SFP28 and QSFP28 ports is not supported. A workaround is to use 1G forced speed configuration on the affected ports. (206778)

Series Affected: 7300X3, CCS-720XP, CCS-750, DCS-7050CX3, DCS-7050SX3, DCS-7050TX3, DCS-7050X3, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7260CX3 Series
- When "reload fast-boot" is performed from a release prior to 4.17.2, switches with BGP peering connections with the switch undergoing "reload fast-boot" could experience BGP KeepAlive timeouts before they see port flaps. This could result in longer than expected data plane downtime during the upgrade. (214700)

Series Affected: DCS-7060X and DCS-7060X2 Series
- A multi-bit ECC error in forwarding ASIC packet memory cannot be automatically corrected. A hitful restart of the slice agent managing the forwarding ASIC or a reboot of the system is required to clear the condition. (214950)

- If ip-ttl is the only enabled field for port channel hashing, after reload hitless, packets may get hashed to a different link than before the reload. (223568)

Series Affected: DCS-7050TX and DCS-7060X Series

- A configured mirroring ACL may not be applied in hardware. This can be confirmed by checking the output of 'show platform trident tcam mirror'. (236908)
- When a VXLAN encapsulated packet is received on a VXLAN VLAN where PBR policy is applied, the switch will not PBR forward the decapsulated packet. (244093)
- Sflow output interface determination will not work when the output interface is a subinterface. (250424)
- Cut-through mode is not supported on SFP+ interfaces. (252731)
Series Affected: DCS-7050X3, DCS-7050X4, DCS-7060X4, DCS-7060X5 and DCS-7260X3 Series
- A switch may not be able to bridge or route VXLAN traffic if it hits one or all of following conditions:
 - Nexthop resources are exhausted as indicated by syslog VXLAN_HWHOP_NEXTHOP_RESOURCE_FULL.
 - MAC table overflow as indicated by show platform trident mac-address-table missing.

To recover from this condition reduce the config scale and flap the VXLAN interface. (253601)

- Flexible port-channel hashing may not hash accurately when the inner L4 header option is configured for IPV6 GRE packets. (267487)
- If a packet stream is policed by multiple policers, the policed rate may be lower than any of the configured policer rates. (271046)
- To perform filtered mirroring of a packet based on VLAN from an interface configured to perform VLAN translation, the ACL must be programmed with the translated VLAN. The mirrored packet will be the same as the packet at the source interface. (278599)
- BGP neighbor sessions may flap when URPF is configured on an interface. (279113)
- To perform filtered mirroring of a packet based on VLAN from a subinterface source, the ACL must be programmed with the internal vlan of the subinterface. The mirrored packet will be the same as the packet at the source interface. (280943)

- Matching on inner VLANs is not supported in mirror ACLs. If an entry is specified matching on inner VLAN, it will be ignored and permit all packets. (284371)
- SSU is supported when upgrading from the associated release in each release train: 4.18.8, 4.19.8, 4.20.7, 4.21.0. (289548)
- Issuing "no shutdown" on a member link of port-channel may send duplicate packets for sub-second on that member link (291514)
- DirectFlow (and features such as the MacroSegmentation Service that use DirectFlow) are unsupported on the 7260, 7300 series of switches. (296715)
- When primary vlan and secondary vlan are part of different MST instances, the hardware programming of lag will not happen. (300119)
- Not all interfaces in the groups Ethernet49-Ethernet52 and Ethernet53-Ethernet56 can be broken out in to 4x25G, 4x10G or 2x50G mode due to MAC resource limitation on the system. Interfaces without MAC resource will be marked inactive and 'STRATA-4-HW_PORT_RESOURCE_FULL' will be logged in syslog.

Speed change on interfaces in the groups Ethernet49-Ethernet52 and Ethernet53-Ethernet56, can free up the MAC resources to make inactive interfaces active in that group. (306121)

SKUs Affected: DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050TX3-48C8-F and DCS-7050TX3-48C8-R

- When MLAG peer MAC routing is enabled, OSPF neighborship over VXLAN overlay may never get established. (349242)
- Ip-length based rules are not supported on Egress ACLs. (353247)
- Packets dropped in the egress pipeline may still be egress mirrored successfully. (357609)
- Routed unicast traffic egressing a SVI for which destination MAC is not learnt are sent to CPU and then flooded on the egress VLAN. (368228)
Series Affected: 7358, 7368, 7388, DCS-7050X4, DCS-7060PX4, DCS-7060X4 and DCS-7060X5 Series
- Ports with ingress MAC ACL still learn MAC addresses although traffic may be dropped (400211)
- Reload hitless from EOS-4.22 or earlier results in /2 and /4 ports on Ethernet 1-12, 21-44 and 53-64 to remain inactive when /1 and /3 are configured in 10G or 25G. In order to activate /2 and /4 ports, configure 40G, 50G or 100G /1, then configure 10G or 25G back. These ports are also activated following a forwarding agent restart. (409057)
Series Affected: DCS-7260CX3 Series

- "reload fast-boot" is not supported with virtual IP address config. (411323)
- Speed change on a port with macsec enabled is not supported and could cause the port to get into a persistently un-authenticated state. Workaround is to disable macsec on the port, change the speed and then re-enable macsec on the port. (415737)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- When a VXLAN VTEP is reached through a resilient ECMP route, (VXLAN bridging or routing) traffic to that VTEP may get dropped in HW. Workaround is to remove the Eesilient ECMP config. (421216)
- Performing SSU from any release prior to 4.22.2 can lead to extended outage and mis-forwarding. (424812)
SKUs Affected: DCS-7050SX3-48YC8-F and DCS-7050SX3-48YC8-R
- IP ACL match over MPLS Encapsulated Packets is not supported. (426413)
- During Mpls PHP operation inner payload type gets ignored. (428546)
- When both "hardware counter feature vni decap" and "hardware counter feature vni encap" are enabled, the egress flexcounter pools will only be released after both features are disabled. (441887)
Series Affected: 7300X3, 7320X, 7368, CCS-720XP, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7260X, DCS-7260X3 and DCS-7300X Series
- When 'hardware nat static access-list resource sharing' is in the switch's config, and a static NAT filter ACL and a security ACL have overlapping rules, the security ACL's hit counter might not increment in all cases. Regardless of the hit counter, both the security ACL rule and the NAT filter ACL rule will be applied correctly.

Workaround is to disable static NAT ACL resource sharing with:
'no hardware nat static access-list resource sharing'
Alternative workaround is to disable port ACL in VFP with:
'platform trident tcam port-acl vfp disabled' (450080)
- When an interface undergoes a hitless speed change, sFlow sample numbers and sequence counts for that interface may be reset. (452949)
Series Affected: 7300X3, 7368, DCS-7050X3, DCS-7060X4 and DCS-7260CX3 Series
- BUM traffic is not sampled by egress sFlow. (459902)
- If the incoming packet is VLAN tagged, then the truncated mirrored packet will be 4 bytes shorter than expected. (463986)
Series Affected: 7300X3, 7368, CCS-720XP, DCS-7050X3 and DCS-7060X4 Series

- If virtual mac address of the switch is learnt against a front panel port, it may not age out. The workaround is to clear the mac address manually by issuing 'clear mac address-table dynamic vlan <VLAN> address <ADDRESS>' command. (470965)
- When IPV4 and IPV6 PBR are both programmed in TCAM, forwarding plane agent restart may result in IPV6 PBR rules not being programmed in hardware. (473582)
- By default, configuring a static NAT rule with an empty access-list is treated the same as no access-list: for source NAT, the switch will ignore the destination IP, and vice versa for destination NAT. With "hardware nat static access-list resource sharing" enabled, configuring a static NAT rule with an empty access-list is treated the same as an undefined access-list: some hardware resources may be consumed, but the traffic will be forwarded untranslated. (473783)

Series Affected: CCS-720XP, DCS-7050CX3 and DCS-7050X3 Series

- The total length field in the IP header is not populated correctly when operating in cut-through mode. The switch should be operated in switch-and-forward mode when mirroring to a GRE tunnel. (475273)
- Mirrored packets that violate the MTU requirements will not be fragmented. (481513)
- When persistent port security is enabled on an interface, MAC addresses will persist when an interface is down. While the interface remains down, traffic destined to these MAC addresses will be dropped.

Persistent port security is enabled by default, and can be disabled with 'switchport port-security persistence disabled'. (485828)

- If dynamic load balancing is enabled for an ECMP group, it's member link flap might cause poor traffic distribution. Traffic distribution may improve over time. (487595)

Series Affected: DCS-7060X4 and DCS-7260X3 Series

- If dynamic load balancing is enabled for ECMP groups while the traffic is flowing, then the load on the ECMP members may not be distributed uniformly. Traffic distribution may improve over time. (490418)

Series Affected: DCS-7060X4 and DCS-7260X3 Series

- When dynamic load balancing is enabled, If multiple traffic flows destined to the same ECMP group arrive at the same time, then the load on the ECMP members may not be distributed uniformly. Traffic distribution may improve over time. (490716)

Series Affected: DCS-7060X4 and DCS-7260X3 Series

- Packets sent from the CPU cannot be egress mirrored unless they are being sent out of an SVI. (495146)
- If the speed of an interface changes with traffic running, it is possible for "show interface counters" and "show interface counters fast-poll" to diverge on that interface. (502252)
- Mirror on drop does not mirror dynamic NAT traffic drops. (506047)
- Egress mirroring of multicast packets that are going to be encapsulated on the switch, independent of mirroring, is not supported. (512880)
- Mirror on drop does not mirror VXLAN IPv6 underlay traffic drops. (516212)
- Issuing "reboot -f" from the shell can cause continuous machine check errors, requiring a power cycle.

Instead, use the "reload" CLI command, or "reboot" in bash. (525314)

SKUs Affected: DCS-7010TX-48, DCS-7010TX-48-DC, DCS-7010TX-48-DC-F, DCS-7010TX-48-DC-R, DCS-7010TX-48-F and DCS-7010TX-48-R

- The mirroring logic can only capture and mirror packets seen by the ingress pipeline. Packets generated by the memory-management unit do not traverse the ingress pipeline and hence cannot be mirrored. This applies to all types of mirroring. (527668)
- When policer is attached to the class-default of an egress policy-map consisting of both IPV4 and IPV6 ACLs, IPV4 and IPV6 traffic hitting class-default would be policed individually with rate configured on default class-map (541843)
- L3 MTU configuration is not honored on a NAT enabled interface for multicast traffic. (546193)
SKUs Affected: CCS-720XP-24Y6, CCS-720XP-24ZY4, CCS-720XP-48Y6, CCS-720XP-48ZC2, DCS-7050CX3-32S, DCS-7050CX3M-32S, DCS-7050SX3-48C8, DCS-7050SX3-48YC12, DCS-7050SX3-48YC8 and DCS-7050TX3-48C8
- The Strata agent restarts after running "platform trident diag psr <port>" command. (546588)
- Routing between VXLAN VLANs with VRRP is not supported. (548160)
Series Affected: DCS-7050QX, DCS-7050SX, DCS-7050TX, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series
- The traffic load across ports comprising a port channel(s) may be uneven for certain traffic patterns and/or port channel configuration and state.

Workaround:

- change the seed in the hashing algorithm until satisfactory result is

achieved by executing configuration CLI command:

```
"arista(config)# port-channel load-balance trident <seed>" (555833)
```

- Implicit ICMPv6 permit rule(s) in TCAM is not seen in output of "show platform trident tcam acl detail" command.

Workaround is to check the entry and packet hits for the entry using output from "show platform trident tcam detail". This command shows all ACL entries in TCAM and packet hits per entry. (563228)

- Enabling L2 protocol forwarding of LLDP on a given port will implicitly enable forwarding of EAPOL and MKA packets on the same. (569478)
- Configure replace operation is not supported with TCAM profile configuration. Changing TCAM profile with configure-replace may cause an unexpected restart of forwarding agent(s) if TCAM based features like ACL, policy QoS are also part of the configuration change. (571912)

Series Affected: DCS-7050CX3 and DCS-7050SX3 Series

- Deny action is not supported on ACLs configured for filtering based packet sampling used for mirroring packets to the collector. (582014)
- BGP packets are ignored by IP counters if hardware IP counters are not enabled. (599275)
- When Port Security configuration is removed from an interface while StrataL2 agent is initializing , we can end up in a state where some MAC addresses on the interface do not get aged out and can also not be cleared via configuration.
A workaround is to flap the interface to remove the MAC addresses. (635439)
- Packets egress mirrored to a Greenspan destination with subinterface nexthops may have the internal VLAN in the inner VLAN portion of the packet. (642351)
- The 32-bit image of EOS does not support the 7050X4, 7358X4, 7060X5, and 7388X5 series switches. Use the 64-bit version instead (646587)

- When MACsec encryption traffic spans across multiple queues encrypted packets may be sent out of order. (669251)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- The outPktCtrl MACsec counters will not increment. (671159)

Series Affected: CCS-750 Series

- When the packets matching FBPS rule are also marked for drop due to other ACL rules, the packets are not mirrored and dropped in hardware. Workaround is to remove the configuration that results in dropping these packets.

(673800)

SKUs Affected: CCS-720XP-48ZC2-F

- Mirroring packets to CPU can cause congestion on CPU queues and disrupt control plane traffic. It is not recommended to leave this feature on for extended periods of time. (675773)
- When attempting to VLAN mirror packets ingressing a VXLAN edge VLAN, the VXLAN encapsulated packets that ingress the core interface and are destined to be decapsulated to the edge VLAN will not be mirrored. (679353)
Series Affected: DCS-7050QX, DCS-7050QX2, DCS-7050SX, DCS-7050SX2, DCS-7050TX, DCS-7050TX2, DCS-7260CX3 and DCS-7260QX Series
- The maximum amount of VLAN sources allowed in a VLAN mirror session depends on the configuration's TCAM usage and the number of VLAN mirror sessions configured. A maximum of up to 3 mirror sessions are supported. Using a configuration that utilizes more VLAN sources than supported will result in VLAN mirroring sessions not being programmed. (680940)
Series Affected: CCS-710P, CCS-720XP, DCS-7010TX, DCS-7050CX3, DCS-7050SX3 and DCS-7050TX3 Series
- DirectFlow redirecting to interface with 'switchport source-interface tx' configured is unsupported. (681960)
- The outPktsUntagged counter will increment for unencrypted traffic egressing an interface even when the interface is not MACsec enabled. (682111)
Series Affected: CCS-750 Series
- The inPktsBadTag MACsec counter will not increment. Instead the drops will be accounted under the inPktsNotValid MACsec counter. (695748)
Series Affected: CCS-750 Series
- A chassis may only contain all 7300X3 or all 7300X linecards. Combining 7300X3 or 7300X linecards may result in repeated restarts of the forwarding agents, and links may stay down. Power cycle the system after removing the incompatible cards to recover. (699880)
Series Affected: DCS-7300X and DCS-7300X3 Series
- If "dst-ip" key-field is removed from "acl vlan ip ingress" feature in TCAM profile then, implicit link-local-multicast TCAM entry will not be programmed for ACL applied to SVI in unshared mode.

Workaround:

To overcome this issue add "dst-ip" key-field to "acl vlan ip ingress" feature in TCAM profile. (704466)

Series Affected: 7300X3, 7368, CCS-710P, CCS-720XP, CCS-750, DCS-7050CX3, DCS-7050SX3, DCS-7060PX4 and DCS-7260CX3 Series

- The system will not boot if the console port is connected to an ethernet switch. To recover the system, remove the cable from the console port and power cycle. (728859)

Series Affected: DCS-7010TX Series

- The CLI "show hardware resource IpostLagLagMember" displays current usage statistics of hardware resources.
That hardware table may appear partially used even after some/all LAG(s) that were active before have been deleted. This is expected behaviour with no other adverse effects. (729756)
- In PTP boundary or generalized PTP mode, configuring the `ptp forward-unicast` feature with interfaces using IPv6 transport is not supported. A workaround would be to use the IPv4 or L2 transport. (730817)
- In cases of failed autonegotiated link up, link training may report as "off" on the primary lane. This is cosmetic. (731948)
- Configuring PTP MACsec bypass using the per-profile "ptp bypass" command is not supported.

Enable the "ptp bypass" command in MAC security configuration mode to have PTP packets bypass MACsec processing. (755578)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8, CCS-722XPM-48ZY8-F, CCS-755-CH and CCS-758-CH

- Only 100G copper downlink and AOC uplinks are supported on this Reverse Airflow platform. (758569)

SKUs Affected: DCS-7060DX5-64S-R

- Per interface L3 MTU CLI configuration is not supported.

Workaround is to use global L3 MTU configuration in "interface defaults" mode. (777699)

Series Affected: CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-722XPM, CCS-750 and DCS-7010TX Series

- If the same ACL is applied on control plane and data plane, the byte counter for the ACL will be incremented for only data plane packet hits while the packet counter will be incremented for both data plane and control plane packet hits.

A workaround is to create different ACLs to apply on control plane and data plane. (788296)

- Mirroring to CPU dataplane traffic using a non-default traffic-class will use the CPU queue of that traffic class, instead of the "Other" CPU queue.
Workaround is to mirror to port instead of CPU destination. (791599)

Series Affected: 7368, DCS-7060X, DCS-7060X2 and DCS-7260X Series

- %ENVMON-3-TEMPSENSORINVALID and %ENVMON-3-TEMPSENSOROK could be reported after overheating of Linecard of Fabric card. The temperature reported by temperature sensors from disabled cards may also report low temperature in "show system environment temperature" CLI command output. (793777)
SKUs Affected: 7320X-32C-LC
- This supervisor module is only supported with 7300X3 linecards. (826371)
SKUs Affected: 7300-SUP2-D
- The 64-bit image of EOS does not support the 7010T series of switches. Use the 32-bit version instead (846011)

CCS-750 Series

- Mirroring PTP packets when PTP is enabled is not supported. (738080)
- PTP two-step E2E and P2P transparent clock mode is not supported. (738097)
- Dropping PTP management message using the per-interface `ptp management drop` configuration is only supported in boundary clock mode. (752342)
- When PTP is configured, PTPv1 packet forwarding is not supported. (752367)

7300X3, 7320X, DCS-7060X and DCS-7260X Series

- When in cut-through forwarding mode, mirror sessions with both ingress mirror sources and egress mirror sources configured to mirror packets to the same mirror destination port will not be programmed in hardware.

The workaround for this is to switch to store-and-forward mode where the limitation is not present. (154721)

- Cut-through mode is not supported on the SFP+ interfaces. (182728)
- Cut-through mode is not supported at 1G speed. (219427)
- Mirror To GRE destinations with multiple ECMP nexthops or LAG is not supported. A single link of the LAG or ECMP nexthops will receive all mirror traffic. (639980)
Series Affected: 7368 and 7388 Series
SKUs Affected: DCS-7060DX4-32 and DCS-7060PX4-32
- The 64-bit image of EOS does not support 7260QX series switches. Use the 32-bit version instead. (723646)
Series Affected: DCS-7260QX Series

DCS-7260X3 Series

- Configuring 100G, 50G, 40G, 25G and 10G simultaneously might result in unexpected packet forwarding behavior. (235075)

7368 and DCS-7060X4 Series

- Multicast replication resources in TH3 are only 10% of their unicast counterparts. Due to this multicast replication in TH3 is limited to 1.2 Tbps. (297317)
- Layer2 source MAC learning may occur on packets received with CRC errors. (325507)
- When 'hardware counter feature gre tunnel interface in' feature is enabled, the gre tunnel interface decap counters are not incremented when subinterface is used as the ingress interface. (464972)
- Per vlan routing and sub interface are not supported together. (486439)
- 3m QSFP-DD passive copper cables are not supported in ports 1-4 and 29-32 (639613)
SKUs Affected: DCS-7060DX4-32
- Double Dot1Q tagged packets with IP payload may not be classified as IP packets in hardware. (813027)

CCS-710P, CCS-720DP and CCS-720DT Series

- Triple-wide TCAMs are not supported. (603246)
Series Affected: CCS-710P Series
- Generation of PFC frames is not supported. (699868)

CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series

- XPN cipher suites are not supported. (590170)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- "show platform trident counters" output is misleading for certain counters with MACsec encryption enabled. The workaround is to use the "show interfaces counters discards" and "show mac security counters interface" commands where possible. (623564)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

CCS-720DF and CCS-720XP Series

- When flow group ACL rules are changed, flow group counters displayed in "show platform trident flow counters" will continue to have counters incremented for active flows corresponding to older flow groups till the flows are aged out. (552450)
- On interfaces with speeds less than 10G, PFC cannot be transmitted. (689087)

CCS-750 Series

- The switch is not capable of generating pause flowcontrol frames. (480688)
Series Affected: CCS-750 Series
- Per interface L3 MTU configuration is not supported.
The workaround is to use the global L3 MTU configuration command instead. (520825)
- The supervisor uplink SCD does not support hitless FPGA upgrade. (528417)
SKUs Affected: CCS-750-SUP100 and CCS-750-SUP25
- The switch is not capable of generating Priority Flow Control (PFC) frames. (555294)
Series Affected: CCS-750 Series
- Shaper below 1 Mbps is not supported. (605407)
- MACsec encrypted flow control packets and priority-flow-control packets that are received will not be visible in "show flowcontrol" or "show priority-flow-control counters". (670539)
- When a mirror session has a LAG member as the source, packets sent to the LAG will not be mirrored. (712186)
- TCAM profile does not effect any saving of TCAM slices / banks for Egress MAC ACL and IP ACL applied to Ethernet, LAG ports (723449)
SKUs Affected: CCS-755-X3-SC and CCS-758-X3-SC
- Ports on this linecard share a total sustainable bandwidth.

Though this linecard can absorb concurrent bursts received on all ports, the total sustained bandwidth is limited to a total of 148 Gbps for an average frame size of at least 276 Bytes.

To optimize sustained bandwidth availability, the recommendation is to distribute the bandwidth equally across the groups of eight ports. The groups are: Ethernet1-8, Ethernet9-16, Ethernet17-24, Ethernet25-32, Ethernet33-40, Ethernet41-48

For example, if there are 24x5G and 24x2.5G devices to connect, it is recommended that 4x5G and 4x2.5G be connected to each of the 6 groups of front-panel ports. (788129)

SKUs Affected: CCS-750X-48ZXP-LC

- Stateful Switchover triggered by removing the active supervisor will cause a long data-plane outage for uplink traffic destined to the removed uplinks. Workaround is to unplug the uplink interfaces, wait for the link to go down

and then remove the active supervisor. (839923)

Series Affected: CCS-750 Series

- If a device is behind a switch or hub and is authenticated via Dot1x MAC Based Authentication it will not be able to authenticate via Dot1x MBA on a different interface.

A workaround is to configure 'dot1x timeout idle-host' (839960)

7300X3, AS7326, AS7726 and DCS-7050X3 Series

- The egress queue and per hop latency of the INT terminator node are always shown as zero. (562000)

Series Affected: DCS-7050CX3 and DCS-7050SX3 Series

7358 and DCS-7050X4 Series

- Sampled counters are not required. (437955)
- Color based DSCP or CoS assignment is not supported. (679840)
- Packets destined for the control plane may appear to be dropped within the ingress pipeline. This will appear in the output of `show platform trident counters` as a RX drop, and is expected due to more drop events available in the pipeline. (685445)
- On interfaces Ethernet33 and Ethernet34, PFC cannot be transmitted. (737611)
SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4-32S-F and DCS-7050PX4-32S-F
- Egress ACL does not match on policy rewritten DSCP nor remarked ECN values. (741424)
Series Affected: DCS-7050X4 Series
- Maximum replications for a non unicast packet is limited to 1023. (750699)
Series Affected: DCS-7050X4 Series

Transceiver Series

- SFP-1G-T transceivers can only advertise one speed at a time. Desired speed and negotiation must be explicitly configured. Recommended commands are:

* "speed auto" / "speed auto 1gfull" / "speed auto 1000full" : enable autoneg, advertise only 1G
* "speed auto 100full" : enable autoneg, advertise only 100M
* "speed 100full" : disable autoneg, force speed to 100M (159401)
SKUs Affected: 1000BASE-T
- 25G transceivers could suffer signal degradation when using the MAM1Q00A-QSA QSFP-to-SFP adapter.

Use a MAM1Q00A-QSA28 QSFP28-to-SFP28 adapter instead. (253324)

- During reboot, a peer may see a link-flap on ports using 1000BASE-T transceivers.
The workaround is to configure the link-debounce period to 30 seconds for the link-up transition on the peer interfaces. (268080)
SKUs Affected: 1000BASE-T
- For SKUs CAB-O-2Q-400G, CAB-O-4Q-400G, CAB-D-2Q-400G and CAB-D-4Q-400G, when operating at 25G or 50G NRZ speeds, forward error-correction should be manually configured (enabled or disabled) on both sides of the link as the default error-correction on both ends may not match. (392124)
- SFP-10G-T transceivers can only advertise one speed at a time. Desired speed and negotiation must be explicitly configured. Recommended commands are:
 - * "speed auto" / "speed auto 10gfull"/"speed auto 10000full" : enable autoneg, advertise only 10G
 - * "speed auto 1gfull" / "speed auto 1000full" : enable autoneg, advertise only 1G
 - * "speed auto 100full" : enable autoneg, advertise only 100M
 - * "speed 100full" : disable autoneg, force speed to 100M (413941)SKUs Affected: SFP-10G-T
- Revision 20 OSFP-400G-2FR4 transceivers may not correctly report the operational value of the configured Tx Input Equalization. This can be observed in 'show interfaces transceiver tuning detail', in which the Operational Tx Input Equalization may not match the Configured value. If the 'Tx Input Equalization' Configured column shows auto, automatic Tx Input Equalization is enabled, despite the Operational value shown. (423045)
SKUs Affected: OSFP-400G-2FR4
- Using a default speed configuration on any interface associated with a 400G optical transceiver may prevent all interfaces associated with the transceiver from linking up.
The workaround is to explicitly configure speed on all interfaces associated with the 400G optical transceiver. (426232)
- FEC should be configured manually when using OSFP/QSFP-DD to QSFP200/SFP50 breakout cables at 25G or 50G-2 speeds (432776)
SKUs Affected: CAB-D-2Q-400G, CAB-D-4Q-400G, CAB-D-8S-400G, CAB-O-2Q-400G, CAB-O-4Q-400G and CAB-O-8S-400G
- There is a hardware incompatibility between 100G AOCs with serial numbers starting with "AEB" and certain Mellanox adapters.

The recommendation is to RMA for AOCs with serial numbers that don't start with "AEB". (457261)

- Interfaces with BCM82391 PHY do not support 1000BASE-T transceivers. (535121)
- QSFP200 cables and optics encoded using the CMIS specification are not recognized in QSFP100 ports. (536318)
- OSFP & QSFP-DD port LEDs are not supported in hardware LED CLI config mode (650615)
- Interfaces with the SFP-10G-MRA-T 10GBASE-T module may link up during system reload if the peer is enabled. During the reload the port will not pass traffic. This will occur on enabled or disabled ports. (651386)
- SFP 10GBASE-T transceivers which do not support Soft RX_LOS or Soft TX_DISABLE are not supported on DSFP ports. (702025)
SKUs Affected: SFP-10G-T
- The rate adapting transceivers SFP-10G-MRA-T, SFP-10G-RA-1G-LX, and SFP-10G-RA-1G-SX will not allow PAUSE or PFC frames to pass to and from the peer device. (881411)
SKUs Affected: SFP-10G-MRA-T
- 400GBASE-ZR transceivers do not support Zero Touch Provisioning. (882209)
SKUs Affected: OSFP-400G-ZR, OSFP-400G-ZRP, QDD-400G-ZR and QDD-400G-ZRP
- Alert state and count of temperature sensor in "show system environment temperature transceiver detail" CLI output is never updated (886972)
- The SFP-1G-T pluggable 1000BASE-T SFP is not capable of advertising PAUSE or ASM PAUSE during auto negotiation. This does not impact the ability to send or receive PAUSE per the specific switch's abilities and flow control configuration. (904116)
SKUs Affected: 1000BASE-T

7300X3, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX and DCS-7050X3 Series

- The DirectFlow "set vlan action" command is not supported. (237046)
- When configuring a DirectFlow action, specifying "interface hardware-loopback" as the destination of "action output" is not supported. (243031)
- IPv6 packets with WESP payload (next header 141) bypass the IPv6 security ACLs and service policies. (243387)
- Directflow configuration with "set vlan", but no "action set output interface" configured is not supported. (345852)
- MAC Egress ACL configurations matching IP packets may not work when the IPv4/v6 per-interface counter feature is also configured. (411243)

- Leaf stateful switch upgrade is not supported on this platform. (415340)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- CLI output of "show platform trident l3 shadow mpls-vc-and-swap-label" will show push as MplsLabelAction even for swap operation. (416134)
- Packets dropped by the MACsec engine may not be reported in the CLI (420085)
SKUs Affected: DCS-7050CX3M-32S
- MPLS packet undergoing swap operation does not honor the MTU configured on ingress/egress interface. (422668)
- Only 2K MPLS routes are supported (425691)
Series Affected: CCS-720XP Series
- If the packet ingresses on one line card and egresses on a different line card, then a truncated egress mirrored packets will be 4 bytes shorter than expected. (445200)
SKUs Affected: DCS-7304 and DCS-7308
- When IPV6 VXLAN underlay core interface is an access port on a SVI, VXLAN encapsulated packets may be sent out with VLAN tag in the outer ethernet header. The work around is to configure the port as a trunk port. (533929)
- Packets mirrored to GRE with a size within 16 bytes of the MTU limit may be dropped on egress due to exceeding the MTU size after the mirrored packet is encapsulated. (592957)
- The UDP header checksum is not set correctly for PTP packets after NAT (601036)
- "Dont fragment" bit is not set in the outer IP header of the VXLAN encapsulated packets when the payload is non-IP. (605003)
- When Postcard Telemetry is enabled, EVPN Multicast will not be supported. (614072)
Series Affected: 7300X3, CCS-720XP, DCS-7050SX3, DCS-7050TX3 and DCS-7050X3 Series
- ACL resource sharing feature is not supported for NAT translations of tunneled packets like VXLAN-tunneled packets.

Workaround is to turn off ACL resource sharing: "no hardware nat static access-list resource sharing" (620222)
- If the destination MAC address of a NAT flow is not learnt on an interface, then the packet will be dropped. Flooding the translated packet on all interfaces belonging to the VLAN is not supported.

Workaround is to adjust the ARP and MAC address timeouts, so that MAC addresses is learnt again. (642094)

- DirectFlow actions on interfaces configured as dot1q-tunnel are not supported. (686339)
- If an egress ACL is filtering on the NAT flow's source IP, then the ACL rules have to permit both original source IP and translated IP to avoid NAT flows getting dropped. (688883)
- With EVPN Multicast with underlay as HER configured, known multicast traffic might not be load balanced across all ECMP links. (709053)
- Ingress mirroring double tagged packets to CPU will have the outer VLAN tag stripped.
Workaround is to use a port mirroring destination instead. (722891)
- Packets with Broadcast source MAC may be copied to CPU instead of being dropped. (745001)
- When IGMP snooping is enabled a single unknown multicast floodset is shared across all VLANs with IGMP snooping enabled. This can cause unexpected flooding to a trunk port that has no multicast routers on a given VLAN, but has multicast router attached to the same trunk port on another VLAN. (822757)
- When ingress sFlow is enabled, DHCP relay functionality may not work when the DHCP packets to be relayed are selected for sampling. (825322)
- Packets that match on a dynamic NAT ACL will be sent to CPU if the packets are egressing out of any other interface that also has NAT rules configured. The workaround is to make the NAT ACL less greedy by either reducing the scope of matching IPs or by adding deny rules for specific destination IPs that should not be NATed. (900348)

7358 and DCS-7050X4 Series

- Greenspan packets where the size of the mirror payload plus the Greenspan header exceeds the configured MTU size of the ingress interface will be fragmented on a second pass of the pipeline.

Note that high rates of fragmentation can result in high CPU utilization. (791223)

- By default IPv6 packets are set to not fragment. If an added IPv6 Greenspan header causes a packet to exceed the MTU size of the destination next hop, then it will be dropped. Workaround is to increase MTU size of the nexthop port. (818665)

- Early shipments of system bundles with this switch card were loaded with the incorrect EOS-4.29.0.2F binary (32-bit) instead of the EOS64-4.29.0.2F binary (64-bit). Unexpected traffic drops may be seen if the 32-bit binary is used. Upgrade to a 64-bit version of EOS to avoid such traffic drops. (909191)

SKUs Affected: 7358X4-SC

Conditions and Impacts

The release notes listed above use shorthand terminology to refer to conditions that arise frequently in connection with bugs. This section explains each condition and its impact in more detail.

Condition: Rib agent restart

Impact: When the Rib agent goes down, all routing sessions are terminated; all BGP sessions drop, all OSPF adjacencies are lost, and neighbors stop forwarding traffic to us. When the Rib agent restarts, adjacencies are re-formed, and, after protocol convergence, traffic flows through us again. In most cases where there is adequate network-level redundancy, application-visible impact should be minimal

Condition: Rib agent hang

Impact: When the Rib agent hangs, routing protocols stop running. Routing peers will generally time out their session with the hung Rib agent, withdrawing routes accordingly. Meanwhile, the device continues to forward packets based on existing routing state. During this time, the device will not respond to routing topology changes. After a heartbeat timer expires (typically 10 minutes), the Rib agent restarts. In networks with sufficient redundancy, application impact of a Rib agent hang is usually low, because there is no data path disruption. However, in conjunction with other network changes (such as link failure or introduction), routing loops may occur.

Condition: kernel crash

Impact: A kernel crash has impact similar to issuing the "reload now" command. All links go down and all forwarding stops. Then, the system reloads, which may take up to 15 minutes. In a network with adequate redundancy, there should be minimal traffic loss associated with this period. After reload, links come up and protocols (LACP, STP, BGP, etc) reconverge. Protocol reconvergence is not necessarily hitless, depending on topology and configuration. For example, a switch may advertise a prefix to a connected subnet before STP has converged. However, any associated outage should be short lived, typically lasting around 30 seconds. The MLAG-SSO feature can dramatically reduce the window of disruption.

Condition: forwarding agent restart

Impact: A forwarding agent restart typically results in the switch ASIC being reset. It clears packet memory, dropping all packets currently in the switch, and causes all links to go down. The restart can take up to 45 seconds, during which time all links are down. Once the restart completes, all links come back up automatically, followed by protocol reconvergence (MLAG, LACP, STP, BGP/OSPF/IS-

IS, etc). Depending on which protocols and options are in use, the reconvergence may take a few seconds up through several minutes. On modular switches, the impact of forwarding agent restart is limited to the affected line card; that is, if the forwarding agent for line card 3 restarts, then all ports on line card 3 bounce and protocols on those ports reconverge, but ports on other line cards are unaffected.