Arista Networks EOS 4.28.0.2F Release Notes

Version: 1.0

12 October 2022

Table of Contents

Rib

MetaWatch

MLAG

MPLS

Multicast

NAT

SwitchApp

vEOS Router / CloudEOS

DCS-7170 Series

DCS-7170 Series

CCS-710P, CCS-720XP and CCS-750 Series

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

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

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

DCS-7280R and DCS-7500R Series

DCS-7280R2 and DCS-7500R2 Series

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

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

DCS-7800R3 Series

DCS-7020 Series

DCS-7280R3 Series

DCS-7500R3 and DCS-7800R3 Series

DCS-7020R, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

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

DCS-7280R and DCS-7280R2 Series

7300X3, 7320X, 7368, 7388, CCS-710P, CCS-720XP, CCS-750, DCS-7010, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060, DCS-7060X, DCS-7060X2,

DCS-7060X4, DCS-7250X, DCS-7260X, DCS-7260X3, DCS-7300X and DCS-7300X3 Series

DCS-7010 Series

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

DCS-7260X3 Series

7368 and DCS-7060X4 Series

7388 Series

DCS-7050X, DCS-7250X and DCS-7300X Series

DCS-7050X2 Series

CCS-710P Series

CCS-720XP and DCS-7010TX Series

CCS-720XP Series

CCS-750 Series

DCS-7050X3 Series

Transceiver Series

7300X3, CCS-710P, CCS-720XP, CCS-750, DCS-7010TX and DCS-7050X3 Series

Limitations and Restrictions in 4.28.0.2F

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

vEOS Router / CloudEOS

CVP

Telemetry

CVX

General

IPSEC

Layer 1

Layer 2

Layer 3

- Multi-agent
- Rib

MLAG

MPLS

Multicast

SwitchApp

vEOS Router / CloudEOS

DCS-7170 Series

- DCS-7170 Series

CCS-710P, CCS-720XP and CCS-750 Series

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

DCS-7130 Series

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

CCS-710P, CCS-720XP and CCS-750 Series

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

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

DCS-7280R and DCS-7280R2 Series

7300X3, 7320X, 7368, 7388, CCS-710P, CCS-720XP, CCS-750, DCS-7010, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060, DCS-7060X, DCS-7060X2,

DCS-7060X4, DCS-7250X, DCS-7260X, DCS-7260X3, DCS-7300X and DCS-7300X3 Series

DCS-7010 Series

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

DCS-7260X3 Series

7368 and DCS-7060X4 Series

DCS-7050X, DCS-7250X and DCS-7300X Series

DCS-7050X2 Series

CCS-710P Series

CCS-720XP and DCS-7010TX Series

CCS-720XP Series

CCS-750 Series

7300X3 and DCS-7050X3 Series

Transceiver Series

7300X3, CCS-710P, CCS-720XP, CCS-750, DCS-7010TX and DCS-7050X3 Series

Conditions and Impacts

EOS 4.28.0.2F 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
- IGMP-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
- 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

- CCS-710P-12
- CCS-710P-16P
- CCS-720XP-24Y6
- CCS-720XP-24Y6-F
- CCS-720XP-24Y6-R
- CCS-720XP-24ZY4
- CCS-720XP-24ZY4-F
- CCS-720XP-48Y6
- CCS-720XP-48Y6-F
- CCS-720XP-48Y6-R
- CCS-720XP-48ZC2
- CCS-720XP-48ZC2-F
- CCS-720XP-96ZC2
- CCS-720XP-96ZC2-F
- CCS-722XPM-48Y4
- CCS-722XPM-48Y4-F
- CCS-722XPM-48ZY8
- CCS-722XPM-48ZY8-F
- DCS-7010T-48
- DCS-7010T-48-DC
- DCS-7010T-48-DC-F
- DCS-7010T-48-DC-R
- DCS-7010T-48-F
- DCS-7010T-48-R
- DCS-7010TX-48
- DCS-7010TX-48-DC
- DCS-7010TX-48-DC-F
- DCS-7010TX-48-DC-R
- DCS-7010TX-48-F
- DCS-7010TX-48-R
- DCS-7020SR-24C2
- DCS-7020SR-24C2-F
- DCS-7020SR-24C2-R
- DCS-7020SR-32C2
- DCS-7020SR-32C2-F
- DCS-7020SR-32C2-R
- DCS-7020SRG-24C2
- DCS-7020SRG-24C2-F
- DCS-7130-48LA-F (HwRev:A3)
- DCS-7130-48LA-R (HwRev:A2)
- DCS-7130-48LA-R (HwRev:A3)
- DCS-7130-48LAS (HwRev:B2)
- DCS-7130-48LAS (HwRev:B3)
- 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 (HwRev:A2)
- DCS-7130-48LB (HwRev:A3)
- 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 (HwRev:A2)
- DCS-7130-48LBA (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 (HwRev:B2)
- DCS-7130-48LBAS (HwRev:B3)
- 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 (HwRev:B2)
- DCS-7130-48LBS (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 (HwRev:B2)
- DCS-7130-48LS (HwRev:B3)
- DCS-7130-48LS-F (HwRev:B2)
- DCS-7130-48LS-F (HwRev:B3)
- DCS-7130-48LS-R (HwRev:B2)
- DCS-7280CR3MK-32P4-F
- DCS-7280CR3MK-32P4-R
- DCS-7280CR3MK-32P4S
- DCS-7280CR3MK-32P4S-F
- DCS-7280CR3MK-32P4S-R
- DCS-7280DR3-24
- DCS-7280DR3-24-F
- DCS-7280DR3-24-M-F
- DCS-7280DR3K-24
- DCS-7280DR3K-24-F
- DCS-7280PR3-24
- DCS-7280PR3-24-F
- DCS-7280PR3-24-M-F
- DCS-7280PR3K-24
- DCS-7280PR3K-24-F
- DCS-7280QR-C36
- DCS-7280QR-C36-F
- DCS-7280QR-C36-M-F
- DCS-7280QR-C36-M-R
- DCS-7280QR-C36-R
- DCS-7280QR-C72
- DCS-7280QR-C72-F
- DCS-7280QR-C72-M-F
- DCS-7280QR-C72-M-R
- DCS-7280QR-C72-R
- DCS-7280QRA-C36S
- DCS-7280QRA-C36S-F
- DCS-7280QRA-C36S-M-F
- DCS-7280QRA-C36S-M-R
- DCS-7280QRA-C36S-R
- DCS-7280SR-48C6
- DCS-7280SR-48C6-F
- DCS-7280SR-48C6-M-F
- DCS-7280SR-48C6-M-R
- DCS-7280SR-48C6-R
- DCS-7280SR2-48YC6
- DCS-7280SR2-48YC6-F
- DCS-7280SR2-48YC6-M-F

- DCS-7020SRG-24C2-R
- DCS-7020TR-48
- DCS-7020TR-48-F
- DCS-7020TR-48-R
- DCS-7020TRA-48
- DCS-7020TRA-48-F
- DCS-7020TRA-48-R
- DCS-7050CX3-32S
- DCS-7050CX3-32S-F
- DCS-7050CX3-32S-R
- DCS-7050CX3-32S-SSD-F
- DCS-7050CX3-32S-SSD-R
- DCS-7050CX3M-32S
- DCS-7050CX3M-32S-F
- DCS-7050CX3M-32S-R
- DCS-7050QX-32S
- DCS-7050QX-32S-F
- DCS-7050QX-32S-R
- DCS-7050QX-32S-SSD-F
- DCS-7050QX-32S-SSD-R
- DCS-7050QX2-32S
- DCS-7050QX2-32S-F
- DCS-7050QX2-32S-R
- DCS-7050SX-128
- DCS-7050SX-128-F
- DCS-7050SX-128-R
- DCS-7050SX-128-SSD-F
- DCS-7050SX-128-SSD-R
- DCS-7050SX-64
- DCS-7050SX-64-F
- DCS-7050SX-64-R
- DCS-7050SX-64-SSD-F
- DCS-7050SX-64-SSD-R
- DCS-7050SX-72
- DCS-7050SX-72-F
- DCS-7050SX-72-R
- DCS-7050SX-72-SSD-F
- DCS-7050SX-72-SSD-R
- DCS-7050SX-72Q
- DCS-7050SX-72Q-F
- DCS-7050SX-72Q-R
- DCS-7050SX-96
- DCS-7050SX-96-F
- DCS-7050SX-96-R
- DCS-7050SX-96-SSD-F
- DCS-7130-48LS-R (HwRev:B3)
- DCS-7130-96 (HwRev:A7)
- DCS-7130-96 (HwRev:A8)
- DCS-7130-96-F (HwRev:A7)
- DCS-7130-96-F (HwRev:A8)
- DCS-7130-96-R (HwRev:A7)
- DCS-7130-96-R (HwRev:A8)
- DCS-7130-96E (HwRev:A7)
- DCS-7130-96E-F (HwRev:A7)
- DCS-7130-96E-R (HwRev:A7)
- DCS-7130-96L (HwRev:A2)
- DCS-7130-96L (HwRev:A3)
- 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 (HwRev:A2)
- DCS-7130-96LA (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 (HwRev:B1)
- DCS-7130-96LAS (HwRev:B2)
- 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 (HwRev:A2)
- DCS-7130-96LB (HwRev:A3)
- 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 (HwRev:A2)
- DCS-7130-96LBA (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 (HwRev:B1)
- DCS-7130-96LBAS-F (HwRev:B1)
- DCS-7130-96LBAS-R (HwRev:B1)
- DCS-7130-96LBS (HwRev:B1)
- DCS-7130-96LBS (HwRev:B2)
- DCS-7280SR2-48YC6-M-R
- DCS-7280SR2-48YC6-R
- DCS-7280SR2A-48YC6
- DCS-7280SR2A-48YC6-F
- DCS-7280SR2A-48YC6-M-F
- DCS-7280SR2A-48YC6-M-R
- DCS-7280SR2A-48YC6-R
- DCS-7280SR2K-48C6
- DCS-7280SR2K-48C6-M-F
- DCS-7280SR2K-48C6-M-R
- DCS-7280SR3-40YC6
- DCS-7280SR3-40YC6-F
- DCS-7280SR3-40YC6-R
- DCS-7280SR3-48YC8
- DCS-7280SR3-48YC8-F
- DCS-7280SR3-48YC8-R
- DCS-7280SR3K-48YC8
- DCS-7280SR3K-48YC8-F
- DCS-7280SR3K-48YC8-R
- DCS-7280SR3K-48YC8A
- DCS-7280SR3K-48YC8A-F
- DCS-7280SR3K-48YC8A-R
- DCS-7280SR3M-48YC8
- DCS-7280SR3M-48YC8-F
- DCS-7280SR3M-48YC8-R
- DCS-7280SRA-48C6
- DCS-7280SRA-48C6-F
- DCS-7280SRA-48C6-M-F
- DCS-7280SRA-48C6-M-R
- DCS-7280SRA-48C6-R
- DCS-7280SRAM-48C6
- DCS-7280SRAM-48C6-F
- DCS-7280SRAM-48C6-R
- DCS-7280SRM-40CX2
- DCS-7280SRM-40CX2-F
- DCS-7280SRM-40CX2-R
- DCS-7280TR-48C6
- DCS-7280TR-48C6-F
- DCS-7280TR-48C6-M-F
- DCS-7280TR-48C6-M-R
- DCS-7280TR-48C6-R
- DCS-7280TRA-48C6
- DCS-7280TRA-48C6-F
- DCS-7280TRA-48C6-M-F
- DCS-7280TRA-48C6-M-R

- DCS-7050SX-96-SSD-R
- DCS-7050SX2-128
- DCS-7050SX2-128-F
- DCS-7050SX2-128-R
- DCS-7050SX2-72Q
- DCS-7050SX2-72Q-F
- DCS-7050SX2-72Q-R
- DCS-7050SX3-48C8
- DCS-7050SX3-48C8-F
- DCS-7050SX3-48C8-R
- DCS-7050SX3-48YC12
- DCS-7050SX3-48YC12-F
- DCS-7050SX3-48YC8
- DCS-7050SX3-48YC8-F
- DCS-7050SX3-48YC8-R
- DCS-7050SX3-96YC8
- DCS-7050SX3-96YC8-F
- DCS-7050SX3-96YC8-R
- DCS-7050TX-128
- DCS-7050TX-128-F
- DCS-7050TX-128-R
- DCS-7050TX-128-SSD-F
- DCS-7050TX-128-SSD-R
- DCS-7050TX-48
- DCS-7050TX-48-F
- DCS-7050TX-48-R
- DCS-7050TX-48-SSD-F
- DCS-7050TX-48-SSD-R
- DCS-7050TX-64
- DCS-7050TX-64-F
- DCS-7050TX-64-R
- DCS-7050TX-64-SSD-F
- DCS-7050TX-64-SSD-R
- DCS-7050TX-72
- DCS-7050TX-72-F
- DCS-7050TX-72-R
- DCS-7050TX-72-SSD-F
- DCS-7050TX-72-SSD-R
- DCS-7050TX-72Q
- DCS-7050TX-72Q-F
- DCS-7050TX-72Q-R
- DCS-7050TX-96
- DCS-7050TX-96-F
- DCS-7050TX-96-R
- DCS-7050TX-96-SSD-F
- DCS-7050TX-96-SSD-R
- DCS-7050TX2-128
- 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 (HwRev:B1)
- DCS-7130-96LS (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 (HwRev:B1)
- DCS-7130-96S-F (HwRev:B1)
- DCS-7130-96S-R (HwRev:B1)
- DCS-7170-32C
- DCS-7170-32C-F
- DCS-7170-32C-M-F
- DCS-7170-32C-M-R
- DCS-7170-32C-R
- DCS-7170-32CD
- DCS-7170-32CD-F
- DCS-7170-32CD-R
- DCS-7170-64C
- DCS-7170-64C-F
- DCS-7170-64C-M-F
- DCS-7170-64C-M-R
- DCS-7170-64C-R
- DCS-7250QX-64
- DCS-7250QX-64-F
- DCS-7250QX-64-R
- DCS-7250QX-64-SSD-F
- DCS-7250QX-64-SSD-R
- DCS-7260CX-64
- DCS-7260CX-64-F
- DCS-7260CX-64-R
- DCS-7260CX-64-SSD-F
- DCS-7260CX-64-SSD-R
- DCS-7260CX3-64
- DCS-7260CX3-64-F
- DCS-7260CX3-64-R
- DCS-7260CX3-64E
- DCS-7260CX3-64E-F
- DCS-7260CX3-64E-R
- DCS-7260CX3-64LQ
- DCS-7260CX3-64LQ-F
- DCS-7260CX3-64LQ-R
- DCS-7260QX-64
- DCS-7260QX-64-F
- DCS-7280TRA-48C6-R
- FAN-7000
- FAN-7000-F
- FAN-7000-R
- FAN-7000H-R
- FAN-7001C
- FAN-7001C-F
- FAN-7001C-R
- FAN-7001D
- FAN-7001D-F
- FAN-7001DH
- FAN-7001DH-F
- FAN-7002H
- FAN-7002H-R
- FAN-7010
- FAN-7010X
- FAN-7011H
- FAN-7011H-F
- FAN-7011H-R
- FAN-7011M
- FAN-7011M-F
- FAN-7011M-R
- FAN-7012H-BLUE
- FAN-7012H-RED
- FAN-7012M-BLUE
- FAN-7012M-RED
- FAN-7130-1U
- FAN-7130-1U-F
- FAN-7130-1U-R
- FAN-7130-2U
- FAN-7130-2U-F
- FAN-7130-2U-R
- PWR-1011-AC-BLUE
- PWR-1011-AC-RED
- PWR-1011-DC-BLUE
- PWR-1011-DC-RED
- PWR-1021-AC-RED
- PWR-1100AC
- PWR-1100AC-F
- PWR-1100AC-R
- PWR-1511-AC-BLUE
- PWR-1511-AC-RED
- PWR-1511-DC-BLUE
- PWR-1511-DC-RED
- PWR-1600AC
- PWR-1600AC-F
- PWR-1611-AC-RED

- DCS-7050TX2-128-F
- DCS-7050TX2-128-R
- DCS-7050TX3-48C8
- DCS-7050TX3-48C8-F
- DCS-7050TX3-48C8-R
- DCS-7060CX-32S
- DCS-7060CX-32S (HwRev: 13.20)
- 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
- DCS-7060CX2-32S-F
- DCS-7060CX2-32S-R
- DCS-7060DX4-32
- DCS-7060DX4-32-F
- DCS-7060PX4-32
- DCS-7060PX4-32-F
- DCS-7060SX2-48YC6
- DCS-7060SX2-48YC6-F
- DCS-7060SX2-48YC6-R
- DCS-7130-16G3 (HwRev:C1)
- DCS-7130-16G3-F (HwRev:C1)
- DCS-7130-16G3-R (HwRev:C1)
- DCS-7130-16G3S (HwRev:D1)
- DCS-7130-16G3S-F (HwRev:D1)
- DCS-7130-16G3S-R (HwRev:D1)
- DCS-7130-48E (HwRev:A4)
- DCS-7130-48E-F (HwRev:A4)
- DCS-7130-48E-R (HwRev:A4)
- DCS-7130-48EH (HwRev:A4)
- DCS-7130-48EH-F (HwRev:A4)
- DCS-7130-48EH-R (HwRev:A4)
- DCS-7130-48EHS (HwRev:B1)
- DCS-7130-48EHS (HwRev:B2)
- 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 (HwRev:B1)
- DCS-7130-48G3 (HwRev:B2)
- DCS-7130-48G3-F (HwRev:B1)
- DCS-7130-48G3-F (HwRev:B2)
- DCS-7130-48G3-R (HwRev:B1)
- DCS-7260QX-64-R
- DCS-7260QX-64-SSD-F
- DCS-7260QX-64-SSD-R
- DCS-7280CR-48
- DCS-7280CR-48-F
- DCS-7280CR2-60
- DCS-7280CR2-60-F
- DCS-7280CR2A-30
- DCS-7280CR2A-30-F
- DCS-7280CR2A-60
- DCS-7280CR2A-60-F
- DCS-7280CR2K-30
- DCS-7280CR2K-30-F
- DCS-7280CR2K-60
- DCS-7280CR2K-60-F
- DCS-7280CR2M-30
- DCS-7280CR2M-30-F
- 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-7280CR3-36S
- DCS-7280CR3-36S-F
- DCS-7280CR3-96
- DCS-7280CR3-96-F
- DCS-7280CR3K-32D4
- DCS-7280CR3K-32D4-F
- DCS-7280CR3K-32D4-R
- DCS-7280CR3K-32D4A
- DCS-7280CR3K-32D4A-F
- DCS-7280CR3K-32D4A-R
- DCS-7280CR3K-32P4
- DCS-7280CR3K-32P4-F
- DCS-7280CR3K-32P4-R
- DCS-7280CR3K-32P4A
- DCS-7280CR3K-32P4A-F
- DCS-7280CR3K-32P4A-R
- DCS-7280CR3K-36A
- PWR-1611-DC-RED
- PWR-1900-DC
- PWR-1900-DC-F
- PWR-1900-DC-R
- PWR-1900AC
- PWR-1900AC-F
- PWR-1900AC-R
- PWR-2021-AC-RED
- PWR-2411-AC-RED
- PWR-2411-DC-RED
- PWR-3001-AC-BLUE
- PWR-3001-AC-RED
- PWR-3001-DC-BLUE
- PWR-3001-DC-RED
- PWR-400-DC-BLUE
- PWR-400-DC-RED
- PWR-400AC-BLUE
- PWR-400AC-RED
- PWR-460AC
- PWR-460AC-F
- PWR-460AC-R
- PWR-460DC
- PWR-460DC-F
- PWR-460DC-R
- PWR-500-DC
- PWR-500-DC-F
- PWR-500-DC-R
- PWR-500AC
- PWR-500AC-F
- PWR-500AC-R
- PWR-501AC
- PWR-501AC-F
- PWR-501AC-R
- PWR-511-AC-BLUE
- PWR-511-AC-RED
- PWR-511-DC-BLUE
- PWR-511-DC-RED
- PWR-621-AC-BLUE
- PWR-621-AC-RED
- PWR-721-DC-RED
- PWR-745AC
- PWR-745AC-F
- PWR-745AC-R
- PWR-747AC

- DCS-7130-48G3-R (HwRev:B2)
- DCS-7130-48G3S (HwRev:C1)
- DCS-7130-48G3S-F (HwRev:C1)
- DCS-7130-48G3S-R (HwRev:C1)
- DCS-7130-48L (HwRev:A2)
- DCS-7130-48L (HwRev:A3)
- 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 (HwRev:A2)
- DCS-7130-48LA (HwRev:A3)
- DCS-7130-48LA-F (HwRev:A2)
- DCS-7280CR3K-36A-F
- DCS-7280CR3K-36S
- DCS-7280CR3K-36S-F
- DCS-7280CR3K-96
- DCS-7280CR3K-96-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
- PWR-747AC-F
- PWR-747AC-R
- PWR-750AC
- PWR-750AC-F
- PWR-750AC-R
- FAN-7000H
- FAN-7000H-F
- FAN-7002
- FAN-7002-F
- FAN-7002-R
- FAN-7002H-F
- FAN-7012MP-RED

Modular

- FAN-7000H
- FAN-7000H-F
- FAN-7002
- FAN-7002-F
- FAN-7002-R
- FAN-7002H-F
- FAN-7012MP-RED
- 7300-SUP
- 7300-SUP-D
- 7300X-32Q-LC
- 7300X-64S-LC
- 7300X-64T-LC
- 7300X3-32C-LC
- 7300X3-48YC4-LC
- 7304X-FM
- 7304X3-FM
- 7304X3-FM2
- 7308X-FM
- 7308X3-FM
- 7308X3-FM2
- 7320X-32C-LC
- 7324X-FM
- 7328X-FM
- 7358-16C
- 7368
- 7368-16C
- 7368-16S
- 7368-4D
- 7368-4P
- 7368-SUP
- 7368-SUP-D
- 7500R2-36CQ-LC
- 7500R2A-36CQ-LC
- 7500R2AK-36CQ-LC
- 7500R2AK-48YCQ-LC
- 7500R2AM-36CQ-LC
- 7500R2M-36CQ-LC
- 7500R3-24D-LC
- 7500R3-24P-LC
- 7500R3-36CQ-LC
- 7500R3-48Y4D-LC
- 7500R3K-36CQ-LC
- 7500R3K-48Y4D-LC
- 7500RM-36CQ-LC
- 7504R-FM
- 7504R3-FM
- 7508R-FM
- 7508R3-FM
- 7512R-FM
- 7512R3-FM
- 7516R-FM
- 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
- CCS-750X-48TP-LC
- CCS-750X-48ZP-LC
- CCS-750X-48ZXP-LC
- CCS-755-CH
- CCS-755-X3-SC
- CCS-758-CH
- CCS-758-X3-SC
- DCS-7304
- DCS-7308
- DCS-7500-SUP2
- DCS-7500-SUP2-D
- DCS-7504
- DCS-7504N
- DCS-7508
- DCS-7508N
- DCS-7512N
- DCS-7516-SUP2
- DCS-7516-SUP2-D
- DCS-7516N
- DCS-7800-SUP
- DCS-7800-SUP1A
- DCS-7800-SUP1S
- DCS-7804-CH
- DCS-7808-CH
- DCS-7816-CH
- DCS-7816-SUP
- DCS-7816-SUP1S
- FAN-752M-RED
- PWR-2421-HV-RED
- PWR-2700-DC
- PWR-2700-DC-F

- 7368X4-SC
- 7388
- 7388-16CD
- 7388-8D
- 7388-SUP
- 7388-SUP-D
- 7388X5-SC
- 7500R-36CQ-LC
- 7500R-36Q-LC
- 7500R-48S2CQ-LC
- 7500R-8CFPX-LC
- 7500R2-18CQ-LC
- 7800R3AK-36DM-LC
- 7800R3AK-36DM2-LC
- 7800R3K-36DM-LC
- 7800R3K-48CQ-LC
- 7800R3K-72Y-LC
- 7804R3-FM
- 7808R3-FM
- 7808R3-FM2
- 7816R3-FM
- CCS-750-SUP100
- CCS-750-SUP25
- PWR-2700-DC-R
- PWR-2900AC
- PWR-3351-AC-RED
- PWR-3K-AC
- PWR-3K-AC-F
- PWR-3K-AC-R
- PWR-3K-DC-BLUE
- PWR-3K-DC-RED
- PWR-3KT-AC-BLUE
- PWR-3KT-AC-RED
- PWR-D1-3041-AC-BLUE

Transceiver

- 1000BASE-BX10
- 1000BASE-BX10-D
- 1000BASE-BX10-U
- 1000BASE-LX
- 1000BASE-SX
- 1000BASE-T
- 100BASE-FX
- 100GBASE-AOC
- 100GBASE-CR4
- 100GBASE-CWDM4
- 100GBASE-DR
- 100GBASE-ERL4
- 100GBASE-FR
- 100GBASE-LR
- 100GBASE-LR4
- 100GBASE-LRL4
- 100GBASE-PLR4
- 100GBASE-PLRL4
- 100GBASE-PSM4
- 100GBASE-SR4
- 100GBASE-SRBD
- 100GBASE-SWDM4
- 100GBASE-XCWDM4
- 100GBASE-XSR4
- 100GE-DWDM2
- 10GBASE-AOC
- 10GBASE-CR
- 10GBASE-DWDM
- 10GBASE-DWDM-T
- 10GBASE-DWDM-ZR
- 10GBASE-ER
- 10GBASE-ERBD
- 10GBASE-ZR
- 200GBASE-CR4
- 200GBASE-SR4
- 25GBASE-AOC
- 25GBASE-CR
- 25GBASE-LR
- 25GBASE-SR
- 25GBASE-XSR
- 40GBASE-AOC
- 40GBASE-CR4
- 40GBASE-ER4
- 40GBASE-LR4
- 40GBASE-LRL4
- 40GBASE-PLR4
- 40GBASE-PLRL4
- 40GBASE-SR4
- 40GBASE-SRBD
- 40GBASE-UNIV
- 40GBASE-XSR4
- 50GBASE-CR
- ADPT-O-Q-100G
- AOC-D-D-400G
- AOC-O-O-400G
- CAB-D-2Q-200G
- CAB-D-2Q-400G
- CAB-D-4Q-200G
- CAB-D-4Q-400G
- CAB-D-8S-200G
- CAB-D-8S-400G
- CAB-D-D-400G
- CAB-O-2Q-200G
- CAB-O-2Q-400G
- CFP2-100GBASE-XSR10
- CFPX-100G-DWDM
- CFPX-200G-DWDM
- H-D400-4Q100
- H-O400-4Q100
- OSFP-400G-2FR4
- OSFP-400G-DR4
- OSFP-400G-FR4
- OSFP-400G-LR4
- OSFP-400G-LR8
- OSFP-400G-PLR4
- OSFP-400G-PLR4-S
- OSFP-400G-SR8
- OSFP-400G-XDR4
- OSFP-400G-XDR4-S
- OSFP-400G-ZR
- OSFP-400G-ZRP
- OSFP-AMP-ZR
- QDD-400G-DR4
- QDD-400G-FR4
- QDD-400G-LR4
- QDD-400G-LR8
- QDD-400G-PLR4
- QDD-400G-PLR4-S
- QDD-400G-SR8
- QDD-400G-XDR4
- QDD-400G-XDR4-S
- QDD-400G-ZR
- QDD-400G-ZRP
- QSFP-200G-AR4
- QSFP-200G-FR4
- SFP-10G-MRA-T

- 10GBASE-ERBD-D
- 10GBASE-ERBD-U
- 10GBASE-ERLBD
- 10GBASE-ERLBD-D
- 10GBASE-ERLBD-U
- 10GBASE-LR
- 10GBASE-LRL
- 10GBASE-SR
- 10GBASE-SRL
- CAB-O-4Q-200G
- CAB-O-4Q-400G
- CAB-O-8S-200G
- CAB-O-8S-400G
- CAB-O-O-400G
- CAB-Q-2Q-100G
- CFP2-100G-DWDM-DCO
- CFP2-100GBASE-ER4
- CFP2-100GBASE-LR4
- SFP-10G-T
- SFP-10GRA-T
- SFP-25G-MR-LR
- SFP-25G-MR-SR
- SFP-25G-MR-XSR
- SFP-25GR-LR
- SFP-25GR-SR
- SFP-25GR-XSR

Reference to MLAG ISSU Upgrade/Downgrade Table

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

Resolved Caveats in 4.28.0.2F

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

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

- Fabric drops, namely DCLfifo drops, can be seen after agent SandFapNi restart due to misconfiguration of fabric balancing.
  Workaround is to power cycle any line cards where SandFapNi agent restarted. (723647)

DCS-7800R3 Series

- Disabling counter features may cause the SandCounters agent to restart continually under certain traffic conditions which will affect counter reporting.
  This is most noticeable when using hardware-accelerated sFlow and the Sflow agent restarts, or when turning the feature on and off.

  A workaround is running `reset platform fap <switch chip name> full`.
  Running this command may affect traffic. (709401)

Known Software Caveats in 4.28.0.2F

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)

- Performing a forwarding mode change after switch initialization with empty linecard slots will cause an extended traffic outage on subsequent SSOs. The workaround is to perform a full system restart after configuring the forwarding mode. (561925)
  SKUs Affected: 7324X-FM and 7328X-FM

- Performing SSU upgrade from source image prior to 4.28.0 may result in a failed SSU and following cold boot if 10GBASE-T is inserted into QSFP ports 49-56. This will lead to an extended (1-2 minutes) traffic outage on all interfaces. (642331)
  SKUs Affected: DCS-7050SX3-48C8, DCS-7050SX3-48C8-F and DCS-7050SX3-48C8-R

- Configuring more than 8 interfaces on QSFP ports 53-56 and performing SSU upgrade from source image 4.21.6-4.21.15, 4.22.0-4.22.3, 4.22.5-4.23.0 to destination image 4.27.2 and newer may result in a failed SSU and a following cold boot. This will lead to an extended (1-2 minutes) traffic outage on all interfaces. (664403)
  SKUs Affected: DCS-7050SX3-48YC8, DCS-7050SX3-48YC8-F and DCS-7050SX3-48YC8-R

BGP EVPN L2/L3 VXLAN/MPLS

- In an EVPN VXLAN deployment with VLAN aware bundle services, configuration a name that is longer than 89 characters may restart the BGP agent. (499958)

- 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)

- In an environment with higher scale (more than 1500 VLANs) deleting the VXLAN interface can cause BGP to restart . (629583)

- 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-7280SR, DCS-7280SR2, DCS-7280SR3, DCS-7280TR, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 Series

- In an EVPN deployment, the MAC/IP advertisement routes for the virtual router MAC address may fail to advertise if the VLAN for the route is assigned to a non-default IP VRF. This may cause traffic that should be flooded to be lost.

  A workaround is to toggle the "vrf" configuration under the "interface Vlan" corresponding to the VLAN in the problem state. (681910)

- Installing an EVPN MAC-IP route with MAC address value of zero may cause BGP process to restart. (696815)

- In an EVPN IRB setup with VXLAN, if a host moves from being local to being remote relative to a switch, traffic from behind the switch towards the now-remote host may not get resolved.

  The workaround is to issue a ping from the switch towards the now-remote host. (704579)

Containerized EOS

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

vEOS Router / CloudEOS

- With DPS enabled, there is a small chance that SFE agent may restart unexpectedly after a DPS path flap event. (527588)

- 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)

- As Microsoft rolled out a new root certificate, all the Azure PayG image Licenses will stop working. The workaround is to request a SWIX file with the fix or upgrade the EOS image. (685843)

- The /var/log/ filesystem on AWS CloudEOS VMs may get filled up with awslogs.*backup files. The workaround is to delete the /etc/cron.d/awslogs_log_rotate file. This workaround is better applied using an on-boot event handler that deletes this file on every reboot of the VM. (713562)

- When DPS path gets deleted, there's a small chance that SFE agent may restart unexpectedly. (713574)

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)

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)

- The OpenStack agent may restart unexpectedly when an OpenStack baremetal port is unbound and rebound to the same switch interface (663654)
- 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)
- If an OpenStack region with over 4000 VMs attempts to sync with the OpenStack service on CVX, it may fail, preventing dynamic VLANs and VLAN to VNI mappings from being provisioned. When this issue occurs, the error message "client intended to send too large body" will appear in /var/log/nginx-error.log on CVX. (704851)
- In a CVX deployment, a switch's mount of CVX state might be stuck in the "InProgress" state if there is link flap between the switch and CVX. The status of the switch mount of CVX state is available through "show management cvx mounts"

  In a MLAG deployment, a workaround is to trigger a failover on the switch with a mount in the "InProgress" state. In a non-MLAG deployment a workaround is to reload the switch. (704903)
- In a Macro-Segmentation Service (MSS-FW) deployment with a firewall manager in a high-availability setup, when a failover event occurs, CVX may not update the security rules.
  As a workaround, restart the MssPolicyMonitor agent on CVX. (730805)

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-7010T-48-DC-F, DCS-7010T-48-DC-R, DCS-7010T-48-F, DCS-7010T-48-R, 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-7050QX-32S-F, DCS-7050QX-32S-R, DCS-7050QX-32S-SSD-F, DCS-7050QX-32S-SSD-R, DCS-7050QX2-32S-F, DCS-7050QX2-32S-R, DCS-7050SX-64-F, DCS-7050SX-64-R, DCS-7050SX-64-SSD-F, DCS-7050SX-64-SSD-R, DCS-7050SX-72-F, DCS-7050SX-72-R, DCS-7050SX-72-SSD-F, DCS-7050SX-72-SSD-R, DCS-7050SX-72Q-F, DCS-7050SX-72Q-R, DCS-7050SX-96-F, DCS-7050SX-96-R, DCS-7050SX-96-SSD-F, DCS-7050SX-96-SSD-R, DCS-7050SX2-128-F, DCS-7050SX2-128-R, DCS-7050SX2-72Q-F, DCS-7050SX2-72Q-R, DCS-7050SX3-48YC12-F, DCS-7050TX-48-F, DCS-7050TX-48-R, DCS-7050TX-48-SSD-F, DCS-7050TX-48-SSD-R, DCS-7050TX-64-F, DCS-7050TX-64-R, DCS-7050TX-64-SSD-F, DCS-7050TX-64-SSD-R, DCS-7050TX-72-F, DCS-7050TX-72-R, DCS-7050TX-72-SSD-F, DCS-7050TX-72-SSD-R, DCS-7050TX-72Q-F, DCS-7050TX-72Q-R, DCS-7050TX-96-F, DCS-7050TX-96-R, DCS-7050TX-96-SSD-F, DCS-7050TX-96-SSD-R, DCS-7050TX2-128-F, DCS-7050TX2-128-R, 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)

- Support for VRRPs in OpenConfig not fully implemented (268736)

- 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 a large number of rules are setup on a single 'ip access-list', the OpenConfig agent may use a lot of CPU time and my restart unexpectedly. (407002)
- If the number of PTP vlan's enabled on a switch using 'ptp vlan' command exceeds 65536, PTP agent will become irresponsive 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)
- When an OpenConfig NETCONF client is used to add new classes to policy map, the new classes are added before any previous classes. (451025)
- 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)
- Source interface configuration in connectivity monitor is only supported for physical interfaces. (564121)
- Local (admin) user is allowed to login with empty password when RADIUS over TLS is enabled (569426)
- Routed traffic loss may occur because an ARP/ND entry is missing from "show arp"/"show ipv6 neighbor" but is present in "show kernel ip arp"/"show kernel ipv6 neighbors" as REACHABLE, STALE, DELAY or PERMANENT.

  Workarounds are to either restart the KernelNetworkInfo agent, shut/no shut the interface, or delete and add the neighbor back with "bash sudo ip neigh delete .../bash sudo ip neigh replace ..." (600394)
- When many thousands of ipv6 routes are pointing out one single interface and that interface is brought down the kernel might run in interrupt

context for a few seconds non-stop and cause delay in packet processing. This can potentially lead to BFD or BGP flaps. (610136)

- "reload peer" configured with a delay such as "reload peer in 1 force" results in both supervisors being reloaded after the specified delay. (621572)

- Under rare circumstances, MLAG agent may get stuck in a "Connecting" negotiation status. This can happen if the MLAG agent on the peer switch is constantly restarting (e.g., due to another bug or a out-of-memory condition).
  The workaround is to resolve the conditions causing the restart on the peer switch and restart the local MLAG. (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)

- If 'neighbor PEER remove-private-as' is applied to a normal peer, it nondeterministically results in confederation segments being stripped from AS paths advertised to confederation peers. (648505)

- Deletion of Interface Reflector configuration may result in deletion of interface entry or unrelated interface attributes at YANG path /interfaces/interface. A workaround is to restart OpenConfig/Octa agent. (661985)

- When copp policy-map copp-system-policy in startup-config has more dynamic copp class configured than platform supports, then Strata slice agent restarts repeatedly on bootup.

  The workaround is to reduce number of dynamic copp classes in copp-system-policy in startup config. (662805)
  Series Affected: 7300X3, 7320X, 7368, 7388, DCS-7010, DCS-7010TX, DCS-7050CX3, DCS-7050QX, DCS-7050QX2, DCS-7050SX, DCS-7050SX2, DCS-7050SX3, DCS-7050TX, DCS-7050TX2, DCS-7050TX3, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060CX, DCS-7060CX2, DCS-7060DX4, DCS-7060PX4, DCS-7060SX2, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7250X, DCS-7260CX, DCS-7260CX3, DCS-7260QX, DCS-7260X, DCS-7260X3, DCS-7300X and DCS-7300X3 Series

- If we have N sub-interfaces all under a common parent interface with a common policy-map applied and a subset of them do not have an IPv4 or IPv6 address, then the policy-map for that address family may end up applying to none of these sub-interfaces.

  The workaround is to flap the policy application on the sub-interfaces which have both IPv4 and IPv6 addresses. Note that if you can get back into the

broken state again if you flap the policy application on a sub-interface that is missing an IP address. (663110)

- Some Check Point Management Server configurations cause API messages that are unrecognized by the MSS feature. In this condition, MSS will end up ignoring some configured security rules.

  Workaround: Evaluate if these constructs can be removed from the policy and modify the policy. (664954)

- Multiple triggers of on-logging event handler within one polling cycle combined with the "trigger on-config disabled" command causes missing some Syslog messages.

  Workaround: Consider removing the "trigger on-config disabled" command (667215)

- ConnectivityMonitor may restart unexpectedly if more than 250 http probes are configured. (672269)

- The OpenConfig or Octa agent may return an error indicating the YANG tree is in an invalid state when VRRP is configured.

  A workaround is to delete the VRRP config or restart the OpenConfig/Octa agent (673284)

- If there is control plane traffic congestion, then PTP event messages may not be sent. This can be confirmed by observing the "egress queue drop" counter incrementing continuously in the output of `show ptp interface counter drops`. (673797)

- De-configuring a PTP-enabled port channel while PTP is in either End to End Transparent Clock Mode or Peer to Peer Transparent Clock Mode may cause the PTP agent to unexpectedly restart. (682913)

- 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)

- System onboarding to CloudVision-as-a-Service with ZeroTouch Provisioning may fail. As a workaround, use a custom bootstrap script or a USB instead for device onboarding. (686904)

- The rasdaemon process may consume a large amount of memory. If this happens, it should be restarted with `sudo systemctl restart rasdaemon` from bash. (696276)

- With "interface defaults; ethernet; shutdown", applying an interface profile to an interface in the shutdown state causes the interface to lose the "shutdown" command in its running-config. If it is saved to startup-config,

the interface will become enabled after reload. In some versions, the profile will also show the "shutdown" command even though it was not added by the user. (704316)

- In rare scenarios under heavy system churn, the system can begin to run out of memory after a period of time. (704662)
- ConnectivityMonitor may report incorrect latency and jitter values at high probe scale. (711946)
- Configuring Rsvp leads to Octa or TerminAttr printing error messages about failing to decode a Rsvp::RsvpLerPathSpecCollForSource and other Rsvp types. It will also lead to continual memory growth in Sysdb, leading to Sysdb crashes due to virtual memory exhaustion, or the OOM killer triggering and resetting TerminAttr or Octa. (726737)

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)

Layer 1

- 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

- A flap in a BASE-T interface with a default speed configuration may cause the forwarding agent to restart. This will cause multiple interfaces to briefly flap and a traffic outage to occur. The workaround is to configure a non-default speed on the affected interface or resolve the cause of flaps. (535720)
  Series Affected: DCS-7050TX3 Series

- Ports with inserted transceivers which only support 40G (such as 40GBASE-SRBD) and a default interface speed configuration may reserve 4 logical ports (hardware interface resources) instead of 1 on boot up. This may cause other interfaces to become inactive and the syslog STRATA-4-HW_PORT_RESOURCE_FULL to be logged.

  The workaround is to configure 40G-only transceivers with the "speed 40g" interface configuration. After the configuration is applied, the forwarding agent must be restarted using the CLI command "platform trident reset". This will cause all interfaces to flap and a brief traffic outage to occur. (554109)
  Series Affected: DCS-7260CX3 Series

- On QSFP-DD/OSFP port EthernetX, if EthernetX/1 and EthernetX/3 are configured for two lane auto-negotation, 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-negotation 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

- Configuring interface "shutdown" followed by "no shutdown" on a port while it is transmitting control plane may result in it continuously transmitting corrupted (FCS-errored) packets. The workaround is to configure "shutdown" followed by "no shutdown" on the affected port again. (596671)

SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-722XPM-48Y4, CCS-722XPM-48Y4-F and DCS-7010TX-48

- Interfaces with 1000BASE-T SFP transceivers may report a false link up when configured with default, "no speed", "speed auto", or "speed auto 1000full". The workaround is to configure affected interfaces with "speed forced 1000full" instead. (614753)
  Series Affected: 7320X, DCS-7060CX, DCS-7060CX2, DCS-7060SX2 and DCS-7260CX Series

- 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 Series

- 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-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 may experience link flaps when running at speeds using PAM4 modulation. (704536)
  SKUs Affected: 7368-4D and 7368-4P

- 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

Layer 2

- If Ebra agent restarts on a VTEP that is running EVPN control plane with route type 5, then it can result in VxLAN traffic loss. (276063)
- 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)
- CVE-2021-27853 and CVE-2021-27861: It is possible to bypass L2 filtering controls. More information is available under Arista Security Advisory 80 (https://www.arista.com/en/support/advisories-notices/security-advisory/16276-security-advisory-0080) (615000)
- 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. (616951)
- If "dhcp server ipv4" and "dhcp server ipv6" are both configured on a SVI and a client class is configured with an assignment and a match criteria matching on layer 2 interface on one address family but not the other, DHCP

services will not run on the address family without the layer 2 interface matching configuration.

Workaround is to create a dummy client class with a dummy assignment matching on the same VLAN and an Ethernet interface not in that VLAN. (678581)

- A single-event upset (SEU) event on the EGR_MASK table can result in the StrataL2 agent restarting, possibly repeatedly.
  A workaround is to reset the affected line card. (680437)
  Series Affected: CCS-750 Series

- Condition: MACsec agent restarts unexpectedly, when speed of MACsec interface is changed to mismatching speed from its peer. Suppose Peer-A and Peer-B, (PeerA <---->PeerB) if PeerA configuration is changed to a speed which is mismatching with the speed of PeerB, MACsec agent restarts unexpectedly.
  Impact: Changing a mismatching speed to one of its peer causes interface status for PeerA in failed state leading to traffic disruption apart from restarting MACsec agent unexpectedly.
  Workaround: If speed of PeerA is again changed back to matching its peerB, PeerA interface status will be in successful state and traffic will continue. (681447)
  SKUs Affected: 7800R3-48CQM2-LC, 7800R3K-36DM-LC, DCS-7050CX3M-32S, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32P4 and DCS-7280SR3M-48YC8

- The BFD protocol configuration under the L2 protocol forwarding profile in startup config is not applied on the switch.
  Workaround is to re-configure BFD under the L2 protocol forwarding profile. (687189)

- By default LPNs will be advertised as zero in MKPDUs for MKA sessions using XPN cipher suites.
  By default non zero LPNs will be advertised in MKPDUs for MKA sessions using non XPN cipher suites. (692286)

- By default LPNs will be advertised as zero in MKPDUs for MKA sessions using XPN cipher suites. By default non zero LPNs will be advertised in MKPDUs for MKA sessions using non XPN cipher suites. (697122)

- Performing SSU in an MLAG setup with STP running and at least one interface in discarding state causes the system to reboot. (708624)

- In a VXLAN deployment using dot1q tunnel on a trunk interface, using the "all" keyword to denote the inner dot1q VLANs (or c-VLANs) when applying the config "switchport vlan translation out <VLAN> dot1q-tunnel all" may result in an unexpected restart of the VxlanSwFwd agent. The restart is continuous till the workaround is applied as described below.

The workaround is to not use the "all" keyword and instead specify the inner dot1q tunnel VLANs (or c-VLANs) as a value or a range value. For example, apply "switchport vlan translation out <VLAN> dot1q-tunnel <VLAN range>" instead of "switchport vlan translation out <VLAN> dot1q-tunnel all". (716937)

- When "dot1x mac based authentication always" is configured on a force-authorized port, Dot1x agent may restart unexpectedly when an EAPOL packet is received on the port.
  A workaround is to remove "dot1x mac based authentication always" config from the port.
  Another workaround is to change port control setting to "auto". (730489)
- 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)

Layer 3

- If a loopback interface is only configured with a link-local IPv6 address, OSPFv3 will not advertise the address. (131369)
- 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)
- In situations involving a route with a next hop IP getting deleted and added back
  without a next hop IP, static ARP entries falling in the same subnet as the prefix of
  this route may not be correctly added.

  Workaround is to restart the ARP agent. (436002)
- 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)
- In EOS, host route computation is facilitated by the Arp agent. In situations where the FIB routing table is at internet scale, particularly if all or most of the routes are tunneled, the Arp agent may encounter a heartbeat timeout, leading to the Arp agent restarting. Typically, the agent will run fine after restart. (484902)
- When changing the OSPF Router Id, OSPF will not flush its existing self-originated LSAs
  from the routing domain before the new Router Id takes effect (499551)
- Isis agent will restart if an interface has ipv4 unnumbered configuration and ipv6 global address configuration and has enabled ipv4 and ipv6 traffic engineering. (501374)
- When the IP reachability of the TI-LFA protected link is leaked from one level into another level on the router in level-1 and level-2, then on link down local convergence delay timer is immediately aborted for the protected prefixes. This may cause the traffic loss due to micro loop. (504629)
- 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)
- After an interface flap, static ARP and ND entries may not be correctly added back, leading to traffic loss for packets being sent to these hosts.

  Workaround is to restart the Arp agent. (580766)
- 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)
- If there are many ARP entries, Arp agent may restart unexpectedly due to a timing condition during network event. (628246)
- Disabling DHCP on the last DHCP-enabled interface in a VRF using "no ip address dhcp" or "no ip address" or deleting a VRF with at least one DHCP-enabled interface or changing the redundancy protocol or doing a switchover on a dual-supervisor switch would cause all the other DHCP-enabled interfaces to lose their IP address. The workaround is to shut the affected interfaces using "shutdown" followed by a "no shutdown" or remove the DHCP config using "no ip address dhcp" followed by enabling it using "ip address dhcp" or making the port non-routed using "switchport" followed by making it routed again with "no switchport" or just waiting for the dhclient to renew the lease with the DHCP server. (632448)
- CLI-configured ARP bindings may be replaced by learned VXLAN or MPLS ARP bindings.

  To prevent this, configure ARP binding using `neighbor-resolution` CLI for each interface that the binding should be present behind. (643211)
- After a system reload, software forwarding or control plane generated packets for routes
  which are leaked, point to tunnels or nexthop-groups, or when routing from one ip
  address-family to another (RFC5549) may not work.

  As a workaround a user can configure 'hardware forwarding id' under a loopback interface.

https://eos.arista.com/eos-4-25-0f/support-for-in-band-traffic-in-a-management-vrf-without-routed-svi-or-physical-interfaces/ (653156)

- If a host route with a matching static ARP/ND entry has a tunnel next hop and changes next hop to an interface that is not a tunnel, the Arp agent may restart. Afterwards, the Arp agent will correctly recover. (653407)
- If a host route with a matching static ARP/ND entry has a tunnel next hop and changes next hop to an interface that is not a tunnel, the Arp agent may not correctly add the static ARP/ND entry.

  Workaround is to restart the Arp agent. (660983)
- After a VRF is deleted and added back in rapid succession, if the KernelNetworkInfo agent is restarted, the Arp agent may unexpectedly restart if another new VRF is added. (663599)
- Cspf agent may restart when BGP ORR config is present with dynamic peers as route reflector clients and the BGP session with the dynamic peers flap multiple times (669848)
- After the Rib agent restarts, static ARP or ND entries may go missing. Workaround is to restart the Arp agent. (673292)
- 'mlag peer-link requests disabled' in the startup-config might not take effect after reloading the switch. A workaround would be to run 'agent DhcpRelay terminate' to restart the DhcpRelay agent (675687)
- With multiple IS-IS instances configured, and after graceful restart of the IS-IS agent, the IS-IS agent may restart again on any IS-IS config change under "router isis" mode. (685181)
- When the ARP / neighbour table becomes full, resolved entries may be garbage collected even if space could be made by freeing only failed entries.

  Workaround is to run 'sudo ip -all netns exec ip -6 neigh flush nud failed' and 'sudo ip -all netns exec ip -4 neigh flush nud failed' periodically. These commands clear 'failed' neighbor entries from neighbor and arp tables respectively. (686096)
- L3 forwarding agent can crash at very high resilient ECMP scale and constant system
  churn. The agent should recover after restart (689387)
- When DHCPv6 all subnets relay is configured, the gateway address will not be changed after any number of requests if the server replies with NoAddrsAvail. (691384)

- 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 IGP topology changes in a network which has broadcast (or LAN) type interfaces. (692209)
- If we remove and add a tunnel table from a tunnel RIB config after a Stateful Switchover (SSO) the Tunnel agent may restart. (693091)
- When an ARP request is sent by a host for the virtual router IP and there is a permanent ARP entry configured, the permanent entry will be overwritten in the kernel neighbor cache. (694068)
- Unicast DHCP/BOOTP replies sent from a DHCP server directly to the client may not get forwarded if 'tunnel requests disabled' is enabled. (698611)
- Cspf agent may sometimes restart when there are two or more routers with duplicate TE router ID in the network (700468)
- ISIS agent may restart when a duplicate xngb tlv in a different fragment in the LSPDB is withdrawn (703092)
- When IS-IS SR dynamic adjacency segments are configured for both IPv4 and IPv6 address-families and if there are multiple interface flaps causing the re-allocation of adjacency segments, then it can cause IS-IS agent to restart (705216)
- If a BGP route is programmed with a colored MPLS next hop, then IpRib agent can restart unexpectedly.
  Workaround is to not add BGP colored extended community for a given BGP prefix. (707313)
- In an EVPN MPLS setup, after an ARP or ND entry added by ArpSuppression is removed from the ARP table for the first time, the ARP or ND entry may be added back.

  Workaround is to manually clear the entry with `clear arp`. (709754)
- IS-IS agent may restart if a prefix segment for a connected route is configured and the interface corresponding to the connected route is enabled with IS-IS later and if there is a IS-IS config change that reinits the instance like: LSP buffer size update, IS level update, address-family update etc (710168)
- IS-IS agent will restart if conflicting prefix segments are received with different algorithms and later prefix segments are changed to configure with the same algorithms. (713452)

- If an interface route is removed and added back while there are multiple routes with less specific prefixes matching the prefix of the deleted route, static neighbor entries associated with the interface route may be removed.

  Workaround is to restart the Arp agent. (714572)

- When the next hop of a host route changes from one dynamic tunnel to another, neighbor entries matching the prefix of the route may go missing.

  Workaround is to restart the Arp agent. (715486)

- The DhcpRelay agent may not subscribe to the All_DHCP_Relay_Agents_and_Servers (FF02::1:2) multicast group on an interface after a reload.
  As a workaround, 'agent DhcpRelay terminate' can restart the DhcpRelay agent and fix the multicast group membership (718485)

- When IS-IS LSP gets flooded across large number of broadcast adjacencies, then IS-IS can purge an LSP pre-maturely due to its life time expiry. This can cause routes to flap.
  To reduce the risk of this happening, configure a high max LSP lifetime, for example: 'max-lsp-lifetime 65535' (726402)

- When a large number of routes are deleted and added back, if IAR occurs afterwards, the Arp agent may be killed with signal SIGQUIT. (730757)

- When IS-IS adjacency is running over an unnumbered interface, then restoring an IS-IS link (bringing to the up state an interface which was previously down, shutdown or in a err-disable state) can result in a flap of the adjacency and if the corresponding link is also part of an ECMP group it can result in a few packet loss. (730944)

- If an interface is flapped after a static ARP or ND entry is added to that interface, the recently added static ARP or ND entry may not be correctly installed.

  Workaround is to restart the Arp agent. (732384)

Multi-agent

- If a VRF is deleted while the Bgp agent is in the middle of processing paths in BRIB, it may restart unexpectedly. (209143)

- Performing a stateful switchover in a switch with a large number of VRFs configured, with BGP config in all VRFs and IGPs configured in a large number of VRFs, may cause the system to restart unexpectedly, leading to traffic loss. (376773)

- During BGP graceful restart, 'update wait-for-convergence' behavior is unnecessarily forced for AFI-SAFIs which do not have graceful restart configured (either specifically per AFI-SAFI or globally). (478038)
- 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)
- If a VRF from which routes are being exported to VPN table is deleted, BGP agent can restart unexpectedly and repeatedly.

  To workaround this issue, reload the router after deleting any VRF from which routes are being exported to VPN table. (621467)
- If a multicast route resolves over Unicast Rib which in turn is reachable via MPLS tunnel, then IpRib agent can restart unexpectedly. (647713)
- The Bgp agent will restart unexpectedly when 32 or more RPKI ROA tables are defined. (665090)
- In some EVPN deployments with large number or VRFs or VLANs or peer flaps, BGP agent could end up restarting unexpectedly due to a steady increase of memory usage (leak). (673399)
- If an RPKI cache server is moved from one table to another, the contents may not be properly detected by the Bgp agent. This will always happen when moving from the default table to a non-default table.

  As a workaround, configure the desired RPKI tables before enabling RPKI validation. (676251)
- If the "show bgp debug policy inbound neighbor" command is run and specifies a prefix which is being denied by a prefix list, the Bgp agent will restart. (680322)
- Redistribution of BGP routes into OSPF/ISIS using route-map in a non-default VRF may stop working if that VRF instance is deleted and re-created. The workaround would be to restart the Bgp agent. (686581)
- Aggregating imported VPN or EVPN routes with match-map configuration can cause the Bgp agent to restart unexpectedly. (688775)
- BGP command 'no bgp bestpath as-path multipath-relax' AS number comparison does not work as intended when ASN value is greater than or equal to 0x80000000 (2147483648) in one path and is less than 0x80000000 for the other path. (691117)

- When maintenance mode is used to quiesce BGP peers, dynamic peers may flap. (691704)
- A StaticRoute agent restart or a switch reload with BFD tracked static routes in multiple VRFs may cause such static routes to be incorrectly removed from the FIB.
  Configuring a new VRF instance, with BFD tracked static routes in a pre-existing VRF, may cause the static routes to be removed from the latter VRF. (692583)
- Multicast route's next hop that recursively resolve over the unicast route may not get updated when the next hop resolution for the corresponding multicast route changes. (694308)
- sFlow agent might restart continuously when there are leaked BGP routes that resolve over MPLS tunnels (695172)
- If IpRib agent is shutdown while in the process of recursively resolving next hops, then IpRib agent may restart in a loop.

  Rebooting the switch will resolve this issue. (705622)
- When MPLS L3 VPN Inter-AS Option B is enabled, back to back changes to the label allocated for a route may not be advertised downstream. (706074)
- The BGP agent can experience a memory leak when receiving path updates with different sets of attributes than previously received for the path. The leak only happens if sFlow or IPFIX is enabled. (713908)
- If the Bgp agent encounters a kernel error while accepting an incoming TCP connection, it will stop listening for further TCP connections from BGP neighbors.

  To recover from this state you must restart the Bgp agent, or do a 'shutdown' followed by 'no shutdown' in the 'router bgp' configuration. (714590)

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)

- 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)

MetaWatch

- The "is duplicate" bit (bit2) of the Metawatch trailer flags field is never set, even for packets detected as duplicates. (705565)

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)

- An extra socket on UDP port 51023 may be observed, causing ARP sync over MLAG to fail. The workaround is to restart the VxlanSwFwd agent, during which there might be a sub-second disruption of handling new ARP/ND traffic, though existing ARP/ND entries will not be affected. (685708)

- In VXLAN MLAG deployment using a configuration session to remove an MLAG configuration can cause VXLAN to be unable to allocate dynamic VLANs. CVP uses configuration session for provisioning.

  A workaround is to run the following commands:
  > switch(config)#mlag config
  > switch(config-mlag)#shutdown (687274)

MPLS

- Modifying the BSID of an SR-TE policy can result in the new BSID route for the policy being unprogrammed. As a workaround, remove and re-apply the policy after a brief delay to have it properly programmed. (615033)
- MPLS LFIB routes and tunnels may be slow to reconverge under heavy system load after a change in the network topology. This is a result of removing the route and tunnel when a downstream LSR is transiently lost. (678467)
- When the primary LSP of an RSVP-TE tunnel is configured at the LER with a pre-signaled secondary LSP for path protection, triggering FRR on the primary LSP may cause the Rsvp agent to restart unexpectedly. (700728)

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-learnt on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIgmp", consider increasing the bandwidth for this Copp class. (356675)
- In an environment with a very large number of MSDP peers, during cleanup the MSDP Agent may unexpectedly restart. (418190)
- After a SandMcast agent restart, it is possible that stale entries are left in the hardware causing packets to be forwarded incorrectly.
  A workaround is to restart all of the SandFap and SandFapNi agents. (461822)
  Series Affected: 7500R3, DCS-7020, DCS-7280CR2, DCS-7280CR3, DCS-7280DR3, DCS-7280PR3, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280SR, DCS-7280SR2, DCS-7280SR3, DCS-7280TR, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 Series
- 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)

- After a MLAG peer state change, IGMP snooping state might go out of sync between the MLAG peers, resulting in loss of multicast traffic to receivers attached to MLAG peers. In this state IGMP Snooping agent in an MLAG setup may use 100% CPU. The workaround is to restart IGMP Snooping agent. (619267)

- 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)

- On devices that have a non default control plane ACL and MLAG with IGMP Snooping and MLD Snooping running might observe traffic loss.

  Work around: allow TCP ports 5541(IGMP) and 5542(MLD)
    permit tcp any any eq 5541 ttl eq 255
    permit tcp any any eq 5542 ttl eq 255 (640273)

- 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

- When there is a churn in group membership or a unicast route to RP causes (*, g) routes to be deleted and re-added, (*, G) routes may not get programmed in hardware. This causes complete traffic loss for those groups.

  Triggers:
  This is applicable when spt-threshold is configured to infinity. Otherwise, (*, G) routes are not programmed in hardware and hence this bug is not applicable.

  Workaround:
  Sometimes 'clear ip mroute' will resolve this issue. Otherwise, 'shut and no shut' of the IIF interface of the (*, G) route will clear this condition. (676677)

- MVPN related multicast route might be programmed incorrectly after BGP restart on a scaled MVPN PE (677280)

- If BGP-MVPN feature is enabled and an new SVI is configured that is being used as PMSI VLAN for a MVPN VRF, then MVPN tunnel decap for that VRF will get affected and the traffic may get dropped.

Workaround is to configure an SVI that is not being used as a dynamic (PMSI) VLAN. (680200)

- 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 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 after issuing a "clear ip mroute" command and the device happens to be a first hop router for certain sources, (S,G) entries for those sources will not get created.
  Workaround is to restart Pimsm agent by issuing "agent pimsm terminate" command. (707691)

- In Multicast EVPN deployments, when "PIM ASM" is configured as underlay multicast protocol,
  ingress VTEP connected to the source will not be able to register to underlay RP.
  As a result, remote/egress VTEPs connected to the receivers may not able to discover and
  join the source VTEP underlay route.

  workaround: Configure 'SSM' instead of 'ASM' as multicast protocol (711559)

- When IGMP Snooping Agent is restarted while using EVPN OISM Multicast feature with higher scale of groups(tens of thousands of multicast groups), may cause multicast traffic loss.

  The workaround is to run this command:
  bash sudo rm /var/shmem/eos/ar/Shmem/multicast/ipv4/irb/membership/join/igmpsnooping/fenix

  Note that with this workaround, IGMP Snooping agent will restart and there may be multicast traffic loss for about 2 minutes for some groups. (712475)

- 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 excuted) being programmed in the hardware resulting in traffic loss.

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

NAT

- Changing the configuration of NAT synchronization connected MLAG peers might cause the NAT agent to unexpectedly restart. (680566)
  Series Affected: CCS-720XP and DCS-7050SX3 Series

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

  A workaround could be either to introduce a small delay between control packets or make sure that all the control packets for a specific connection arrive on the same peer. (681428)
  Series Affected: DCS-7170 Series

- When NAT sync is configured, a TCP connection between a client and server might never be established
  if the SYN and SYN-ack arrive on different peers. This is because the SYN-ACK is not sent to CPU
  due to an incorrect programming of the trap tables. (699544)

SwitchApp

- LAGs may not be correctly programmed after reloading SwitchApp. The workaround is to flap all port-channel ports once SwitchApp is loaded. (661352)

- A 7130 with a SwitchApp FPGA image running a L3 profile may fail to ARP for hosts out of an interface. The command 'show fpga switch ip route` will show that the connected prefix for the interface has action DROP rather than CPU_LOW.

  Shut/no shut of the interface should resolve the problem. (695100)

- Fastdrops installed for tagged packets received on trunk ports will be incorrectly installed in the native VLAN of the trunk port. This will cause unwanted packets to keep reaching CPU.

  A workaround is to manually install a static multicast route to fastdrop the unwanted packets. (697149)

- Routing traffic may not work on certain VLANs and bridging packets may be incorrectly flooded
  if the destination MAC (including router MAC) on those VLANs are programmed in the hardware
  with the hardware index value larger than 8192. The hardware index for a given MAC address and
  VLAN is show in the output of the exec-mode command "platform fpga switch diag table MacVlan". (701039)

- For IGMP snooping enabled VLANs, multicast forwarding may not function after a device reboot.

  One workaround is restarting the MakoMcast agent. This may result in a brief disruption for multicast traffic forwarded by the switch. (703051)
  Series Affected: DCS-7130 Series

- Multicast routing may not be correctly enabled for a VLAN.
  A workaround is to remove and then reapply pim configuration for the VLAN. (703235)

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)

- When multiple IPSec connections are created, some with AES encryption and some with NULL encryption and NULL authentication, bessd will crash when decrypting the IPsec packets encrypted with NULL encryption and NULL authentication. (378590)

- 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 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 workaround 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)

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)

- Using virtual router mac address with mask can result in routing loop for ip[v6] packets. (682400)

- Bfn* fowarding agents may restart due to oom in certain conditions of churn (682993)

- Under some race conditions, some of the NAT flows may age out even in the presence of active traffic. This could potentially cause a new translation to be created for the flow while in the middle of an active session since the NAT entry may be removed from the hardware prematurely.

  The workaround is to configure the aging time for NAT longer than the expected lifetime of NAT flows. (697259)
    Series Affected: DCS-7170 Series

- When running a profile which supports IPv6 in the VXLAN underlay, IPv6 will be used
  by default by the EVPN control plane if the VXLAN source interface has an IPv6
  address configured. As a workaround, remove the IPv6 address configuration from the
  VXLAN source interface, which will make EVPN use IPv4 in the underlay. (702852)

- On DCS-7170B running baremetal profile, when VXLAN is used in a MLAG configuration, L2 flooding loop may occur (707378)

DCS-7170 Series

- PTP transparent mode is currently not supported. Attempting to configure PTP transparent mode with the `ptp mode e2etransparent` or `ptp mode p2ptransparent` command will cause the BfnSlice agent to unexpectedly restart. (609143)
  Series Affected: DCS-7170 Series

CCS-710P, CCS-720XP 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 be 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)

- When hardware flow tracking is configured with "record format ipfix standard timestamps counters" (i.e. sw export), the IPFIX data packets generated by the hardware may get dropped on the CPU queue. The workaround is to increase the export interval by configuring "record export on interval <interval>". (552698)

- 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

- Performing SSO when 802.1X with Port Security is configured may result in extended traffic outage on the affected ports. (615942)

- When MACsec is enabled, Single Event Upsets (SEU), that result in ECC or parity errors in a device's switch chip may result in traffic bypassing encryption and thus egressing in cleartext. The peer may drop these packets or allow them, depending on the MACsec policy implemented.
  This issue can be identified by seeing both of the following conditions:
  1. The following syslog: "%HARDWARE-6-PERR_CORRECTED" with table "MODPORT_MAP_SUBPORT" and index of less than or equal to 60.
  2. "show mac security counters" is not showing OutPktsEncrypted or OutOctetsEncrypted rising as expected for data plane unicast traffic. (It may be incrementing for multicast, broadcast, unknown unicast or control plane traffic).

  Any of the following options can be used to work around this issue:
  1. run "shutdown", wait a few seconds and run "no shutdown" in interface config mode for the affected interface.
  2. remove and restore the "mac security profile" in interface config mode for the affected interface.
  3. run "platform trident reset" to reset the switch chip. Warning this affects all interfaces using the affected switch chip.
  4. reload the switch (620974)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- A flap in any BASE-T interface may cause the forwarding agent to restart. This will cause multiple interfaces to briefly flap and a traffic outage to occur. (626355)
  SKUs Affected: CCS-722XPM-48Y4 and CCS-750X-48TP-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

- Prior to an SAK being installed, traffic will be forwarded unprotected regardless of the "traffic unprotected" configuration in the MACsec profile. Unprotected traffic may be dropped or forwarded by the peer depending on its configuration. (654087)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- During MACsec rekey operation, encrypted packets may be dropped on the peer due to replay protection checks.
  To workaround this, please configure "replay protection disabled" on the peer. (669840)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- When removing a linecard or switchcard module, an internal operation may fail resulting in a "%PM853X-3-PORT_BIND_FAILURE" log message. Once this operation has failed, the card slot from which the module was removed may remain in a failed state indefinitely. In this failed state, when a module is reinserted into the affected card slot it may not initialize correctly.

  This can be worked around by restarting the MicrosemiAgent agent. (672051)
  Series Affected: CCS-750 Series

- With MACsec configured if StrataSec and Strata forwarding plane agents are restarted at the same time, MACsec might get deconfigured.
  Workaround is to restart the StrataSec forwarding plane agent. (674591)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- When MACsec is deconfigured via "shutdown" in "mac security" mode, egress traffic on a previously MACsec enabled ports may be silently dropped in the switch.

Workaround is to toggle with "no shutdown" and "shutdown" in "mac security" mode. (675935)
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

- With MACsec configured if Strata forwarding plane agents are restarted, MACsec might get deconfigured.
  Workaround is to restart the StrataSec forwarding plane agent. (691416)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- If StrataCounters agent is restarted while MACsec is enabled, MACsec may not be enabled on interfaces.

  Workaround is to restart StrataSec agent. (696301)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- A flap on Et1-Et24 BASE-T interfaces may cause the forwarding agent to restart. This will cause multiple interfaces to briefly flap and a traffic outage to occur (702037)
  SKUs Affected: DCS-7010TX-48 and DCS-7010TX-48-DC

- Physical removal of Active Supervisor while SSO is configured is not supported and may result in a failed switchover that does not recover.

  A workaround is to manually force a switchover using CLI "redundancy manual switchover" before physically removing the card. (713334)
  Series Affected: CCS-750 Series

- Packets received on MACsec enabled ports with Dot1q tunneling configuration may not be processed by control-plane. (732258)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

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

7368, 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

- Replacing certain linecards on modular systems with one that has a lesser switch count while SSO is enabled will result in agent restarts on the standby supervisor. This affects replacing a 7300X-32Q-LC with a 7300X-64S-LC and replacing a CCS-750X-48ZXP with either a CCS-750X-48TP or CCS-750X-48ZP.

  The workaround is to configure the system for RPR mode, replace out the cards, and then reconfigure for SSO mode. (586499)
  SKUs Affected: 7300X-32Q-LC and 7300X-64S-LC

- Shutting down a Fabric or Linecard module before switching over in SSO mode can prevent the card from reaching PCIe switchover-ready state again, after a power-up on the new active supervisor. (635502)
  Series Affected: CCS-750 and DCS-7800R3 Series

- If an active supervisor card removal results in an SSO (Stateful Switchover), and the supervisor card is not plugged back into the old active supervisor slot and the SSO operation fails, the switch reboot may not occur automatically.

To recover a reload operation must be issued to the only supervisor inserted. (698735)

7500R3, DCS-7020, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 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 Sand systems IPv4/IPv6 DoS attack can result in neighbor resolution entries in the kernel getting flushed out frequently. Reducing the shape rate for copp-system-l3dstmiss can help mitigate the issue. (151168)
- 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-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 linecards are powered on/off, there may be a traffic loss for a short duration on the interfaces where vlan mapping feature is configured. (254476)

- Unknown unicast and multicast packets are unnecessarily copied to CPU after VXLAN decapsulation when there is an SVI configured on the VXLAN VLAN. (258990)

- 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 the 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)

- When the outgoing interface of a VXLAN underlay route changes from one forwarding chip to another, it can negatively impact VXLAN overlay route convergence. (440233)

- 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)
- A configuration session where multiple class-maps/ACLs used by a primary PBR policy are created/modified may lead to the policy to not be programmed correctly in hardware if the interface(s) where this policy is applied is/are also covered by a fallback PBR policy.

  A workaround is to only update a single class-map or a single ACL in a config session if a primary PBR policy already includes the class-map/ACL to be updated.

  An alternate workaround is to use three different config pushes:
  (1) Remove the primary policy from all applied interfaces

(2) Update/create all required ACLs/class-maps
(3) Re-apply the primary policy to all applicable interfaces. (573311)

- 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)

- When the IP address on a IPsec tunnel on J2 is changed, the B52 agent can restart unexpectedly, causing all the front panel interfaces to go down momentarily (635711)

- With 'ip policy hardware persistent' configured, port subinterface PBR policies will fail to be installed from the startup-config on boot, after SandFapNi restart, and after linecard hotswap.

  To recover, either reapply the PBR policy to the port subinterface or restart the SandAcl agent. (652435)

- Adjusting TI-LFA tunnel counters priority to make TI-LFA tunnels more preferable when in a multi-level hierarchy of BGP-LU over SR-TE policies may cause the forwarding agent to restart. As a workaround, either modify the configuration to ensure BGP-LU tunnels are the most preferred for those hierarchies or remove the specific per-tunnel type counters priority configuration. (660868)

- When `hardware counter feature mpls tunnel` is enabled, issuing the `clear traffic-engineering segment-routing policy counters` command followed by a SandCounters agent restart may cause the SandCounters agent to restart continuously. (660971)

- The HARDWARE-4-CARD_EJECTOR_OPEN_SHUTDOWN log message may appear and fabric card shut down even though the ejector is closed. (663113)
  SKUs Affected: 7508R-FM

- When the EVPN VXLAN VLAN aware bundle configuration is changed, the SandL3Unicast agent may restart. (678131)

- When inserting a linecard, the adjacency update may cause momentary mis-forwarding of packets. (679075)

- With routed LAG ports configured, concurrent restarts of the SandTunnel and SandFap agents may result in incorrect forwarding on some of the routed LAG ports.

  An additional restart of the SandTunnel agent will fix this issue. (681560)

- If MAC learning on a VLAN containing an L2 subinterface is toggled, multiple features including ACL and Interface Policing may not work on the L2 subinterface.
  Workaround is to shut and un-shut the affected L2 subinterface. (686387)

- For a port-channel subinterface with multiple LAG members and a only tx-queue config, when the member that hosts the anchor interface goes down. The anchor interface may not automatically migrate to another LAG member.

  Removing and re-applying the QoS config should assign a new anchor interface. (693849)

- In the presence of rapid VLAN-to-VNI mapping changes, where a VNI is mapped/unmapped to multiple VLANs in quick succession, one or more mappings could go missing in hardware. This would lead to dropping of packets received
  with the corresponding VNIs. Workaround is to restart the SandTunnel agent. (703726)

- PTP internal time synchronization could be disrupted. When this happens, the message 'HARDWARE-3-TIMESYNC_ERROR' will be logged. The internal time sychronization and PTP will recover shortly without any user intervention. (711771)

- Configuring PTP may cause the Smbus and Babbage agents to crash intermittently (715238)
  SKUs Affected: DCS-7280CR3-96-F and DCS-7280CR3K-96-F

- Continuous shut/no shut on a Port-Channel with 1000s of subinterfaces can cause the SandTunnel agent to restart unexpectedly. (718924)

- 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)
- 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)
- 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)
- When operating in 8QAM modulation format, removal and reinsertion of a transceiver from an even port may cause a TX power spike from the transceiver in the odd port. The workaround is to administratively disable all Ethernet interfaces associated with the transceivers before physically removing them. (707181)
  SKUs Affected: 7500R-8CFPX-LC
- When a CFPX-200G-DWDM transceiver is inserted physically or using the 'no transceiver diag removed' CLI command, the actual transmit power might briefly be higher that the configured transmit power. (707195)
  SKUs Affected: 7500R-8CFPX-LC
- Unpadded packets of 35 or fewer bytes are silently dropped. Most packets of this length are padded to 60B plus a CRC. However, if we terminate a tunnel

of 25 or more bytes, then the packet might be unpadded and therefore dropped. (725131)

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)

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)

- PBR, QOS, CPU policies applied to VLAN interfaces may not function correctly if VLAN interface memberships change during SandFapNi restart, Linecard insertions or removals, or during interface speed changes.

  The workaround is to restart the SandTunnel agent, or reapply the policy configurations. (488958)

- 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)

- If a LAG with some subinterfaces configured with ACLs is deleted while SandAcl agent is shutdown, the ACL applied on such subinterfaces will continue to apply even after removing the ACL configuration on such subinterfaces. (519537)

- 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)
- Overlay VRF is not supported on IPsec tunnels (641507)
  SKUs Affected: DCS-7280CR3K-32D4 and DCS-7280CR3K-32P4
- On a modular system with multiple linecards running the SandFapNi agent, if all of the SandFapNi agents are manually restarted via a CLI command at the same time it creates a possibility that one of the agents will restart one or more times over the course of a few minutes before recovering. (661378)
- VXLAN encapsulated IPv4 link local packets may not get flooded in the VLAN after decapsulation if both IPv4 multicast routing and IGMP snooping is enabled
  in the VLAN. (673606)
- One IPsec tunnel going down or its config getting removed can impact traffic on other IPsec tunnels which share the same ACL config and are on the same overlay VRF.
  Workaround is to shut and unshut the impacted tunnel interfaces. (721254)

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

- Realtime, traffic-based Weighted Fair Queuing fails to initialize. (684153)

DCS-7800R3 Series

- After deleting a sub-interface with shaping or removing shaping configuration from a sub-interface, traffic will stop egressing its parent port.

  The work-around is to restart the SandBcmSec agent. Note that this action will cause all MACsec sessions to flap. (681135)

- 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)

- Affected linecards may report incorrect temperature readings for the on-chip temperature sensors 7 and 8. This may cause these linecards to be incorrectly shut down during large temperature changes. The show command "show system environment temp" may also display "Invalid" for these sensors. (689043)

- Parity interrupts on some MACsec hardware table entries may lead to continuous interrupts as well as traffic drops on non-MACsec ports. (703710)

- Static SAK with XPN ciphers incorrectly allows the programming of RX association numbers of 2 or 3 on 1-bit association number systems. (705628)
  SKUs Affected: 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3AK-36DM-LC and 7800R3AK-36DM2-LC

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)

- When IP packets small enough to require ethernet padding are forwarded into an IPsec tunnel, the padding bytes are not stripped. This may cause these packets to be dropped when decapsulated and decrypted by other IPsec products. (603398)
  SKUs Affected: DCS-7020SRG-24C2, DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

DCS-7280R3 Series

- A forwarding agent restart while traffic is being actively enqueued or dequeued to a chip may lead to ingress drops with VsqQbVsqASramBuffersMaxSizeRjct, VsqQbVsqCSramBuffersMaxSizeRjct, VsqQbVsqDSramBuffersMaxSizeRjct reject reason set.

  The following command can be used as a workaround to stop the drops but is not persistent across reboots:
  platform fap diag mod CGM_PB_VSQ_RJCT_MASK 0 1 VSQ_RJCT_MASK=0. (703177)

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-7020R, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

- 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)

- 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)

- When a shape rate is set on a tx queue of a subinterface, the SandFap agent may quit unexpectedly under some rare circumstances. (594292)

- In a EVPN VXLAN to VXLAN DCI deployment, an L2 loop may form between the gateway VTEP in the local domain and the gateway VTEP in the remote domain
  when there are no Type-2 MAC routes advertised by the remote domain gateway.

  Workaround is to add a dummy static MAC behind impacted remote domain VTEPs. (669987)

- Traffic policy matches on VLAN tags may incorrectly match untagged packets by interpreting part of the untagged packet as a VLAN tag. (671474)

- Configuring VPLS causes repeated restart of the SandL3Unicast agent, which may result in traffic loss.
  Work around is to remove VPLS configuration. (687971)

- Egress IPv6 Port ACLs on L3 Port channel sub interface can stop working when we shutdown all members on any given Fap. The workaround is to remove the shutdown members from the port-channel. (693173)

- If there is a PFC pause storm and the interface priority has ECN and priority-flow-control configured, the action of 'drop' can cause buffer corruption for the ingress chip with interrupt IQM DramDynSizeRollOver. (722624)
- 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)
- When the VXLAN layer2 ECMP feature and the 'ip hardware fib hierarchical next-hop urpf' CLI are both enabled, the SandL3Unicast agent restarts indefinitely. (725020)
- Egress sFlow is unable to sample certain tunnel terminated packets, e.g., GRE packets decapsulated via a tunnel interface, MPLS decap packets. (734781)

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

- Under sustained heavy load, when more than one interface is receiving traffic at linerate capacity, the SandCounters process may restart repeatedly until traffic is reduced. (543931)
- When a port is "shutdown" and "no shutdown" the traffic for a traffic class may stop egressing from the port. (549451)
- 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)
- When an interface traffic-policy containing both IPv4 and IPv6 match rules is applied to one or more L3 interfaces with a valid IPv4 and IPv6 addresses later has all IPv4 or IPv6 match rules removed results in the traffic-policy no longer being applied to the interfaces.

  Workaround is to re-install the policy on all interfaces (574937)
- Traffic policy matches on VLAN tags may incorrectly match untagged packets by interpreting part of the untagged packet as a VLAN tag. (597565)
- With hardware offload for sampled flow tracking (IPv4) enabled and flow tracker configured on subinterfaces, agent SandFlowTracker may unexpectedly restart if there are short-lived IPv4 flows on the subinterfaces. (607104)

- Latency based drop feature is not compatible with features using dynamic rate classes
  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)

- IPv4 packets routed via a IPv6 next hop will not be logged if the egress L3 subinterface had an ACL applied with a deny rule that also logged such packets. (621198)

- 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)

- If 4K L2 subinterfaces are configured on the same fap, then all of them may not come up. (634640)

- Egress IPv6 ACL logging do not work on packets subject to layer 3 MPLS termination. (640491)

- 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)

- 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 an ACL 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)

- An interface with a link down debounce of less than 3 seconds configured may not respond to a local fault deassertion at the physical layer fast enough causing the link to spuriously flap.

  A workaround is to configure a longer link down debounce time. (650423)

- IPv4 link local packets received on front panel ports or sourced by CPU do not
  get headend replicated to remote VTEPs when IPv4 multicast routing is enabled
  in the VXLAN VLAN. (670370)

- When using the "mpls-evpn" TCAM profile, the CPU-generated IPv6 NS packet will not be flooded to the remote PE(s).
  The workaround is to create a copy of system profile mpls-evpn without "feature tunnel vxlan". (672266)

- Egress MAC ACLs applied on front panel ports might not match routed traffic (676843)

- When the device is configured with 'vxlan qos dscp ecn propagation decapsulation disabled', the inner DSCP for VXLAN decapsulated packets is not retained if the ingress port is in DSCP trust mode and global DSCP rewrite is enabled. (677549)

- Linecard hotswaps or SSO while traffic is being actively enqueued or dequeued to a chip may lead to ingress drops with VsqQbVsqASramBuffersMaxSizeRjct, VsqQbVsqCSramBuffersMaxSizeRjct, VsqQbVsqDSramBuffersMaxSizeRjct reject reason set. If PFC is not enabled for any of the ports on the FAP, following command can be used as a workaround to fix the drops:
  platform fap <FapName> diag mod CGM_PB_VSQ_RJCT_MASK 0 1 VSQ_RJCT_MASK=0. (681696)

- On applying or modifying an ACL on a GRE or IPsec tunnel interface, all the traffic coming in on the tunnel interface may start getting dropped or misforwarded. (681977)

- When
  (1) A traffic-policy is applied in the ingress direction of an interface
  (2) The same traffic-policy is applied in the egress direction of an interface, and
  (3) The egress application in (2) fails, then

  the traffic-policy will be removed from the ingress interface as well.

To work around this issue, create two distinct traffic-policies if it will be applied in both directions. (683463)

- When multicast routing is enabled in the overlay, certain VXLAN encapsulated IPv6 link local packets (e.g. IPv6 NS) are not properly forwarded after decapsulation. (686345)
- When multicast routing is enabled in the overlay, VXLAN encapsulated IPv6 link
  local packets with hop-by-hop extension do not respect the VXLAN split horizon
  rules. As a result, such packets might loop between VTEPs. (686611)
- When multicast routing is enabled in the overlay, VXLAN encapsulated IPv6 link
  local OSPF packets do not respect the VXLAN split horizon rules. As a result,
  such packets might loop between VTEPs. (686612)
- TCAM banks may fill up with dbIpv6MpTtBasic when both VRF selection policy with IPv6 match rules and IPv6 decap groups are configured.

  Workaround is to restart SandTunnel agent after removing VRF selection policy and IPv6 decap group configurations. (688133)
- When config-replace of configuration having VPLS commands is performed back-to-back in quick succession, the SandMact agent can restart unexpectedly. (690059)
- Successive addition and removal of port-channel members that changes the set of
  Linecard chips where the traffic can egress this port-channel might result in sustained traffic loss. The Workaround is to restart the SandL3Unicast agent. (700290)
- After system reload, all packets on an interface may be dropped with InDiscards incrementing in the output of "show interfaces counters discards".

  The workaround is to run "shutdown" followed by "no shutdown" on all problematic interfaces together. (700299)
- If QoS policy-maps are applied on L2 subinterfaces alongside L3 subinterfaces or VLAN interfaces, the traffic on the subinterfaces or the VLAN interfaces may match on policy-maps not attached to the said interfaces. (706148)

- A linecard hot swap in chassis with some linecards running tap-aggregation-extended TCAM profile and other configured with VXLAN VARP configuration can cause SandL3Unicast agent to restart. (712611)
- ECC errors in exact match tables may result into mis-forwarding of the packets. (726910)
- Repeated forwarding plane restart and/or linecard insertion/removal may cause a memory leak in the SandAegis agent. (733762)

DCS-7280R and DCS-7280R2 Series

- VXLAN packets that are decapped and bridged are incorrectly subject to the IP ACL applied on the ingress port. (221457)
- 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)

7300X3, 7320X, 7368, 7388, CCS-710P, CCS-720XP, CCS-750, DCS-7010, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7250X, DCS-7260X, DCS-7260X3, DCS-7300X and DCS-7300X3 Series

- Traffic streams matching DirectFlow flows with VLAN or MAC rewrite actions could face brief disruption across a StrataL3 agent restart. (111833)
  Series Affected: DCS-7050X 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-7010, DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7260X, DCS-7260X3 and DCS-7300X Series

- If nexthop resources are exhausted as indicated by syslog VXLAN_HWHER_NEXTHOP_RESOURCE_FULL, forwarding of BUM traffic to that VTEP will be affected. (246391)
  Series Affected: DCS-7050X and 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)

- 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)

- Some ports (esp. Et73-Et96 and Et103-Et104) might blackhole traffic or fail to linkup. (358361)
  SKUs Affected: DCS-7050TX-128-F, DCS-7050TX-128-R, DCS-7050TX-128-SSD-F and DCS-7050TX-128-SSD-R

- 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, DCS-7050X2 and DCS-7050X3 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-7050X, DCS-7050X2, 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. Workaround is to 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, DCS-7050X2 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)
  SKUs Affected: DCS-7060SX2-48YC6-F and DCS-7060SX2-48YC6-R

- 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)
  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)

- In scale policy-map configuration, in some cases, a config replace can lead to restart of Strata slice agent. (430654)
  Series Affected: DCS-7260CX Series

- During Leaf-SSU or SSO, interfaces with 100GBASE-SRBD transceivers may experience a single link flap. (440451)

- Hitless upgrade may fail if there are errdisabled interfaces.

  Workaround is to fix the speed misconfiguration to bring the interfaces out of errdisabled state. (445466)

- Changing Acl configuration rapidly and multiple times may cause the forwarding agent to repeatedly restart unexpectedly. (447584)

- The configuration applicable for user defined LAG/ECMP hash features may not be applied to the hardware during the switch (re)start and/or linecard insertion on modular devices. The following features may not work correctly

if configured:
 - IPv4/TTL and IPv6/Hop limit packet fields hashing for LAG and ECMP
 - user defined arbitrary LAG/ECMP hashing, feature is configured with the CLI:
      arista(config)#port-channel load-balance trident udf ...
- RRoCe protocol hashing

Workaround:
After the startup of the switch, change (either enable or disable depending on initial configuration) the setting for IPv4/TTL hash feature, then change it back, for example:
arista(config)#port-channel load-balance trident fields ip ip-tcp-udp-header ttl
arista(config)#port-channel load-balance trident fields ip ip-tcp-udp-header (450035)

- 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-7010, DCS-7050X, DCS-7050X2, DCS-7060X, DCS-7060X2, DCS-7250X, 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)

- 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)

- During an SSO failover, Strata Linecard agent may not detect hardware, resulting in a fatal error reboot with increased traffic outage. (512403)
  Series Affected: DCS-7300X Series

- 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)

- 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.

  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. (529148)
  SKUs Affected: CCS-720XP-24Y6, CCS-720XP-48Y6 and DCS-7010TX-48

- 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-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8, CCS-722XPM-48ZY8-F and DCS-7010TX-48

- 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

- When using a config session that deletes and recreates an existing ACL applied to one or more SVIs, and detaches the ACL from its applied SVIs in the same config session, the affected SVIs may be left in a state that blocks all traffic.

  Reapplying the ACL to the affected SVIs will unblock them. Using "hardware access-list update default-result permit" will prevent blocking the SVIs during ACL modification. (555485)

- VXLAN rule in ACL applied on SVI will cause Strata slice agent to continuously restart. Workaround is to remove VXLAN rule from ACL applied on SVI. (559747)

- 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

- Port channel(s) may fail to forward non-unicast traffic correctly under the following scenario:
  - port channels with interface members are configured
  - a new switch configuration is applied where interface(s) is moved from one port channel to another

  Workaround options:
  - restart StrataLag agent using CLI "agent StrataLag terminate"
  - reload the switch (586414)

- Performing SSU may result in a cold reboot due to DMA cancellation failure resulting in extended traffic loss. (589308)
    SKUs Affected: DCS-7050SX-64, DCS-7050SX-64-F, DCS-7050SX-64-R, DCS-7050SX-64-SSD-F and DCS-7050SX-64-SSD-R

- When many VLANs and/or LAGs are configured, port channels may flap during SSU. (592390)

- DHCPv6 unicast packets which are not destined to switch are copied to CPU, which may forward the packet to destination along with one copy forwarded in hardware resulting in duplication of unicast DHCPv6 packets. (595070)
- A SSU reload may result in the SCD watchdog causing an extended traffic outage. (605183)
- PIM over GRE interface does not support over 1K multicast routes (605368)
- When PTP is enabled, the slice agent may unexpectedly restart during a hitless restart. This will result in a hitful restart of the slice agent. (605632)
  Series Affected: DCS-7050TX, DCS-7050X, DCS-7260CX, DCS-7260QX and DCS-7300X Series
- MACsec sessions may fail to come up due to TCAM programming failures.

  Workaround is to deconfigure some features, and restart StrataSec forwarding plane agent. (614696)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- 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
- When the switch is reloaded, StrataL3 forwarding agent may restart unexpectedly. (636201)
- The StrataL3 Agent may unexpectedly restart when running the "show platform trident l3 software mgroups" on a switch with a large number of multicast nexthops. (636683)
- In a EVPN VXLAN Single-Active Multihoming deployment, VXLAN encapsulated BUM traffic ingressing on the Designated Forwarder's ES peer may get dropped. (637154)
- If a Port Channel with an inactive member interface is part of a VLAN, BUM traffic on that VLAN might get blackholed. Workaround is to remove the inactive interface from the Port Channel configuration. (638303)
  Series Affected: 7300X3, CCS-710P, CCS-720XP, DCS-7010TX, DCS-7050CX3, DCS-7050SX3, DCS-7050TX3, DCS-7050X3, DCS-7060DX4, DCS-7060PX4, DCS-7260CX3, DCS-7260X3 and DCS-7300X3 Series

- After a reload or after restart of Strata agent, CPU queues that were disabled with a zero shaper value become re-enabled.

  Configuring the CPU queue to have a non-zero shaper value then back to a zero shaper value will appropriately disable the CPU queue again. (641617)

- 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)

- When the switch reloads or when StrataL3 forwarding agent restarts and when the syslog "STRATA-3-VIRTUALPORT_RESOURCE_FULL " is also seen, StrataL3 forwarding agent may not come up resulting in continuous traffic loss. Workaround is to reduce the scale of remote VTEPs to prevent virtual port resource exhaustion. (645475)

- 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)

- After an invalid MMU queue 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 profile (661229)

- Multicast routes with same OIF sets are programmed non optimally which may lead to quicker exhaustion of hardware resources.

  There is no workaround. (662810)

- 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)

- Strata forwarding plane agent may spuriously detect control plane forwarding stuck conditions and unexpectedly restart. This results in a brief traffic outage. A known workaround to prevent the false positive stuck condition, that results in the agent reset, is to enable sFlow with a low sampling rate such as 1/16K. (666153)

- 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)

- Rapid changes in PCS Rx link status and/or MAC/RS link fault status (local/remote fault) might result in link staying down indefinitely or link coming up but reporting FCS errors.

  Workaround is to run `shutdown` followed by `no shutdown` on the affected interfaces. (671983)
  Series Affected: DCS-7050X, DCS-7050X2, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260X, DCS-7300X and DCS-7300X3 Series
  SKUs Affected: DCS-7010T-48, DCS-7010T-48-DC, DCS-7010T-48-DC-F, DCS-7010T-48-DC-R, DCS-7010T-48-F, DCS-7010T-48-R and DCS-7050SX3-96YC8

- All traffic to MACsec enabled ports may be dropped indefinitely, if both StrataSec and Strata forwarding agents unexpectedly restart at the same time.

  The workaround is to restart the Strata forwarding agent(s). (677396)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- An interface with a link down debounce configuration may stop forwarding traffic in following cases:

  1. When the interface becomes errdisabled for a period that is shorter than the configured link-debounce time on the interface.
  2. When the interface is reconfigured via CLI which does not result in a link down.
  3. When the transceiver is hot-swapped within a time interval shorter than the configured link-debounce time on the interface.

  The workaround is to run "shutdown" followed by "no shutdown" on the affected interface or their connected peer, with a gap between the two command greater than the configured link-debounce time. (678518)
  Series Affected: 7388, CCS-710P, DCS-7010TX and DCS-7050TX3 Series
  SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-24ZY4-F, CCS-720XP-48ZC2, CCS-720XP-48ZC2-F, CCS-720XP-96ZC2, CCS-720XP-96ZC2-F, CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- An interface with a link down debounce configuration may stop forwarding traffic in following cases:

  1. When the interface becomes errdisabled for a period that is shorter than the configured link-debounce time on the interface.
  2. When the interface is reconfigured via CLI which does not result in a link down.

3. When the transceiver is hot-swapped within a time interval shorter than the configured link-debounce time on the interface.

The workaround is to run "shutdown" followed by "no shutdown" on the affected interface or their connected peer, with a gap between the two command greater than the configured link-debounce time. (678519)
SKUs Affected: CCS-720XP-24Y6 and CCS-720XP-48Y6

- An interface with a link down debounce configuration may stop forwarding traffic in following cases:

  1. When the interface becomes errdisabled for a period that is shorter than the configured link-debounce time on the interface.
  2. When the interface is reconfigured via CLI which does not result in a link down.
  3. When the transceiver is hot-swapped within a time interval shorter than the configured link-debounce time on the interface.

  The workaround is to run "shutdown" followed by "no shutdown" on the affected interface or their connected peer, with a gap between the two command greater than the configured link-debounce time. (678520)
  SKUs Affected: DCS-7050CX3M-32S

- An interface with a link down debounce configuration may stop forwarding traffic in following cases:

  1. When the interface becomes errdisabled for a period that is shorter than the configured link-debounce time on the interface.
  2. When the interface is reconfigured via CLI which does not result in a link down.
  3. When the transceiver is hot-swapped within a time interval shorter than the configured link-debounce time on the interface.

  The workaround is to run "shutdown" followed by "no shutdown" on the affected interface or their connected peer, with a gap between the two command greater than the configured link-debounce time. (678521)
  Series Affected: DCS-7260CX3 Series
  SKUs Affected: DCS-7050CX3-32S

- An interface with a link down debounce configuration may stop forwarding traffic in following cases:

  1. When the interface becomes errdisabled for a period that is shorter than the configured link-debounce time on the interface.
  2. When the interface is reconfigured via CLI which does not result in a link down.
  3. When the transceiver is hot-swapped within a time interval shorter than the configured link-debounce time on the interface.

The workaround is to run "shutdown" followed by "no shutdown" on the affected interface or their connected peer, with a gap between the two command greater than the configured link-debounce time. (678522)
Series Affected: 7388 Series
SKUs Affected: DCS-7050SX3-48C8, DCS-7050SX3-48YC12 and DCS-7050SX3-48YC8

- An interface with a link down debounce configuration may stop forwarding traffic in following cases:

  1. When the interface becomes errdisabled for a period that is shorter than the configured link-debounce time on the interface.
  2. When the interface is reconfigured via CLI which does not result in a link down.
  3. When the transceiver is hot-swapped within a time interval shorter than the configured link-debounce time on the interface.

  The workaround is to run "shutdown" followed by "no shutdown" on the affected interface or their connected peer, with a gap between the two command greater than the configured link-debounce time. (678523)
  Series Affected: DCS-7060DX4 and DCS-7060PX4 Series

- When "hardware counter feature vtep encap" CLI is configured, the deletion of VTEPs can trigger an unexpected restart of the StrataL3 agent. (678643)

- Committing a CLI config session containing "hardware access-list update default-result permit" command and router access-list updates may result in permanent loss of all data plane and control plane traffic for the VLAN corresponding to router access-list.

  Workaround is to reload the device. (680577)

- 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-7050X, 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

- The "show platform trident forwarding-table partition" command might cause a deadlock and unexpected restart the forwarding agent. (690958)

- When a flooded L3 packet is sampled by sFlow, the packet may not be forwarded as well as sampled. (691287)

- Switch may run out of resources while adding ports to VXLAN VLANs resulting in "STRATA-3-VLANXLATE_RESOURCE_FULL" syslog messages and not HER forwarding the traffic from those ports. Workaround is reduce the no of VLANs or no of ports associated with those VLANs. (696429)
  Series Affected: 7320X, DCS-7050QX, DCS-7050SX, DCS-7050TX, DCS-7050X, DCS-7060, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260CX, DCS-7260CX3, DCS-7260QX and DCS-7260X Series

- When configuring MLAG before the system is fully warm, the system can restart unexpectedly due to a kernel panic. (701820)

- A port channel may be brought down and later up during SSU upgrade, which may cause temporary traffic loss. (701831)

- If an active supervisor card removal results in an SSO (Stateful Switchover), the new active supervisor may get partial state update of unrelated card removal. This may cause SSO to fail, and the switch will reboot to recover. (703676)
  Series Affected: 7300X3, 7320X, DCS-7300X and DCS-7300X3 Series

- Changing the PTP mode may not get applied to PTP enabled LAG interfaces.

  Workaround is to disable then enable PTP on the LAG interface after changing the PTP mode to have the config applied. (708333)

- Multicast scale has been reduced from 16k entries to 8k entries. (709296)
  Series Affected: DCS-7300X3 Series

- 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-7060X4 Series

- Rapid changes in PCS Rx link status and/or MAC/RS link fault status (local/remote fault) might result in link staying down indefinitely or link coming up but reporting FCS errors.

  Workaround is to run `shutdown` followed by `no shutdown` on the affected interfaces. (716821)
  Series Affected: 7368 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)

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

  Workaround is to run 'shutdown' followed by 'no shutdown' on the affected interfaces. In some cases, the workaround may be needed to be applied multiple times. (721884)
  Series Affected: CCS-710P Series

- 1000BASE-T ports Et25-Et48 may fail to link up after a reload or forwarding agent restart.

  Workaround is to run 'shutdown' followed by 'no shutdown' on the affected interfaces. In some cases, the workaround may be needed to be applied multiple times. (723050)
  SKUs Affected: CCS-722XPM-48Y4 and DCS-7010TX-48

- 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)

DCS-7010 Series

- 802.1x pause and guaranteed bandwidth cannot be configured at the same time. (91141)

- IPV6 traffic received over IPV4 tunnel interface may be dropped.

  The workaround is to configure IPv6 address on all the ingress L3 underlay interfaces. (545486)

- The switch will not boot if the real-time clock (RTC) battery is depleted. (675835)
  Series Affected: DCS-7010 Series

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

- If RPR or SSO is configured, Strata fabric agent restarts may cause switchover to occur. (368082)
  SKUs Affected: DCS-7304 and DCS-7308

- 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)
- Load balancing across ECMP members will not work when DLB is enabled and there are more than 31 members present.
  When issue occurs "RX - Port bitmap zero drop condition" counters will increment in below command:
  #show platform trident counters | nz

  Work-around is to limit the number of ECMP members upto 31. (671066)

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)

7388 Series

- StrataMmu agent has high CPU usage when PFC watchdog is configured with a low polling interval, such as 0.010. (636394)
- After an invalid MMU queue 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 profile (644909)

DCS-7050X, DCS-7250X and DCS-7300X Series

- A port operating at 100Mbps with a 1000BASE-T SFP adapter may not transmit packets to its link partner or may transmit corrupt packets failing FCS checks at the link partner, even if it successfully links up.

  Workaround is to run a shut/no shut on the affected port. (80970)

- Configured as MLAG, link flap of one or more LAG peer-link member ports could cause brief flooding of packets over the peer-link and can also lead to double delivery of packets on the MLAG interfaces. (117870)

- Altering the priority of transmit queues of a port will result in some traffic loss on that port. (194460)

- A single-event upset (SEU) on the tables REPLICATION_FIFO_BANK* may not be corrected, which may lead to traffic loss. (391625)

- After a reload or after an upgrade, an interface can stop transmitting despite having packets queued on the interface's egress queues. This can cause buffer buildup on the CPU port and eventually impact control plane transmission across all the interfaces.
  A workaround is to flap the port by running "shutdown" followed by "no shutdown" on the affected interface. (400610)

- In the presence of large amount BUM traffic across several interfaces, discards may be observed on some CPU queues. (563505)

DCS-7050X2 Series

- When an interface which is transmitting at linerate is disabled and re-enabled, it might end up in a condition where it is unable to transmit any more packets. Disabling the interface again causes the Strata-FixedSystem agent to restart following which the interface continues to transmit properly. (245501)

- Vxlan multicast decapsulation will not work for packets that are sourced from unknown source VTEP. (538407)

- Continuous flapping of L3 interfaces could cause the StrataL3 forwarding agent to restart, resulting in a momentary traffic loss. (550872)

- When the per VLAN routing feature is enabled, BFD packets will be dropped in HW resulting in existing sessions going down or preventing new sessions from getting established.
  Workaround is to disable per vlan routing feature. (568298)

CCS-710P Series

- A single-event upset (SEU) on the table L3_DEFIP_PAIR_128_DATA_ONLY may lead to a forwarding agent restart, after which the switch forwards traffic normally. (638982)
- An SEU in the L3_DEFIP table may cause L3 traffic to be misforwarded.

  Workaround: restart the StrataL3 after such an SEU (653348)
- Scaled IP Locking configuration may result in an unexpected restart of Strata forwarding agent. (678926)
  SKUs Affected: CCS-710P-12 and CCS-710P-16P
- A Singe Event Upset in the L3_DEFIP table may cause mis-redirected packets.

  Workaround: restart the StrataL3 agent (701736)

CCS-720XP and DCS-7010TX Series

- Ethernet25-Ethernet48 interfaces will continuously flap their links when configured with the default speed or "speed auto" commands, if the interfaces link up at 10Mbps or 100Mbps rates with half-duplex mode of operation. This might happen when the peer does not support the BASE-T auto-negotiation or when auto-negotiation resolves to 10Mbps or 100Mbps rates with half-duplex mode of operation.

  The workaround is to configure the affected interfaces using the "speed auto <SPEED>" command with the desired speed and duplex. (662833)
  Series Affected: DCS-7010TX Series
  SKUs Affected: CCS-722XPM-48Y4 and CCS-722XPM-48Y4-F
- During a MACsec rekey, a few packets may be dropped and show in 'show mac security counters detail'. (669187)
- Configuring "qos random detect allow non-ect" will cause Strata forwarding plane agent to unexpectedly restart. (709781)
  SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

CCS-720XP Series

- Setup with only Storm-control and scaled IP Locking configuration may cause slice agent to restart continuously. (599935)
- PoE ports may shut off unexpectedly or get stuck in the failed state when multiple PoE ports shut off simultaneously

Run "poe disabled" and "no poe disabled" to recover affected ports (672412)
SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F

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)

- 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)

- Stateful switchover may result into a higher than expected traffic outage. (677047)

- When multiple forwarding plane agent restarts or when switch is reloaded, VXLAN packets to be decapsulated and forwarded or Layer3 unicast traffic to some host routes may be dropped continuously.
  The workaround is to restart the StrataL3 forwarding agent. (721838)

- Packets received on downlink ports with Dot1q tunneling configuration may not be processed by control-plane. (732257)
  Series Affected: CCS-750 Series

DCS-7050X3 Series

- If removable 10GBASE-T transceivers are present, while some ports are disabled due to logical port exhaustion, and there is a forwarding agent restart, then there may be additional unexpected forwarding agent restarts which may cause loss of traffic. (570906)
  SKUs Affected: DCS-7050SX3-48C8, DCS-7050SX3-48YC8 and DCS-7050TX3-48C8

- Copper-based modules may experience FEC/FCS errors or link failure due to suboptimal low-level link configuration. (702176)
  SKUs Affected: DCS-7050SX3-48YC8

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

- XcvrAgent may restart repeatedly after the removal and insertion of a linecard that has QSFP or SFP transceivers installed in QSFP200 or QSFP-DD ports. (662431)

- When an SFP-10G-MRA-T is installed and no speed is configured, the SandFapNi agent may restart unexpectedly and repeatedly.
  The workaround is to configure a valid speed on the interface such as 'speed auto 1000full'. (684412)
  SKUs Affected: SFP-10G-MRA-T

- If the XcvrAgent unexpectedly restarts around the same time as a transceiver is reseated, the XcvrAgent may continuously restart. The workaround is to remove the reseated transceiver. The transceiver can be inserted again once the XcvrAgent stops continuously restarting. (701044)

- If a 400GBASE-FR4 transceiver is removed while the XcvrAgent is administratively shutdown or while the XcvrAgent is restarting unexpectedly, then the XcvrAgent will continuously restart unexpectedly.
  The workaround is to reinsert the removed 400GBASE-FR4 transceiver. (701047)
  SKUs Affected: OSFP-400G-FR4 and QDD-400G-FR4

7300X3, CCS-710P, CCS-720XP, 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

- 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)

- When address-only NAT is configured and the default NAT profile (NAPT) is configured, after a host is assigned a NAT IP address, ACL check is skipped for subsequent flows from that host. Such flows gets translated even if they match a deny ACL entry. (618773)

- If dynamic NAT rules with the same ACL are configured on more than 64 interfaces, the StrataNat agent may restart continuously.

Workaround is to use different ACL names such that no more than 64 interfaces share the same ACL name. (621189)

- Modifying the ACL configuration applied to the NAT rules with ECMP configuration is not supported.

  Workaround is remove the ACL, verify the rules are deleted from hardware and then re-apply the ACL. (623961)

- 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)

- If the forwarding agents are restarted with dynamic NAT configured, the StrataNat agent may repeatedly restart.

  The workaround is to cleanup dynamic connections before restart the forwarding agents. (640532)

- Dynamic NAT flows configured on a non-default VRF may not work after config-replace .

  Workaround is to restart the NAT agent. (642083)

- 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)

- When device is reloaded while DirectFlow config containing output interface or L2 rewrite actions is configured, StrataL3 agent may unexpectedly restart. (645665)

- When DirectFlow flow configured with CPU output interface and L2 rewrite action, StrataL3 may continuously unexpectedly restart. (647160)
- After a MAC moves to another MBVA enabled port, first subsequent reauthentication may result in the supplicant's applied ACLs cleaned up unexpectedly.
  Following subsequent reauthentications will recover the ACL. (672195)
- During config replace with NAT configuration, TCAM churn may result in forwarding agent restart. (676117)
- If 'dot1x mac based access-list' for a port is not disabled while dot1x is disabled globally, then subsequent Mac traffic policy configurations will not take effect (686449)
- Removal of VRF segment configuration along with deletion of the VRF may lead to stale segment hardware IDs and may result in running out of segment IDs for new segment configuration.

  The workaround is to restart StrataL3 agent. (696284)
- When 'ip address virtual ...' configured for a scaled SVI configuration with VXLAN, restart of StrataL2 agent may result in continuous restarting of StrataL2 forwarding agent.
  The work around is to reload the switch (716360)
- The "priority-flow-control priority ..." configuration is accepted even though generation of PFC frames is not supported. Applying this configuration could lead to the Strata agent crashing repeatedly.
  A workaround is to remove the unsupported "priority-flow-control priority ..." configuration. (718540)
  Series Affected: CCS-710P and DCS-7010TX Series
- In an EVPN VXLAN environment, IPV6 NS packets with destination MAC as broadcast when received on an ingress port will get flooded back out of the same ingress port. (729001)

Limitations and Restrictions in 4.28.0.2F

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-7050SX-128-F, DCS-7050SX-128-R, DCS-7050SX-128-SSD-F, DCS-7050SX-128-SSD-R, DCS-7050SX-64-F, DCS-7050SX-64-R, DCS-7050SX-64-SSD-F, DCS-7050SX-64-SSD-R, DCS-7260CX3-64, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F and DCS-7260CX3-64E-R

- Aborting 'reload fast-boot' command with Ctrl-C is unsupported (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)

vEOS Router / CloudEOS

- 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)

CVP

Telemetry

- Sampled flow tracking (IPFIX) and sFlow export flow records with incorrect nexthop for the flows using tunnel nexthops with hierarchical FECs (HFECs) having multiple Vias. (691829)

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 Po1
        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 Po1103    >>>>>>>> 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)
- 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 skips 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-7050SX-128-F, DCS-7050SX-128-R, DCS-7050SX-128-SSD-F, DCS-7050SX-128-SSD-R, DCS-7050TX-128-F, DCS-7050TX-128-R, DCS-7050TX-128-SSD-F, DCS-7050TX-128-SSD-R, DCS-7250QX-64-F, DCS-7250QX-64-R, DCS-7250QX-64-SSD-F, DCS-7250QX-64-SSD-R, 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)
- 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)
- vlan A cannot be both source and target of 'floodset expand vlan' configuration at same time, i.e. have 'floodset expand vlan <vlan B>' and be used in 'floodset expand vlan <vlan A>' to configure vlan C (497773)
- Installed licenses without corresponding licensing features configured might not be visible from the CLI (504793)
- 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 delcared 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)
- "match-prefix-set" in OpenConfig does not allow matching against IPv6 prefix lists (613502)
- "show reload cause" may show "Unknown" with no debug information if the reload reason is an internal CPU error. (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
- Match on multiple DSCP aliases under a QoS class-map is not supported. (633678)
- 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

- Sampled flow tracking (IPFIX) and sFlow export flow records with incorrect nexthop for the flows using tunnel nexthops with hierarchical FECs (HFECs). (677481)

- "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

- OpenConfig does not support BFD configuration for VLAN interfaces. BFD configuration under VLANs must be removed for any gNMI set operation to succeed. (682920)

- 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)

- If a flow traverses a tunnel, the extended router data nexthop in the sFlow record will be incorrectly set to the tunnel nexthop instead of the directly connected (IGP) nexthop. (694404)

- The "install bios" command can fail to identify some versions of the running BIOS and consider normally applicable updates as incompatible. (701500)

- Previous releases of EOS included EosSdkRpc .proto files that documented interfaces not yet supported by the implementation. The shipped .proto files from here on will only show those features implemented in the release. Documentation of future interfaces can be provided separately upon request. (715138)

- 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-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)

IPSEC

- We do not support NAT on Ipsec 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
- 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

- Cable diagnostics is not supported on any BASE-T ports. (669975)
  Series Affected: DCS-7010TX Series

- 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-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

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)

- There is no Accelerated Software Upgrade(ASU) support for 802.1X dynamic ACL assignment feature. (725100)

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-7050X, DCS-7050X2, 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)

- 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)

- In RSVP tunnels with pre-signaled secondary paths that undergo make-before-break (MBB) because of bandwidth changes, the bypass for the pre-signaled secondary LSP may flap. (700255)

- 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)

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)
- 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)

Rib

- 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)

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)

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-router interface" command in "config-mld-snooping-vlan" mode shows the unsupported "Switch" interface extension. (602413)
- Multicast routes are not created when a source is transmitting only fragemented multicast packets.
  Switch to using SFE for multicast software forwarding. (605326)

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)

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), 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)

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
- 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 tp 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)
- 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
- The 32-bit image of EOS does not support the 7170B series switches. Use the 64-bit version instead (646589)

DCS-7170 Series

- On the 7170 series of switches, when IPv6 in 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-720XP and CCS-750 Series

- 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

- 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

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)

- 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

- 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

7368, 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)
- On modular switches, the "show platform bios history" command can fail if the BIOS update history is missing from the standby supervisor. (674601)
- On modular switches configured in SSO redundancy mode, the Tpm agent does not start on the standby supervisor. (677335)

CCS-710P, CCS-720XP 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-7280R, DCS-7280R2, DCS-7280R3, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 Series

- The L3 MTU on interfaces is not enforced on SVIs. (11659)
  Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, 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)
- 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 +/-2 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)
- IPv6 Egress ACL deny logging is not supported on subinterfaces. (146487)
- PBR policy is supported only on interfaces in the default VRF. If a PBR policy is configured on an interface in a user defined VRF, the behavior is undefined. (148895)
- 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)
- 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)

- 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)
- 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)

- If a security-ACL is configured on a sub-interface beyond the supported limit of ACLs,
  SandAcl might restart continuously until the number of configured ACLs is brought back within the supported limits (398815)

- 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)

- Egress MAC ACLs are matched on the incoming packet's ethernet header and not on the rewritten ethernet headers. (458724)

- 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-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

- 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-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)

- 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)

- 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 matche 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 L2 MPLS EVPN configuration does not support PE to PE overlay ping. There is no available workaround. (617720)

- The SandAcl agent may restart repeatedly when more mirroring ACLs are appled than can be configured on any given FAPin 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)

- "show inventory" and "show module" commands only show Fabric1-Fabric6. Fabric7-Fabric12 are not shown. A workaround is to use "modulargenprefdl" in bash except during SSO events to query the module serial numbers or "show hardware eeprom" in newer EOS releases. (664931)
  SKUs Affected: 7816R3-FM, DCS-7816-CH and DCS-7816-SUP

- BGP LU route with swap label action, resolving over SR-TE is not supported with MPLS tunnel counter feature. (669633)

- 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)

DCS-7280R and DCS-7500R Series

- 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

- 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-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)
- 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)

- "hardware next-hop fast-failover" is not supported with "ip load-sharing prefer local" configuration. (650521)

- VXLAN is not supported with greater than 256-way ECMP in the underlay. (652266)

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)

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)

- 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)

- UDP packets with ports 4754 or 4755 may not be forwarded correctly. (648240)

7500R3, DCS-7280R3 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-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-7020R, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

- 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)

- 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)
- 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)
- 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
- 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)
- 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)
- 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)

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

- 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)
- 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)

- When MAC address is changed on the interface gratuitous ARP is not being sent. (536257)
- When MAC address is changed on the interface neighbor advertisement is not being sent. (539590)
- 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 if 100% throughput of data traffic (591523)
- 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)
- 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)

- 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)

- When 8-to-1 prefix compression is configured along with hardware next-hop fast-failover, the forwarding agent may restart unexpectedly.
  Workaround options:
  - Disable 8-to-1 prefix compression using "no ip hardware fib optimize"
  - Disable hardware next-hop fast-failover. (646417)

- 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)

- 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 edited or filtered on egress as if there were forwarded regularily.

  Do not use mirroring to multiple destinations outside of Tap Aggregation mode. (670046)

- L2 Protocol Forwarding profiles will not apply on selective Q-in-Q trunk ports. (674442)

- GUE tunnel decapsulation does not support UDP destination port 2152. (675837)

- 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)

- 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)
- 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

DCS-7280R and DCS-7280R2 Series

- CBF does not work in AlgoMatch mode on some custom TCAM profiles (603206)

7300X3, 7320X, 7368, 7388, CCS-710P, CCS-720XP, CCS-750, DCS-7010, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7250X, DCS-7260X, DCS-7260X3, DCS-7300X and DCS-7300X3 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)

- Forwarding-table partition mode 4 is not supported. (100862)
  Series Affected: DCS-7010 Series

- 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, systems will start sharing resources that are required for programming access-group/ access-list. In this mode only 24 port range checkers are available for programming L4 port ranges. After the 24 port range checkers are consumed, system will use port number and mask combination to program the rules in the access-list. Doing so would result in a given access-list consuming more TCAM entries in this mode than it would in the default mode. However, as the access-list is shared among multiple VLAN interfaces the effective TCAM usage can be lower.

  This affects ip access-list, ip standard access-list, and ipv6 access-list. ipv6 standard access-list
  is not impacted. (105498)
  Series Affected: DCS-7010, DCS-7050X and DCS-7050X2 Series

- "hardware access-list resource sharing vlan in" configuration will not share resources required by IPv6 standard access-lists. (105502)
  Series Affected: DCS-7010 and DCS-7050X Series

- 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)

- A large ACL and PBR policy configuration that consumes all TCAM resources may not fit after doing a config-replace, ACL agent restart or switch reload. After performing a config-replace, ACL agent restart or a switch reload, it

is possible that only the ACL or the PBR policy will be able to fit. (105667)
Series Affected: DCS-7010 and DCS-7050X Series

- 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)

- Configuring "hardware access-list resource sharing vlan in" in a config replace session or in config session is not supported (114291)
  Series Affected: DCS-7010 and DCS-7050X Series

- 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)

- Toggling "hardware access-list resource sharing vlan in" configuration could cause ACLs that were already programmed on VLAN interfaces to not get programmed. This could traffic loss on the VLANs where the ACLs were originally applied. (115348)
  Series Affected: DCS-7010 and DCS-7050X Series

- 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-7050X2, 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)
- If hardware tables for VLAN translation overflows, the entries need to be un-configured and re-configured after the overflow condition is removed due to appropriate configuration changes. (157196)
- 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, DCS-7050X2 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-7050X, 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, DCS-7050X 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)

- 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_HWHER_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)
- Some ports (esp. Et73-Et96 and Et103-Et104) might blackhole traffic or fail to linkup. (358362)
  SKUs Affected: DCS-7050TX2-128-F and DCS-7050TX2-128-R

- 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

- 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-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7250X, 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

- With VXLAN dataplane learning disabled, locally learned MAC addresses will not age out if VXLAN encapsulated packets from remote VTEPs are received with inner packet's source mac as the locally learned MAC. The locally learned MAC entry needs to be cleared manually by issuing 'clear mac address-table dynamic vlan <VLAN> address <ADDRESS>'. (460607)

- 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-7050X, 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)

- In VXLAN routing deploy, when VRRP is configured for an active/standby gateway, the VXLAN encapsulated packets destined to the VRRP MAC does not get decapsulated and forwarded. (557430)
  Series Affected: DCS-7050QX, DCS-7050SX, DCS-7050TX, DCS-7050X, DCS-7060, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7300X Series

- 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
- End EOS support for DCS-7050X, DCS-7050X2, DCS-7250X, and DCS-7300X platforms. (691787)
- 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-7060DX4, 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
- 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)

DCS-7010 Series

- IEEE 1588 PTP transparent clock does not decrement TTL for PTP routed packets (55078)

7300X3, 7320X, DCS-7060, DCS-7060X, DCS-7260X and DCS-7300X3 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 multple 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)

- 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)

- 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)

- "show version compatibility" CLI command reports image incompatibility when image is compatible. (547566)
  SKUs Affected: DCS-7060DX4-32-F and DCS-7060PX4-32-F

- 3m QSFP-DD passive copper cables are not supported in ports 1-4 and 29-32 (639613)
  SKUs Affected: DCS-7060DX4-32

DCS-7050X, DCS-7250X and DCS-7300X Series

- Packets sourced from the CPU to be VXLAN encapsulated will not have the right DSCP marking in the outer header based on the global QOS map configuration. (139248)
- Tcam entries used by IPv6 standard access list is not shared across interfaces (159827)
- In MLAG VTEP configuration, the switch incorrectly drops VXLAN encapsulated packets that are destined to the VTEP IP address with the destination mac as peer switch's bridge MAC address. The workaround is to configure VARP in such scenarios. (173716)
- Multicast traffic destined to boundary or prefix mroutes configured may be forwarded out of order if fabric priority flow control is enabled for the given priority. (185981)
- If an SVI is used as a core interface to reach remote Vteps, any L3 EVPN traffic using that SVI will be dropped if the underlay mac becomes unlearnt. A workaround is to configure MAC timeout to a higher value than ARP timeout to ensure macs don't get aged out. (192419)
- During "reload hitless", ECMP traffic flows may get hashed to a different ECMP path than before the reload. (218700)
  Series Affected: DCS-7050X Series
- Direct flow on T2+ has the following limitations
  1. Only IP packets can have their source mac , destination mac and vlan changed.
  2. When a packet has its smac, dmac and vlan changed then it can only be forwarded to one interface.
  3. When a packet has its vlan changed it can only be forwarded untagged on access ports.
  4. An output interface must be specified whenever a packet's smac,dmac and/or vlan is changed. (251477)
- SFP+ and QSFP+ interfaces with a MAC loopback configuration (traffic-loopback or routing recirculation-interface) may erroneously allow the peer device to link up and blackhole traffic. To work around the issue, either unplug the cable or administratively shutdown the interface on the link partner. (494452)

DCS-7050X2 Series

- When IEEE 1588/PTP is configured on interfaces running at 100Mb/s speeds, the master offset and mean path delay calculations are inaccurate due to a limitation in the internal PHY (122816)

- Switch tries to perform VXLAN decapsulation of a packet if the destination IP of the packet matches the local Vtep IP address and the destination UDP port matches the configured VXLAN Port, even though the L2 Destination Mac address does not match any of its local Mac addresses. This can result in the packet getting dropped. This issue is more likely to be encountered when the switch happens to be an MLAG peer and the MLAG peers have the same VTEP IP configured. The MLAG switch may drop the received VXLAN encapsulated packets addressed to its MLAG peer's Mac address when the destination IP matches the shared VTEP IP configured on the MLAG Peers. (145099)
  Series Affected: DCS-7050SX2 Series

- Egress VLAN translation will not work on routed multicast packets when the VLAN to be translated is part of the outgoing interface list. (151093)

- VLAN translation is not supported on vxlan core ports (154309)
  Series Affected: DCS-7050X2 Series

- Due to an increase in the number of default entries programmed in the TCAM, 'Out of TCAM entries', might be seen while configuring ACLs. This happens due to the TCAM hardware resources being full. The workaround is to unconfigure any unused feature that requires TCAM resources. (182178)

- L2 protocol forwarding is not supported (470496)

CCS-710P Series

- Triple-wide TCAMs are not supported. (603246)
  Series Affected: CCS-710P Series

- Generation of PFC frames is not supported. (699868)

CCS-720XP 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

- Generation of PFC frames is not supported. (699866)

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

7300X3 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

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 BASE-T transceivers are not supported on DSFP ports. (702025)
  SKUs Affected: 1000BASE-T, SFP-10G-MRA-T and SFP-10G-T

7300X3, CCS-710P, CCS-720XP, 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)
- Cut-through mode is not supported on SFP+ interfaces. (252731)
- Directflow configuration with "set vlan", but no "action set output interface" configured is not supported. (345852)
- The following DirectFlow configurations are not supported: match input interface hardware-loopback
  action output interface <> (more than one)
  action output interface hardware-loopback
  action output flood|all
  action set cos (355927)
- 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)
- "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)
- 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)

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.