Arista Networks EOS 4.23.14M Release Notes

Version: 1.0

27 May 2023

Table of Contents

EOS 4.23.14M Release Highlights

- New Platforms and Hardware

SNMP MIBs

SNMP Traps

Supported Hardware

Reference to MLAG ISSU Upgrade/Downgrade Table

Resolved Caveats in 4.23.14M

- General

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

Known Software Caveats in 4.23.14M

- Accelerated System Update
- BGP EVPN L2/L3 VXLAN/MPLS
- vEOS Router / CloudEOS
- CVP
  - Network Provisioning
  - Telemetry
- CVX
- General
- IPSEC
- Layer 1
- Layer 2
- Layer 3
  - Multi-agent
  - Rib

MLAG

MPLS

Multicast

NAT

SwitchApp

vEOS Router / CloudEOS

DCS-7150 Series

DCS-7170 Series

- DCS-7170 Series

CCS-720XP Series

Chassis Series

DCS-7160, DCS-7170, DCS-7280R and DCS-7280R2 Series

7368, DCS-7300X, DCS-7500R and DCS-7800 Series

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

- DCS-7500E Series
- DCS-7280R and DCS-7500R Series
- DCS-7280R2 and DCS-7500R2 Series
- 7500R3, DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series
- DCS-7020 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, CCS-720XP, DCS-7010, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7250X, DCS-7260X, DCS-7260X3 and DCS-7300X Series

- DCS-7010 Series
- 7300X3, 7320X, DCS-7060X and DCS-7260X Series

- DCS-7260X3 Series
- 7368 and DCS-7060X4 Series
- DCS-7050X, DCS-7250X and DCS-7300X Series
- DCS-7050X2 Series
- CCS-720XP Series
- DCS-7050X3 Series
- 7300X3 and DCS-7050X3 Series

Transceiver Series

7300X3, CCS-720XP and DCS-7050X3 Series

DCS-7160 Series

Limitations and Restrictions in 4.23.14M

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

Containerized EOS

vEOS Router / CloudEOS

CVP

- Telemetry

CVX

General

Layer 1

Layer 2

Layer 3

- Multi-agent
- Rib

MLAG

MPLS

Multicast

SwitchApp

DCS-7150 Series

DCS-7170 Series

- DCS-7170 Series

CCS-720XP Series

DCS-7160, DCS-7170, DCS-7280R and DCS-7280R2 Series

7368, DCS-7300X, DCS-7500R and DCS-7800 Series

CCS-720XP Series

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

- DCS-7500E Series
- DCS-7280R and DCS-7500R Series
- DCS-7280R2 and DCS-7500R2 Series
- 7500R3, DCS-7280R3, DCS-7500R3 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, CCS-720XP, DCS-7010, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7250X, DCS-7260X, DCS-7260X3 and DCS-7300X Series

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

Transceiver Series

7300X3, CCS-720XP and DCS-7050X3 Series

Conditions and Impacts

EOS 4.23.14M Release Highlights

New Platforms and Hardware

- None

SNMP MIBs

- SNMPv2, SNMPv3
- RFC 3635 EtherLike-MIB
    - obsoletes RFCs 1650, 2358, 2665
- RFC 3418 SNMPv2-MIB
    - obsoletes RFCs 1450, 1907
- RFC 2863 IF-MIB
    - obsoletes RFCs 1229, 1573, 2233
- RFC 2864 IF-INVERTED-STACK-MIB
- RFC 4292 IP-FORWARD-MIB
    - obsoletes RFC 1354
- ARISTA-SW-IP-FORWARD-MIB
    - IPv4 only
- RFC 4363 Q-BRIDGE-MIB
- RFC 4188 BRIDGE-MIB
- ARISTA-BRIDGE-EXT-MIB
- RFC 4113 UDP-MIB
    - obsoletes RFC 1213
- RFC 4022 TCP-MIB
    - obsoletes RFC 1213
- RFC 4293 IP-MIB
    - obsoletes RFC 1213
- HOST-RESOURCES-MIB
- LLDP-MIB
- LLDP-EXT-DOT1-MIB
- LLDP-EXT-DOT3-MIB
- ENTITY-MIB
- ENTITY-SENSOR-MIB
- ENTITY-STATE-MIB
- RMON-MIB
    - rmonEtherStatsGroup
- RMON2-MIB
    - rmon1EthernetEnhancementGroup
- HC-RMON-MIB
    - etherStatsHighCapacityGroup
- RFC 3636 MAU-MIB
    - ifMauDefaultType and ifMauAutoNegStatus are writeable

- SNMP-TLS-TM-MIB
    - RFC 6353
- SNMP-TSM-MIB
    - RFC 5591
- ARISTA-ACL-MIB
- ARISTA-SNMP-TRANSPORTS-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-SMI-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-PRODUCTS-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-QUEUE-MIB
    - Excluding 7150
- RFC 4273 BGP4-MIB
- RFC 4750 OSPF-MIB
- ARISTA-CONFIG-COPY-MIB
- ARISTA-CONFIG-MAN-MIB
- ARISTA-REDUNDANCY-MIB
    - 7500R, 7500E, 7300X, 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- MSDP-MIB
- PIM-MIB
- IGMP-MIB
- IPMROUTE-STD-MIB
- VRRPV2-MIB
- ARISTA-QOS-MIB
- ARISTA-ENTITY-SENSOR-MIB
- ARISTA-BGP4V2-MIB
- ARISTA-VRF-MIB
- ARISTA-DAEMON-MIB
- ARISTA-IF-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-MAU-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-PFC-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-FIB-STATS-MIB
- ARISTA-GENERAL-MIB
- ARISTA-HARDWARE-UTILIZATION-MIB
- ARISTA-NEXTHOP-GROUP-MIB
- ARISTA-SW-IP-FORWARDING-MIB
- ARISTA-TEST-MIB
- ARISTA-XCVR-DWDM-MIB
- ARISTA-XGS-MIB
- IEEE8021-PFC-MIB
- IEEE8023-LAG-MIB
- MPLS-LDP-STD-MIB
- NOTIFICATION-LOG-MIB

- OSPF-TRAP-MIB
- OSPFV3-MIB
- PTPBASE-MIB

SNMP Traps

- RFC 2863 IF-MIB
    - linkUp, linkDown
- LLDP-MIB
    - lldpRemTablesChange
- RFC 3418 SNMPv2-MIB
    - coldStart
- NET-SNMP-AGENT-MIB
    - nsNotifyRestart
- ENTITY-MIB
    - entConfigChange
- ENTITY-STATE-MIB
    - entStateOperEnabled, entStateOperDisabled
- OSPF-MIB
    - ospfNbrStateChange, ospfIfConfigError, ospfIfAuthFailure, ospfIfStateChange
- BGP4-MIB
    - bgpEstablished, bgpBackwardTransition
- ARISTA-REDUNDANCY-MIB
    - aristaRedundancySwitchOverNotif only for: 7500R, 7500E, 7300X, 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-CONFIG-MAN-MIB
    - aristaConfigManEvent
- SNMPv2-MIB
    - authenticationFailure
- VRRPv2-MIB
    - vrrpTrapNewMaster
- MPLS-LDP-STD-MIB
    - mplsLdpSessionUp, mplsLdpSessionDown
- MSDP-MIB
    - msdpEstablished, msdpBackwardTransition
- PIM-MIB
    - pimNeighborLoss
- ARISTA-BRIDGE-EXT-MIB
    - aristaMacMove, aristaMacLearn, aristaMacAge

Supported Hardware

Fixed System

* CCS-720XP-24Y6 * DCS-7060SX2-48YC6-F * DCS-7280SE-64-R

* 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
* 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-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-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-32
* DCS-7050QX-32-F
* DCS-7050QX-32-R
* DCS-7060SX2-48YC6-R
* DCS-7150S-24
* DCS-7150S-24-CL
* DCS-7150S-24-CL-F
* DCS-7150S-24-CL-R
* DCS-7150S-24-CL-SSD-F
* DCS-7150S-24-CL-SSD-R
* DCS-7150S-24-F
* DCS-7150S-24-R
* DCS-7150S-52-CL
* DCS-7150S-52-CL-F
* DCS-7150S-52-CL-R
* DCS-7150S-52-CL-SSD-F
* DCS-7150S-52-CL-SSD-R
* DCS-7150S-64-CL
* DCS-7150S-64-CL-F
* DCS-7150S-64-CL-R
* DCS-7150S-64-CL-SSD-F
* DCS-7150S-64-CL-SSD-R
* DCS-7150SC-24-CLD
* DCS-7150SC-24-CLD-F
* DCS-7150SC-24-CLD-R
* DCS-7150SC-64-CLD
* DCS-7150SC-64-CLD-F
* DCS-7150SC-64-CLD-R
* DCS-7160-32CQ
* DCS-7160-32CQ-F
* DCS-7160-32CQ-R
* DCS-7160-48TC6
* DCS-7160-48TC6-F
* DCS-7160-48TC6-R
* DCS-7160-48YC6
* DCS-7160-48YC6-F
* DCS-7160-48YC6-R
* 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-7280SE-68
* DCS-7280SE-68-F
* DCS-7280SE-68-R
* DCS-7280SE-72
* DCS-7280SE-72-F
* DCS-7280SE-72-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-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-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-7050QX-32-SSD-F
* DCS-7050QX-32-SSD-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-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-7050SX3-48YC12
* DCS-7050SX3-48YC12-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-7260QX-64
* DCS-7260QX-64-F
* 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-7280TRA-48C6-F
* DCS-7280TRA-48C6-M-F
* DCS-7280TRA-48C6-M-R
* 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-7011H
* FAN-7011H-F
* FAN-7011H-R
* FAN-7011M
* FAN-7011M-F
* FAN-7011M-R
* FAN-7012H-RED
* 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
* PWR-1611-DC-RED

* DCS-7050SX3-48YC8
* DCS-7050SX3-48YC8-F
* DCS-7050SX3-48YC8-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-7050TX2-128-F
* DCS-7050TX2-128-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-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
* DCS-7280CR3K-96-F
* DCS-7280CR3MK-32P4
* DCS-7280CR3MK-32P4-F
* DCS-7280CR3MK-32P4-R
* DCS-7280DR3-24
* DCS-7280DR3-24-F
* DCS-7280DR3-24-M-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
* PWR-1900-DC
* PWR-1900-DC-F
* PWR-1900-DC-R
* PWR-1900AC
* PWR-1900AC-F
* PWR-1900AC-R
* PWR-2411-AC-BLUE
* PWR-2411-AC-RED
* PWR-2411-DC-BLUE
* PWR-2411-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-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
* PWR-747AC-F
* PWR-747AC-R
* PWR-750AC
* PWR-750AC-F

* 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-7280QRA-C36S
* DCS-7280QRA-C36S-F
* DCS-7280QRA-C36S-M-F
* DCS-7280QRA-C36S-M-R
* DCS-7280QRA-C36S-R
* DCS-7280SE-64
* DCS-7280SE-64-F
* PWR-750AC-R
* FAN-7000H
* FAN-7000H-F
* FAN-7002
* FAN-7002-F
* FAN-7002-R
* FAN-7002H-F

Modular

* FAN-7000H
* FAN-7000H-F
* FAN-7002
* FAN-7002-F
* FAN-7002-R
* FAN-7002H-F
* 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
* 7316X-FM
* 7320X-32C-LC
* 7324X-FM
* 7328X-FM
* 7368
* 7368-16C
* 7368-SUP
* 7368-SUP-D
* 7368X4-SC
* 7500E-12CM-LC
* 7500E-12CQ-LC
* 7500E-36Q-LC
* 7500E-48S-LC
* 7500E-48T-LC
* 7500E-6C2-LC
* 7500E-6CFPX-LC
* 7500E-72S-LC
* 7500R-36CQ-LC
* 7500R-36Q-LC
* 7500R-48S2CQ-LC
* 7500R-8CFPX-LC
* 7500R2-18CQ-LC
* 7500R2-36CQ-LC
* 7500R2A-36CQ-LC
* 7500R2AK-36CQ-LC
* 7500R2AK-48YCQ-LC
* 7500R2AM-36CQ-LC
* 7500R2M-36CQ-LC
* 7500R3-24P-LC
* 7500R3-36CQ-LC
* 7500R3K-36CQ-LC
* 7500RM-36CQ-LC
* 7504E-FM
* 7504R-FM
* 7504R3-FM
* 7508E-FM
* 7508R-FM
* 7508R3-FM
* 7512R-FM
* 7516R-FM
* 7800R3-36P-LC
* 7800R3-48CQ-LC
* 7800R3-48CQM-LC
* 7800R3K-48CQ-LC
* 7808R3-FM
* DCS-7304
* DCS-7308
* DCS-7316
* DCS-7500-SUP2
* DCS-7500-SUP2-D
* DCS-7500E-SUP
* DCS-7500E-SUP-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-7808-CH
* PWR-2700-DC
* PWR-2700-DC-F
* PWR-2700-DC-R
* PWR-2900AC
* 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
* 100GBASE-AOC
* 100GBASE-CR4
* 100GBASE-CWDM4
* 100GBASE-DR
* 100GBASE-ERL4
* 100GBASE-FR
* 100GBASE-LR
* 100GBASE-LR4
* 100GBASE-LRL4
* 100GBASE-PSM4
* 100GBASE-SR4
* 100GBASE-SRBD
* 100GBASE-SWDM4
* 100GBASE-XSR4
* 100GE-DWDM2
* 10GBASE-AOC
* 10GBASE-CR
* 10GBASE-DWDM
* 10GBASE-DWDM-T
* 10GBASE-DWDM-ZR
* 10GBASE-ER
* 10GBASE-ERBD
* 10GBASE-ERBD-D
* 10GBASE-ERBD-U
* 10GBASE-ERLBD
* 10GBASE-ERLBD-D
* 10GBASE-ERLBD-U
* 10GBASE-LR
* 10GBASE-LRL
* 10GBASE-SR
* 10GBASE-SRL
* 10GBASE-ZR
* 25GBASE-AOC
* 25GBASE-CR
* 25GBASE-LR
* 25GBASE-SR
* 40GBASE-AOC
* 40GBASE-CR4
* 40GBASE-ER4
* 40GBASE-LR4
* 40GBASE-LRL4
* 40GBASE-PLR4
* 40GBASE-PLRL4
* 40GBASE-SR4
* 40GBASE-SRBD
* 40GBASE-UNIV
* 40GBASE-XSR4
* 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-D-400G
* CAB-O-2Q-200G
* CAB-O-2Q-400G
* CAB-O-4Q-200G
* CAB-O-4Q-400G
* CAB-O-8S-200G
* CAB-O-O-400G
* CAB-Q-2Q-100G
* CFP2-100G-DWDM-DCO
* CFP2-100GBASE-ER4
* CFP2-100GBASE-LR4
* CFP2-100GBASE-XSR10
* CFPX-100G-DWDM
* CFPX-200G-DWDM
* OSFP-400G-2FR4
* OSFP-400G-DR4
* OSFP-400G-LR8
* OSFP-400G-SR8
* QDD-400G-LR8
* SFP-10G-T

Reference to MLAG ISSU Upgrade/Downgrade Table

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

Resolved Caveats in 4.23.14M

General

- 'ptp monitor threshold offset-from-master drop' must be configured after PTP has synced from its master. If configured prior, the switch will drop all PTP packets indefinitely until the command is unconfigured. (543981)

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

- 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-7060SX2, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260CX, DCS-7260CX3, DCS-7260QX and DCS-7260X Series

- Performing SSU upgrade may result in black holing of VXLAN VLAN traffic.

  The workaround is to reload the system or restart Strata forwarding agent. (704372)

Known Software Caveats in 4.23.14M

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

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 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 a EVPN MLAG IRB deployment, reloading one of the MLAG peers may trigger the loss of some dynamic VLANs allocated for certain VRF to VNI mappings.

  The workaround is to restart the VXLAN agent. (548097)
- In an EVPN deployment, ARP entries learned from MAC/IP advertisement routes may remain stale after removing the associate MAC-VRF or changing the BGP AS number.

  The workaround is to restart the BGP agent. (578375)
- 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 a Multi-VTEP EVPN VXLAN MLAG IRB deployment, reloading both peers may result in the lost of dynamic VLANs on one or both peers required by VRF to VNI configurations.

  A workaround is to restart the VXLAN agent. (586826)
- In an EVPN multi-homing deployment, with multiple interfaces in an Ethernet Segment and a large number of allowed VLANs, the BGP agent may unexpectedly due to a configuration change. The configuration changes that can trigger the restart include but are not limited to 1) updates to the Ethernet Segment

Identifier; 2) administratively shutting down or enabling a member interface.

As a workaround, if a port has a large number of allowed VLANs its configuration should be updated separate from that of other ports. (609100)

- In a EVPN MH A-A deployment, traffic to and from the CE behind the MH peer maybe blackholed when one of the MH peers enters BGP Maintenance Mode.

  A work around is to disable first disable the CE facing link on the MH peer that is enter BGP Maintenance Mode prior to entering BGP Maintenance Mode. (629109)

- In an EVPN VXLAN All-Active Multihoming (MH) deployment, pings between any overlay physical IPs configured on local SVIs on the MH peers will fail to resolve the ARP entry for the overlay physical IPs and therefore fail.

  The workaround is to configure static ARP entries for the overlay SVI physical IP addresses on the relevant MH peers.

  If the multi-homed VTEPs are acting as an any-cast centralized gateway for one or more L2-VTEPs and are configured with a VIRTUAL VTEP IP (VVTEP IP), then the ARP resolution for a host behind a L2-VTEP initiated from one of the MH VTEPs may fail to synchronize the ARP reply on all the any-cast gateway MH VTEPs.

  The workaround would be to use 'ip virtual-router' instead of 'ip address virtual' to configure the any-cast gateway IPs on the MH VTEPs. (639721)

- In an EVPN VXLAN deployment, clearing a blacklisted MAC after the associated MAC-VRF has been removed may restart the BGP agent. Possible triggers for clearing blacklisted MAC entries include using the CLI command "clear bgp evpn host-flap", and deleting the associated SVI (663048)

- 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)

- In an EVPN VXLAN IRB deployment, a remote ARP (or IPv6 neighbor) entry learned over EVPN might not be installed.
  The issue can be identified when the entry appears up under "show arp remote" but not under "show arp agent ipv4(or ipv6) input evpn-vxlan" , and the ARP (or IPv6 neighbor) entry is missing from the switch ARP (or IPv6 neighbor) cache.

  The workaround is to install a CLI static ARP entry or alternately, the VxlanSwFwd agent could be restarted. (792566)

- Customers may sometimes observe that a few MAC only routes are not advertised over EVPN despite these MAC address being locally learnt. It is a rare event that may hit if the system is sufficiently loaded while trying to distribute the routes over EVPN.

  Learning a new MAC entry in the VLAN should fix the problem and trigger the missing routes to be advertised over EVPN. Alternatively BGP agent could be terminated to fix this issue. (794205)

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)
- Reusing a VNI previously used by VRF-A on VRF-B will affect traffic on VRF-A.

  The workaround is to restart the software forwarding engine (Sfe) agent. (523066)
- CloudEOS VM running on-prem with NICs x520/x540/82559, some of the interfaces may stay in down state due to a race condition in the ixgbe driver. This can be triggered upon VM image upgrade/downgrade or VM reload/ reboot or SFE agent restart. This problem can be seen on any NIC that is controlled by ixgbe driver. (523705)
- With DPS enabled, there is a small chance that SFE agent may restart unexpectedly after a DPS path flap event. (527588)
- Changing the mapping of Vrf to VNI (VXLAN Network Identifier) twice will cause the software forwarding engine (Sfe) agent to restart. There will be a brief moment of traffic loss which will recover eventually. (539064)
- When the VXLAN VNI mapping for a VRF is modified, paths towards one or more DPS peers may get updated incorrectly, resulting in loss of traffic towards such peers.

  Removing and re-adding the VNI mapping may fix this problem. (549493)
- 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)

- A bug in Google Cloud KVM implementation causes Live Migration performed during a
maintenance event to reload the vEOS router.

  Until a new KVM instance is deployed the only workaround is to modify /mnt/flash/kernel-params to disable hyperthreading using the nosmt parameter. Please contact Arista for detailed steps to apply the workaround. (579590)

- In certain CloudEOS deployments the VM may be oversubscribed and thus subject to processing delays while sending packets to the CPU. This primarily affects BFD control packets and result in the BFD detecting a link flap.

  One workaround is to increase the interval, min-rx and multiplier to larger values for each affected interface. For example, use the command: "bfd interval 500 min-rx 500 multiplier 5" to increase the timeout to 2.5 seconds in the interface sub-mode. (581475)

- When SFE agent restarts or device reboots, some DPS paths may incorrectly set the loss rate to 50%. This will impact the load balancing of traffic by not utilizing the paths with 50% loss rate to their real capacity. The same issue may also be seen in the scenario where traffic is stopped while there is already significant amount of traffic loss over DPS paths. Auto correction of this issue may take a while to recover the loss rate back to 0 or normal value.

  It can be worked around by flapping the corresponding WAN interfaces on the DPS peers of the problematic paths. (618258)

- 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)

- 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)

- Kernel to network traffic may be dropped in the CloudEOS platform due to scheduling delays that cause the kernel to send more than 16 packets before the SFE agent can refill the kernel buffers. This may be diagnosed using the ethtool -S <intf> command and noting if the tx_queue_0_descdropped value is non-zero. (736836)

- Security Advisory 85 for CVE-2023-24545.
  See https://www.arista.com/en/support/advisories-notices/security-advisory/17240-security-advisory-0085 for more details. (743423)

CVP

Network Provisioning

- OpenConfig configurations involving route targets may be incorrectly ignored. (528374)

- 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)

Telemetry

- Configuring an IPv4-mapped-IPv6 static route may result in OpenConfig/Octa agent continuous restart.
  Workaround is to unconfigure the IPv4-mapped-IPv6 static route. (796314)

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 a group on NSX-T containing more than 1000 hosts, policy rules for the group on CVX will only generate ACLs for the first 1000 hosts in that group. (459644)

- The Policy Control Service (PCS) fails to apply any security policy configured on NSX-T on a switch interface after the switch configuration is modified to remove all tags from an interface followed by re-adding tags, using the configuration replace feature. (463958)

- When 'client state preserve' feature is configured on CVX, state of the primary switch of an MLAG VTEP will not be preserved upon disconnection. If the primary switch did not reboot and disconnected from CVX due to network faults, it may cause disruption of VXLAN control plane leading to VXLAN traffic loss.

  As a workaround, reboot the MLAG primary switch. (533925)

- In a Macro-Segmentation Service (MSS) deployment with Panorama, uncommitted changes to the list of managed devices may be used by MSS to program rules on the switches. (536412)

- 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)

- In a CVX HA deployment, the NetworkTopologyAggregator agent may restart during a cluster failover.

  A workaround if the NetworkTopologyAggregator is stuck in a restart loop, is to trigger a failover on the CVX node so a new master is elected. (621677)

- In a VXLAN routing deployment using VCS to distribute MAC and ARP bindings with MSS enabled, if there is an application using a virtual IP address

whereby the MAC address bound to the virtual IP address changes as it migrates to behind a different VTEP, then the VTEP may fail to update the ARP entry for the virtual IP to the newer MAC address.

The workaround is to either disable MSS during the migration, or clear the ARP entries for the virtual IP on the VTEPs prior to migrating the virtual IP. (624581)

- The JsonApi agent's memory usage may increase over time, causing new OpenStack updates to not be processed after the agent exhausts available memory

  The issue may be temporarily resolved by restarting CVX or the JsonApi agent (648018)

- The ManagementActive agent can restart unexpectedly when a virtual IP address is configured on CVX. (652036)

- 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)

- The MCS agent might restart after multiple addSenders POST requests for the same senders that exist in the setup. (767011)
- In a VXLAN deployment using VXLAN Controller Service (VCS) to distribute MAC addresses, the MAC address of an orphan host behind a secondary MLAG switch may not get published to VCS. (770762)

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)

- OpenFlow Flows may not be programmed in the right priority order after changing OpenFlow Table Profile config to "auto" mode. (118767)
  SKUs Affected: DCS-7260CX3-64, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F and DCS-7260CX3-64E-R
- 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-32-F, DCS-7050QX-32-R, DCS-7050QX-32-SSD-F, DCS-7050QX-32-SSD-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-7150S-52-CL-F, DCS-7150S-52-CL-R, DCS-7150S-52-CL-SSD-F, DCS-7150S-52-CL-SSD-R, DCS-7160-32CQ-F, DCS-7160-32CQ-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-7280SE-64-F, DCS-7280SE-64-R, DCS-7280SE-68-F, DCS-7280SE-68-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)
- The ZeroTouch agent does not properly handle 301 HTTP return codes when downloading the bootstrap script. This can lead to the device leaving ZTP mode with a default startup-config. (342098)
- 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-7280E, DCS-7280R, DCS-7280R2, DCS-7500E, 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)
- Mlag agent may crash if the system is being put in Maintenance Mode if interface profile has shutdown max-delay enabled. Workaround is as follows:
  1. Create a maintenance interface profile which does not have shutdown max-delay configured with profile name being lexicographically largest.

  ```
  profile interface zzz_no_shutdown_delay
     rate-monitoring threshold 4294967295
  ```

  2. Create an interface group constituting all the physical ethernet interfaces of peer-link and attach the above interface profile that does not have shutdown max-delay configured.
  3. Create a unit with this group and quiesce it. (380387)
- Routing protocols such as OSPF, IS-IS are disabled when VLAN or interface config is changed via gNMI or NETCONF. (401590)
- 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)
- Any configured QoS mappings for traffic class to COS and traffic class to DSCP bits are not applied on packets that are reflected via an Interface Reflector that is configured in MAC swap and direction out mode. The source and destination MAC addresses are swapped as expected, but the COS and DSCP bits of the reflected packets are copied unmodified from the original ingressing packet. (476507)
  Series Affected: DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series
- On system reload, the kernel may crash with the "PANIC: double fault" message. (493516)
- EosSdk agents using a certain combination of Sysdb mount paths may cause the Sysdb agent to restart on connecting to Sysdb. If the agent is configured and stored in startup-config this can cause Sysdb to get stuck in a restart loop. (496371)
- Adding and deleting Nexthop groups via EosSdk APIs, may increase process memory usage over a period of time. (498368)
- 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)
- Value for mlag leaf might not be mapped in certain circumstances. (518968)
- The "arp reply relay" feature does not support a VXLAN domain with L2 VTEPs. (540369)
- Modular switch may restart in case of Fabric or Linecard failure. (554503)
- In certain circumstances, when a process takes large amount of memory, it may trigger a kernel panic. (556742)

- Adding or modifying Route-map via OpenConfig may not result in applying the new Route-map in hardware. (575325)
- 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)
- A gNMI subscriber may not receive all delete messages to an eos_native / Smash path that has had all its elements deleted.

  The workaround is for the gNMI subscriber to resynchronize by starting a new Subscribe. (609497)
- 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)
- OpenConfig agent can unexpectedly restart when subject to a high scale of subscriptions and system updates. (615715)
- "reload peer" configured with a delay such as "reload peer in 1 force" results in both supervisors being reloaded after the specified delay. (621572)
- A constant very high rate of events triggering an event handler can cause memory buildup leading to an out-of-memory situation in the system

  Workaround: use the "repeat" event-handler configuration command to limit the rate of event trigger (624314)
- When a port is PTP enabled and `ptp vlan` is configured, a configure replace or configure session operation that results in no modification to the PTP configuration may cause PTP churn on the first instance.

  A workaround would be to perform a configure replace with the exact same running configuration to induce PTP churn on the affected port once after performing any of the following:
  - Reloading the switch.
  - Configuring additional allowed switchport VLANs (e.g. `switchport trunk allowed vlan`) on a PTP port with `ptp vlan` configured.
  - Configuring new VLANs to be used with `ptp vlan`. (625345)
- Under periods of high system churn, Octa or OpenConfig agents may restart with a broken pipe panic. The Sysdb agent log will contain a message

"Closing connection to Octa ... (AttrLog buffer is full)". The agent should recover when the churn subsides. (627426)

- If ingress sFlow is applied to a mirroring source port which is also used as a nexthop interface for a GRE tunnel used as a destination in any mirror session, then ingress sFlow on that port could stop working.

  A workaround is to disable and then re-enable sFlow on that port; An alternative is to remove the port as a source from its monitor session and then add it back. (629718)
  Series Affected: 7500R3, DCS-7020, DCS-7280CR2, DCS-7280E, DCS-7280PR3, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280SR3, DCS-7280TR, DCS-7500E, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 Series

- 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)

- When deleting VRFs with IPv6 Multicast configurations a repeating kernel warning can appear:

  [161364.865111] unregister_netdevice: waiting for lo to become free. Usage count = 5. Process: kworker/u8:0:11534

  This is the result of memory not being freed, which can eventually result in a failure.
  The switch should be rebooted at your earliest convenience. (636026)

- An RPC request containing both OpenConfig model based data and CLI commands (mixed-schema) may result in applying only the OpenConfig model based data changes and ignore some CLI commands.

  To workaround this problem, split up the request in two separate RPC calls, with the CLI commands in one RPC and the OpenConfig model-based data in another. (638247)

- Under certain conditions, it is possible for the YANG tree representation of the system's VLAN translation config to differ from the actual config. If an RPC to modify or update the system configuration is executed while the system is in this state, existing VLAN translation configuration that's not part of the RPC request can be erroneously removed.

  The system can enter this inconsistent state immediately after reload or on change in port speed settings or on a forwarding agent restart or after a

line card insertion/ removal event.

To workaround this problem, restart the OpenConfig agent when the system enters this erroneous state (verified by performing a GET operation and comparing with the running configuration). (641536)

- When configuring interfaces in stateful DHCPv6 environments, ZeroTouch agent will select the longest matching prefix from the Neighbor Discovery router advertisement rather than assigning a default prefix-length of /64. Defaulting the prefix-length to /64 can lead to reachability problems resulting in the agent unable to communicate with the ZTP bootstrap-server. (645404)

- OpenConfig agent is unable to populate the tunnel counters at "interfaces/interface[name=Tunnel*]/state/counters". A workaround is to configure Octa agent and add the following CLI config:
```
management api models
   provider smash
      path interface/counter/tunnel/fastpath/current/counter
! (647110)
```

- 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)

- Octa may restart unexpectedly when a gNMI client frequently syncs eos_native paths with more than 100,000 elements. These frequent syncs could be caused by a SAMPLE subscription or many ON_CHANGE or TARGET_DEFINED subscriptions. (649675)

- 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, DCS-7010, DCS-7050CX3, DCS-7050QX, DCS-7050QX2, DCS-7050SX, DCS-7050SX2, DCS-7050TX, DCS-7050TX2, DCS-7050TX3,

DCS-7050X, DCS-7050X2, DCS-7050X3, 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)

- Subscribers to gNMI eos-native can see sensitive data. (664159)

- Subscribers to gNMI eos-native can see sensitive data. (664160)

- 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 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)

- 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)

- If the default MTU is changed to something other than 1500 with configuration like:
  interface defaults
  mtu 9000
  Then interfaces which have their MTU explicitly set to 1500 will not have that MTU included in the 'show running-config'. This can cause the setting to be lost on a switch reload. (713182)

- OpenConfig or Octa's memory use may permanently increase as linecards are removed and re-inserted.
  Work around the issue by restarting OpenConfig or Octa. (750505)
- OpenConfig/Octa agent might reject incoming SET requests if the YANG container '/interfaces/interface/arista-vxlan' exists for any interface other than Vxlan1. Enable toggle OCSkipYANGValidation as a workaround. (783226)
- OpenConfig SET transactions fails when there is a Tunnel Interface configured and it is assigned to non default VRF. A workaround is to remove Tunnel interface from the VRF. (788203)
- Removing an interface and reassigning its IP addresses to another interface is not possible via OpenConfig in a single transaction. A workaround is to split the transaction into two or use a mixed-mode (CLI + YANG tree) OpenConfig transaction (791299)
- Setting an IPv4-mapped-IPv6 prefix in an IPv6 prefix-list can trigger an error in OpenConfig which can cause that prefix to not be saved in the YANG tree. (795935)
- Disabling errdisable detection for a given cause using the `no errdisable detect cause` command has no effect: subsequent detection of the given cause will cause the interface to be errdisabled again. (802655)
- OpenConfig SET transactions fail when there is a class map configured with matching for other L2 parameters except matching VLAN. A workaround is to remove the corresponding class map or by adding a match VLAN config on the map. (807799)
- If a logical interface associated with a non default VRF is deleted and then recreated in default VRF, packets matching the deny log rule of egress IPV6 ACL applied to the interface are software forwarded instead of being logged. (807976)

IPSEC

- Under some circumstances IPsec rekey time may show up as 'N/A'. (561575)

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

- 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 and DCS-7060X4 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: 7500R3-24P-LC, 7800R3-36P-LC, DCS-7280CR3K-32P4-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F and DCS-7280PR3K-24-F

- The B52 agent may restart, resulting in all interfaces experiencing a temporary link down event. (606651)
  SKUs Affected: DCS-7280CR3MK-32P4-F and DCS-7280CR3MK-32P4-R

- 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-7060SX2 and DCS-7260CX Series

- On 100G ports with the CRT50216 PHY, large changes in the recovered clock from the received signal without an associated loss of signal may result in the port remaining link down. This most commonly occurs with the port connected to line system equipment that has not been configured to forward faults. (626446)
  SKUs Affected: 7368-16C, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3K-48CQ-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

- Port active state is not guaranteed to be preserved across a reboot or a forwarding agent restart when more than the maximum number of logical ports are configured. Ports that were previously active may become inactive and vice-versa.

  A workaround is to avoid configuring more than the maximum logical ports on a pipe. If the maximum logical ports must be exceeded, unexpected activeness changes can be avoided by preemptively restarting the forwarding agent using the CLI "platform trident reset". (682740)
  SKUs Affected: DCS-7050SX3-48YC8

- 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-7060SX2, DCS-7260CX, DCS-7260CX3 and DCS-7260QX Series

- Interfaces with BASE-T SFP transceivers and link-debounce configuration may experience a link down debounce period double than the configured period. The workaround is to configure "phy link detection aggressive" on affected interfaces (this may lead to an interface flap at the time of applying this configuration). (702920)

- 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. (706064)
  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. (706065)
  SKUs Affected: DCS-7050SX3-48YC8

- 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-7050TX3 and DCS-7260X3 Series

- "phy link training" in startup-config will not be present in running-config after startup and will not be applied during startup.

  A workaround is to configure "phy link training" manually, or reapply startup-config with "configure replace startup-config". (750212)

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)
- In a VXLAN routing environment using EVPN IRB, when there is a mix of L2 AND L3 VLANs, the ARP and IPv6 ND forwarding traffic can overload the CPU resulting in dropped ARP/ND packets.

  As a workaround, the following configuration adds an ARP/ND proxy deny rule for the default route and reduces the CPU overload, preventing ARP/ND packet drops.

  ```
  !
  router l2-vpn
     arp proxy prefix-list disable-default-route-v4
     nd proxy prefix-list disable-default-route-v6
  !
  ip prefix-list disable-default-route-v4
     seq 10 deny 0.0.0.0/0
  !
  ipv6 prefix-list disable-default-route-v6
     seq 10 deny ::0/0
  !
  ```

  The above config disables proxy reply for all ARP/ND hosts, and all ARP/ND traffic is forwarded in hardware. The remaining EVPN IRB features (other than proxy reply) continue to function as before. (599556)
- In VXLAN setup using either data plane learning and/or CVX, when anycast gateway is configured with a IPv6 address, then there may be multiple responses to the Network Solicitation for the anycast gateway IP address send from a host. The number of responses will be proportional to the number of anycast gateways in the flooding domain for the VLAN. (609827)
- 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)
- "mac address-table flushing interface vlan aggregation disabled" may not work with the LAG port. When some allowed VLAN gets removed from a LAG port, the whole LAG port may be flushed instead of the VLAN that was removed. (631846)

- The "mac address-table flushing interface vlan aggregation disabled" command does not persist after device reload.

  After reload, use "no mac address-table flushing interface vlan aggregation disabled" followed by "mac address-table flushing interface vlan aggregation disabled" to disable aggregation. (640490)

- In the scenario of a phone sending LLDP packets after a finite delay post it's successful Dot1x authentication, it's MAC address can get advertised in the older (non-phone) VLAN via EVPN. The workaround is to restart Dot1x agent. (645511)

- 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: DCS-7050CX3M-32S and DCS-7280CR3MK-32P4

- 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 an EVPN VXLAN deployment, receiving an IPv6 IMET route on a device that doesn't support VXLAN IPv6 underlay can cause its L3 forwarding agent to restart unexpectedly.

  Configuring a IPV6 VTEP address in the VLAN flood-set can also cause its L3 forwarding agent to restart unexpectedly. (732972)

Layer 3

- During reload or configuration changes, BFD session may be stuck in init state on one side. Restart Bfd agent to work around the issue. (195545)

- The command "no ip routing ipv6 interfaces vrf <vrfName>" can disable IPv4 addressless forwarding for interfaces in all configured VRF instances including the "default" VRF. (267335)
- The Ldp agent may restart if configured with a large number of MPLS ECMP routes with many next hops. (273767)
- 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)
- Removing LDP mappings in a large scale LDP configuration may result in LDP neighbor sessions flaps. (283972)
- 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)
- If a decap group with multiple IPv6 addresses are configured in either a config-session or config-replace, forwarding agent may restart unexpectedly, or the decap-group may not get programmed correctly. (421010)
- 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)
- When a IS-IS/BGP/OSPF peer is torn down on the receipt of the BFD session down notification on a link that is still physically up, there could be traffic loss on traffic destined to prefixes which have ECMP paths that also include the aforementioned link. (434576)

- 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)

- When a series of back-to-back configuration changes are applied to a BFD session, the session may not have the latest configured value. To work around the issue, un-configure and re-configure the BFD session. (458955)

- 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)

- When there're a lot of NAT translations synchronized through NAT Sync in a very short time period, NAT Sync may hang and stop synchronizing. Restarting Nat agent with "agent Nat terminate" in CLI Enable mode may help. Retrying the command several times may be needed. (467300)

- DHCP relay agent might exit unexpectedly if several VRFs are continuously deleted and re-created. (471169)

- With 'ipv6 dhcp relay install routes' configured on more than 100 SVIs, DHCP relay agent may exit unexpectedly when a IPv6 DHCP helper address is configured simultaneously on all SVIs. (471338)

- In situations in which a large number of routes are forwarded over the same dynamic tunnel interface, the Arp agent may exit unexpectedly if the forwarding information of these routes are changed. The Arp agent will restart normally afterwards. (472837)

- When adding a VRF during config replace or quickly deleting and re-adding a VRF, the internal state of the ARP agent may go out of sync and ARP entries may be missing from "show arp." Workaround is to delete the VRF and re-add it. (474469)

- Cspf agent may restart when TI-LFA protection is configured and the number of segments requiring protection is over 2K (481703)

- In EOS, host route computation is facilitated by the Arp agent. In situations with scaled neighbor entries, when a route is deleted or added, particularly when all or most of the host routes are covered by the same route, the Arp agent may restart. Typically, the Arp agent will run fine after the restart. (484901)

- 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)

- TI-LFA backup paths will be computed slowly when there are large number of node-segments/anycast-segments/adjacency-segments that are protected. (489097)
- An LSR which is the penultimate hop in an RSVP LSP may incorrectly report a "bandwidth unavailable" error if the tunnel destination IP is an interface IP that does not match the last hop in the ERO.

  A workaround is to use a loopback interface as tunnel destination or the interface IP that is present in the ERO. (493578)
- BFD sessions may continuously flap after a speed change on an interface. A workaround is to restart StrataBfd. (499185)
- When there are a large number of attached routes on an interface and the interface's VRF is changed with attached routes moving to new VRF quicker than the attached routes deleted from the old VRF mayb result in some of the attached routes not programmed in hardware in the new VRF (501058)
- 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)
- Protocol messages exceeding 1500 bytes are ignored, in particular BUNDLE and SREFRESH messages, which is the case for jumbo frames. The symptom can be dropped LSPs or long convergence times. (521363)
- An RSVP RESV message is immediately sent as a reply to a NACK object, which can lead to a packet storm if an RSVP neighbor incorrectly uses NACK objects as an error reporting mechanism.
  Configuring RSVP with `refresh method explicit` will prevent this issue but disable Refresh Overhead Reduction. (522041)
- 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)
- If neither Graceful Restart nor Hold Timer are negotiated for a peer, BMP will not send any Route Monitoring Messages for that particular peer. (539641)

- KernelFib agent may restart on interface flap, if there are a large number of route resolving through that interface (539757)

- When any form of dynamic NAT is enabled, the first packet for translation is sent to CPU and
  forwarded in software and after the translation is programmed in the forwarding ASIC,
  the subsequent packets are forwarded by the ASIC. If the NAT destination is reachable
  via ECMP paths, the path chosen by the first packet by software forwarding may not be
  the same as subsequent packets forwarded in hardware. This can potentially cause issues
  if the NAT destination is to load-balancer or an anycast address that maintains NAT flow
  state is maintained per receive port.

  The workaround for this case is to have a single ECMP path for the NAT destination. (543692)
    Series Affected: DCS-7170 Series

- RSVP tunnels might not get established or may use non-best paths when IS-IS system ID or OSPF router ID is changed for any node in the IGP network and the best path goes via that node. One workaround is to change the TE router ID also for the same node. (546904)

- KernelFib agent may restart, if EVPN config is toggled in quick succession. (557358)

- On scale setup during certain network events, ARP agent's internal state may go out of sync, leading to ARP agent repeatedly deleting and adding the same entry to the kernel. This may cause the CPU usage to spike to 100% usage.

  Workaround is to restart ARP agent. (559855)

- SrTePolicy agent may restart if the top label of the Segment List is resolved over LFIB entry with outgoing label as Explicit Null Label. (561061)

- When TE router ID of one of the two routers which have the same TE router ID is changed to be same as that of a third router, Cspf agent may sometimes restart (565678)

- RSVP tunnels might not get established or may use non-best paths when IS-IS system ID or OSPF router ID of any two neighbouring nodes in the IGP network is changed back-to-back and the best path goes via that node. (572080)

- KernelFib crashes upon restart when a third party route added through bash or an external agent, is quickly added and deleted before the interface pointed to by the route is UP (574133)
- If a route received from a BGP peer is leaked from one VRF to another using VrfLeak agent (https://eos.arista.com/eos-4-22-1f/eos-vrfleak-agent/) and the leaked route in the target VRF is again leaked to another VRF using route-target import/export configuration in BGP (https://eos.arista.com/eos-4-21-3f/inter-vrf-local-route-leaking/), the Bgp agent can restart unexpectedly.

  As a workaround, use these mechanism to leak distinct set of routes. If a route has to be leaked across multiple VRFs, use either VrfLeak agent or route-target import/export configuraiton for leaking the route to all the target VRFs. (577257)
- Rib agent or the multiagents (Isis/Ospf) can restart with a large prefix-list configuration (578037)
- 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)
- An IP configured on an interface in a non-default VRF may cause an RSVP LSP to not come up if the same IP appears in that LSP's path for a different router. (591762)
- KernelFib agent may restart when interface flaps lead to a large number of routes (entire routing table) being withdrawn (597557)
- The VxlanSwFwd agent may crash when processing a DAD packet from a host that is destined to a virtual link local address. (602176)
- When auto-cost reference-bandwidth is configured in OSPF and OSPF receives a speed change event with bandwidth value 0 for a port channel interface, OSPF may set the cost of the interface to the default cost (10). This could result in traffic loss since routes may be installed to the port channel that has no bandwidth. (606650)
- 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)
- When there are least two IS-IS adjacencies between the routers and two or more adjacencies have the same associated SRLG value configured. IS-IS agent may restart on one or more nodes when one of the adjacency goes down. (617367)

- PBR nexthop rules set without the 'recursive' may still become resolved recursively. (627178)
- Rib/IS-IS agent may restart after a link flap if IS-IS adjacency is formed over a GRE tunnel and the tunnel is reachable via both IS-IS Level-1 and Level-2 reachabilities (635006)
- When 'ip hardware fib next-hop sharing' is configured on a device, some tunnel adjacencies that should be shared will no longer be shared. This happens even if the
  routes' contents were not changed and can lead to resource overflow.

  If UCMP is not enabled, toggling between "ip hardware fib next-hop weight-deviation
  0.01" and "ip hardware fib next-hop weight-deviation 0.02" will alleviate the resource usage and cause the routes to be shared again. (643733)
- The layer3 forwarding agent may unexpectedly restart when the ECMP groups table is full and adjacency sharing is enabled. (645324)
- 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)
- In scenarios in which FEC in place update occurs while routes are being added or updated, ARP/ND entries matching the subnet of these routes may not be added or may be incorrectly removed.

  Workaround is to restart the Arp agent. (660491)
- After the Rib agent restarts, static ARP or ND entries may go missing. Workaround is to restart the Arp agent. (673292)
- 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)
- 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)

- 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)

- 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)

- 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)

- IS-IS agent may restart if BGP VPN routes are evaluated in IS-IS either because of redistribute bgp command or because of set-attached-bit command and if the VPN routes are undergoing churn (731103)

- 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)

- When a session expires between two LDP peers used for pseudowire (link flap, etc.), the pseudowire might stay stuck in a "No label" state. (734555)

- If an IS-IS interface between L1L2 router and and L2 router is configured with circuit type as level-1-2 then if IS-IS L2 router is gracefully restarted and if the scale of L1 database on the L1L2 router is high then it is possible that the L2 database of the restarting router will never be synced.
  Workaround is to configure the interface on the L1L2 node with IS-IS circuit-type level-2 (735890)

- When IS-IS has large scale of self originated LSPs, Graceful Restart might fail resulting in traffic drop. (737118)

- When an OSPF topology has a large number of nodes and a large number of paths to a destination in the topology, the OSPF agent may restart when running SPF. (747859)

- Extra sockets on UDP port 51025 bound by VxlanSwFwd may be observed, causing ND 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. (755127)

- Under rare circumstances occurring due to configuring a large number of static ARP entries while simultaneously running "show tech-support" or "show arp agent ipv4 input", some static ARP entries may not be displayed in "show running-config" and may not get deleted from the system even after their configuration is removed.

  A workaround is to configure the same ARP entry again and remove it a second time. (759617)

- IS-IS may not advertise TE attributes after traffic-engineering is globally disabled and enabled if TE Router ID is same as Router ID configured in router general configuration mode\nWorkaround is to unconfigure and configure TE Router ID again (770528)

- When there is only one TLV present in an IS-IS LSP fragment, any update to the TLV might cause a purge of the fragment. (776597)
- The Sflow agent will select a random port within the ephemeral port range, 32768 - 60999, to bind to. Since the VxlanSwFwd agent is statically bound to a port within that range, 51023, the Sflow agent might select that same port, causing the VxlanSwFwd agent to repeatedly restart until the Sflow agent chooses another port.

  Restart the Sflow agent. (780707)
- IS-IS agent may restart if a peer undergoing graceful start advertises hello packet that contains interface addresses different from the one that it advertised before undergoing graceful start (785394)
- VRF names with "." may result in repeated restart of Launcher agent while starting process with name like 'Rib-vrf-foo.bar'. (792987)
- The OSPF agent may restart when processing stale self-originated LSAs as a Graceful Restart Speaker. (794004)
- Router Solicitation (RS) packets with the source address that is not the unspecified address and no source link-layer address option, causes a Router Advertisement (RA) to be sent with the destination address as unicast address and the destination link-layer address as all-nodes multicast address. (802490)
- RSVP FRR node protection tunnels are optimized after every change in topology including any bandwidth change, which can cause CSPF agent to become very busy when there are continuous bandwidth changes in topology. The workaround is to increase CSPF throttle timer values using 'cspf delay' CLI. (804284)

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)
- When a large number of updates are backed up in the transmit socket buffers, it may sometimes lead to the BGP neighbors session getting stuck in Idle(PendingNotification) state when the session reset is attempted and BGP is unable to send a notification. Not being able to send out updates may lead to traffic misforwarding or drops. (466575)

- Sometimes BGP (multi-agent) would advertise an incorrect originator-ID when sending updates to iBGP clients when configured as route-reflector. The problem only manifests when different paths (could be paths belonging to the same prefix or different prefixes) have the exact same path attributes, including NextHop. The problem corrects itself if the paths are resent to the route-reflector from the source whose paths had the incorrect originatorId. This problem can manifest with add-path configuration as well when the additional paths have the same path attributes. (472025)
- 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)
- On a reboot of a modular chassis without linecards, a user configured (non-DEFAULT) VRF exclusively associated with management interfaces has no management connectivity (563022)
- The Bgp agent may restart unexpectedly if two threads simultaneously apply input policy and the result is a rejected path. (652370)
- On a switch with a BGP peer that is activated for multiple address families, the Bgp agent may restart when attempting to transmit end of rib (EOR) messages to the peer. (661143)
- 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)
- 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)
- 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)
- The switch might transiently drop packets when the ECMP membership changes. (720703)
- Graceful restart with dynamic BGP peers does not work when switch reload using 'fastboot' is performed. The restarting speaker may drop traffic from

peers until the peering sessions are re-established based on hold timer expiry. (765816)

- BGP agent will not forward unknown Tunnel Encapsulation Attribute Sub TLVs correctly for VXLAN, DPS, and IPsec TLVs. (799441)

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)

- When the Tunnel agent goes down, stale entries in the tunnel RIB may not be cleaned up upon restart of the agent. (586928)

- RIB agent leaks memory when the color extended community or NextHop attribute changes. With a high scale of BGP routes with color extended community, even without SRTE configured, the RIB agent may restart due to the leak which leads to running out of memory. (615108)

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)
- MACs learnt locally may get deleted if System unit enters and exit maintenance without MLAG peer getting reloaded. (635453)
- 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)

MPLS

- When the refresh overhead reduction extension is used, some RSVP messages may never be sent out (shown as "State: path-only" under "show mpls rsvp session").
  As a workaround, the extension can be turned off and on again. (441063)
- Excessive toggling of the 'mpls ip' or 'ip routing' commands may cause the LDP label range to be exhausted.

  Wait for the LDP agent to shutdown after disabling 'mpls ip' before reenabling it to ensure the LDP labels are released. (539317)
- Upon reload of the switch, neighbors may reject Rsvp messages with "invalid sequence number" if Rsvp message authentication is enabled and the clock had not synced (e.g. via NTP) when the Rsvp agent started.

  A workaround is to restart the Rsvp agent once the system clock has synced, for example with "shutdown" + "no shutdown" of the RSVP feature. (614022)
- When attempting to program an MPLS over GRE adjacency (excluding MPLS over GRE next-hop groups), the forwarding agent may restart. (638559)
- 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)
- Change ISIS SRGB range configuration may result in continuous forwarding agent crashes if any global segment index is out of the configured label

range. Revert the configuration and restart MPLS agent resolves the issue. (734651)

- If a malformed PATH message is received which specifies a bandwidth priority value larger than 7, the RsvpProtocol agent may restart unexpectedly. (767244)

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)

- Fast reload fails when multicast is configured on a non-default VRF that is not PIM enabled. The default VRF is always PIM enabled, but a non-default VRF is PIM enabled when it contains at least one PIM interface. Workaround is to enable a PIM interface on each non-default VRF. (398256)

- 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-7280E, DCS-7280PR3, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280SR3, DCS-7280TR, DCS-7500E, 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)

- When using Sfe based forwarding configuration, the BessMgr agent may restart unexpectedly when multicast packet size is greater than 1500 bytes (481192)

- If MLD Snooping is configured, IPv6 multicast packets may not be bridged as expected in a PIM SM or PIM SSM non DR VLAN. (493226)
  Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

- BessMgr agent could restart unexpectedly when resolving a large number of unresolved packets at once. (495863)

- Removing router pim bsr also remove router pim sparse mode configuration. Workaround is to just remove specific configuration under router pim bsr mode. (516188)
- Using "software-forwarding sfe" for software forwarding of multicast packets does not work on x86_64 version of EOS. (537616)
- On devices running multicast with large number of active VLANs one might observe increased CPU utilization and latency in receiving multicast stream when the state of VLAN changes from down to up. (558317)
- On devices running PIM, on system boot up, might suffer loss or slower re-convergence of multicast traffic flows. (570930)
- 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)
- Periodic PIM assert messages might be sent continuously when the number of winner
  asserts per interval takes more than 0.1 seconds. There is no workaround. (613310)
- 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)
- In instances when CLI configuration between different aspects of PIM configuration (example main vs clear config) for a given VRF are mismatched, PimBsr and PimReg agents may continuously restart.

  Workaround is to restart the system. (633679)
- 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 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)
- 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)
- IGMP Snooping L2 querier doesn't function properly after enabling/disabling the primary VLAN of private VLANs.
  Workaround: Remove and re-add querier command for the primary VLAN. (740132)
- In an MVPN setup, traffic at ingress PE might get dropped if multiple flaps occur on the source interface. (777298)
- The multicast stream may be delayed by 30 - 60 seconds when the only receiver for the stream is present in another PIM domain and multicast source discovery needs to happen between the sending PIM domain and the receiving PIM domain. (801468)
- When a Cisco router is the upstream router to a Arista router for Source and RP of a given group, then a traffic drop might be experienced for that group.
  The workaround is to avoid using a topology where upstream interface of (s,g) and (*,g) are the same when upstream router is a Cisco switch. (810925)

NAT

- Address-Only mode within the NAT module has the potential to restart NAT agents during operation and possibly cause Kernel panics. Sometimes this can lead to a system reset. This only happens in Address-Only mode. (684229)
- %NAT-3-PEER_ERROR: Unexpected error (-17) encountered while programming the Kernel is a generic error indicating there is an attempt to create a

NAT translation that conflicts with an existent translation in kernel.

This bug tracks the case where this was happening due to stale TCP entries on a nat synch peer. TCP entries were not being deleted (they are after 99days in this case) due to a regression on the NAT synch code. If the TCP session was being closed from the same peer where the connection was created the close event was missed. The problem could be detected by check the output of "show ip nat translation kernel" together with the output of "conntrack -L". The stale entries would be on TIME_WAIT, FIN_WAIT or LAST_ACK for a long time. While seeing conntrack entries on this state is normal; they should disappear in minutes. (696011)

- Discovered entries are not programmed in kernel or in hardware for TCP connections with NAT sync. Due to races in the way the TCP state is propagated by NAT synch sometimes the peer doesn't receive the new state and the entry is never programmed. This bug can be detected by checking the difference between:
  "show ip nat synchronization discovered-translations" and "show ip nat translation". While it is normal to see some differences temporarily; discovered-translations appear on the translation list after some seconds. When this bug is present; the discovered-translations never make it to the translation list. (696583)

- On rare occasions during reboot the NAT agent hits a sigfault and restarts. The NAT agent should restart without further issues. (696865)

- Error: Nat: %NAT-3-PEER_ERROR: Unexpected error (-17) encountered while programming the Kernel
  This error indicates there is a conflicting NAT translation in kernel already. The error can be triggered by multiple issues. This bug has fixes for some races during NAT synchronization that causes the peer to not delete an entry that was deleted by the originator of the translation. Subsequently the nat translation is reused but causes a conflict with the stale entry.

  Cleaning the error requires manual intervention. The command "show ip nat translation kernel" can be used to identify conflicting entries. While it is possible to clear the affected connection it is better to apply patches to affected switch and reboot it. (697275)

- Under some racy conditions when a flow is being setup for NAT, if the flow is removed (e.g. FIN/RST in case of TCP), an incorrect conntrack entry without translation rules is programmed in the kernel thereby causing any subsequent packets for that flow to be not translated. (739359)

- NAT Sync feature might not recover its peer properly after the restart of a host.
  Workaround is to remove the NAT Sync configuration and re-apply it. (770609)

- Hardware translations may not properly programmed if Dynamic NAT profile is changed while high amount of NAT entries are created. (792730)

SwitchApp

- VRRP switchover does not work with MLAG interfaces. (773492)
- After MakoL3Unicast and/or MakoSharedResources agent restarts routes for connected hosts may remain programmed in the FPGA (and be visible via the 'show fpga switch ip route') even though the route and ARP entry have been removed (no longer shown in 'show ip route' and 'show arp').

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

vEOS Router / CloudEOS

- Security Advisory 85 for CVE-2023-24513.
  See https://www.arista.com/en/support/advisories-notices/security-advisory/17240-security-advisory-0085 for more details. (764777)

DCS-7150 Series

- If a multicast route (S1, G1) shares the same OIL (outgoing interface list) with
  another multicast route (S2, G2), (S1, G1) doesn't age out if source S1 is removed,
  but source S2 is still actively sending traffic.

  Workaround: None. This is a hardware limitation. (122655)
- If a multicast boundary is applied on an SVI which also has IGMP snooping enabled, IGMP snooping will not take effect. (138610)
- On DCS-7150 series, FocalPointV2 agent might restart if a misbehaving host creates a large number of new connections in quick succession, without sending data.
  This issue can be prevented by configuring NAT connection limits for such hosts. (149634)

DCS-7170 Series

- On the DCS-7170 series MLAG protocol traffic shares the same CPU queue with traffic destined to CPU. A burst of such traffic, destined to CPU, can lead to MLAG flaps. It will recover automatically when the burst is over. (449993)
- On the DCS-7170 series L3 protocol traffic shares the same CPU queue with traffic destined to CPU. A burst of such traffic, destined to CPU, can lead to

protocol flaps. It will recover automatically when the burst is over. (464981)

- On DCS-7170 series, a policer configuration with burst unit in mbytes will result in continuous crash of BfnQos agent. Remove the configuration to recover, and the issue can be worked around by specifying the burst in kbytes. (469790)
- IPv4 Multicast routes may be prematurely aged out (472845)
- Hardware errors in the switching ASIC may get silently ignored after an initial reporting. (484277)
- Monitor session with TX source to a port-channel may not be rate limited correctly. Link flap of the port channel may cause the mirrored traffic to stop. Workaround is to remove non-working port from port-channel, or remove egress mirror rate limit. (516828)
- Under heavy l3 multicast route flaps, replication update may fail, leading to traffic drops. (620107)
- Under rare circumstances, a kernel panic may occur displaying the following in the backtrace:

  ```
  [ 1194.143530]  remove_wait_queue+0x17/0x52
  [ 1194.143534]  ep_unregister_pollwait.isra.1+0x66/0x85
  [ 1194.143535]  ep_remove+0x24/0xc6
  [ 1194.143537]  eventpoll_release_file+0x59/0x80
  [ 1194.143540]  __fput+0xc9/0x1a7 (643777)
  ```

  Series Affected: DCS-7170 Series
- 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
- NAT VIP conflict for UDP traffic can occur on a NAT device configured with NAT sync. The error leads to one NAT entry being deleted and the device will continue to create new NAT translations but the flow being deleted will be affected.

The error can occur on bursts of new UDP connections happening for a period longer than 30 seconds. (712298)

- TCP NAT entries might be recreated on a NAT device configured with NAT-sync after reload. The entry will age eventually and be removed from the NAT table and kernel.
  The scenario that leads to this problem is the following:
  - peer1 is the NAT device advertising the TCP NAT entry.
  - peer2 receives a FIN or RST for the TCP connection.
  - peer1 reloads and peer2 starts advertising the connection.
  - When peer1 comes back it recreates the NAT entry.
  The entry eventually age after the TCP timeout due to lack of activity. (719526)

- Custom profile users may experience deadlock in BfnSliceAgent when using GRPC combined with counters less than 64-bits wide. (811394)

DCS-7170 Series

- When a monitor destination is continually congested due to fan-in/high rates of traffic, and LAG is used in the system, the slice agent may restart when there is LAG membership change or port status change. (516260)

CCS-720XP 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)

- When two or more hardware flow trackers are configured the forwarding agent may restart unexpectedly. In this case even if the extra flow tracker is removed the forwarding agent may continue to restart.

  Workaround is as follows:
  1) Do "shutdown" on the "flow tracking hardware"
  2) Remove all flow trackers except one in the flow tracker config. Only one flow tracker can be present at a time.
  3) Do "no shutdown" on the "flow tracking hardware"

  In case the above doesn't work you need to remove flow tracking hardware by doing:

(config)# no flow tracking hardware
and reconfigure flow tracking hardware and add only one flow tracker. (426694)

- When a device behind phone/hub/switch connected to Arista switch is authenticated via Dot1x MAC Based Authentication, it cannot be moved to different port if dynamic VLAN is not received or dynamic VLAN is same as native VLAN. (477280)
  Series Affected: CCS-720XP Series

- Authentication of any data device behind a tagged phone will fail after the phone is authenticated and authorised via EAP in phone VLAN on that port. Also, if a data device authenticates before tagged phone doing EAP, the authorisation of phone in phone VLAN will fail. (490365)

- 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)

- NAS sends two EAPOL identity requests in response to EAPOL start from supplicants. This can cause authentication timeouts for some Windows supplicants. (498610)

- Executing ASU with phone trunk ports configuration can lead to Dot1x agent being killed. (503859)

- Some front panel BASE-T interfaces might fail to forward traffic after a reload or linecard insertion (for modular systems), if they are configured with `speed 100full` command. The issue only affects the following 2.5GBASE-T interfaces:
  Et1,Et6,Et9,Et14,Et17,Et22,Et25,Et30,Et33,Et38,Et41,Et46,Et49,Et54,Et57,Et62,Et65,Et70,

  Workaround is to run `shutdown` followed by `no shutdown` on the affected interfaces. (519774)
  SKUs Affected: CCS-720XP-24ZY4 and CCS-720XP-48ZC2

- Some BASE-T ports might report linkup; but fail to forward traffic after a reload or power-cycle.

  Workaround is to reload or power cycle the switch. (523649)
  SKUs Affected: CCS-720XP-24Y6 and CCS-720XP-48Y6

- While doing ASU, Dot1x Agent may restart unexpectedly if there were 802.1X supplicants present before ASU. (558476)

- Front panel BASE-T ports might stop forwarding traffic even though they are reported as linkup.

  Workaround is to run `shutdown` followed by `no shutdown` on the affected

interfaces. (572686)
  SKUs Affected: CCS-720XP-24Y6 and CCS-720XP-48Y6

- Interfaces configured for MAC loopback are limited to 10 Mbps instead of the configured speed. (585395)
  SKUs Affected: CCS-720XP-24Y6 and CCS-720XP-48Y6

- Some link partners might receive bad traffic at 100M rate on the BASE-T ports. (631123)
  SKUs Affected: CCS-720XP-24ZY4 and CCS-720XP-48ZC2

- On a phone trunk untagged port with a large number of device, trying to change multiple VLAN membership on a switch port may lead to restarting of the Ebra agent. (793733)

Chassis Series

- In a chassis with redundant supervisors configured with the SSO redundancy protocol, after a switchover has been performed, a forwarding agent restart may cause the active supervisor to experience a fatal CPU error and restart, which will cause another switchover. (586096)
  Series Affected: DCS-7300X and DCS-7300X3 Series

DCS-7160, 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

- Under certain conditions where small packet are bursting across many interfaces, chip can get in a state of dropping half packets coming into the chip. (488904)

- 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)

- 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

- For some card types, if switchover fails during SSO, the card may not automatically recover via a card power cycle. (618704)

- A restart of the Forwarding plane agent renders the switch blackholing the incoming traffic. A slice restart resolves the issue. The following platforms are impacted: DC-7300X3 (621224)
  Series Affected: 7300X3 and DCS-7300X3 Series

- 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: DCS-7800R3 Series

- A Linecard with link issues on the control plane may cause the Supervisor to reboot (670003)

- A misbehaving card may continue to power cycle indefinitely causing unexpected forwarding agent restarts and link flaps. (689789)

- 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)

- Users can login to the standby supervisor as root on switches with dual supervisor (security advisory 0082). (723401)

- Correctable PCIe errors may be seen on a forwarding chip control plane link. (779604)
- If ManagementActive interface is configured with redundancy enabled, then ManagementActive agent can restart unexpectedly during supervisor switchover. (805315)

7500R3, DCS-7020, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7500E, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800 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-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, 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-7500E and 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)

- When a bidirectional PIM multicast receiver is connected to only one MLAG peer switch, the IGMP reports from the receiver may get dropped resulting in multicast traffic drops towards the receiver. (259308)

- 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)

- After applying or removing shaping configuration on an L2 subinterface, the L2 subinterface must be flapped for the configuration to take effect. (278730)

- 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)

- When both port-channel load balance for multicast and selective ingress replication is configured, groups that get ingress replicated incorrectly replicate packet to each member of port-channels that are member of the group.
  As a workaround one of the features has to be disabled. (344650)

- In the presence of virtual MAC with mask configuration, packets that undergo recirculation may not be properly forwarded. The workaround is to eliminate the configuration that causes the recirculation to happen. (387535)

- If VXLAN is configured, PBR and QOS policies won't apply to traffic ingressing on a MLAG peer-link. (390804)

- Traffic subject to CPU traffic policy _policer_ action will not be _forwarded_. This

includes multicast control traffic which is replicated to both the CPU but also
forwarded (e.g. PIM, OSPF HELLOs, etc.). Unroutable unicast traffic is also flooded,
and so will not be forwarded if matched by a CPU traffic policy policer action. An
example policy of concern is

```
traffic-policies
   traffic-policy EXAMPLE
      match PIM ipv4
         protocol pim
         actions
            police rate 1024 kbps
```

In this example, PIM control packets would be policed to 1024 kbps to the kernel, but
dropped completely on any outbound interface.

NOTE: If the traffic subject to policing is only bound for the CPU port (e.g. unicast control plane traffic such as BGP packets), there is no impact. (395561)

- 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)

- Insertion of a fabric module may cause brief traffic loss on the device (402508)
  SKUs Affected: 7504R3-FM and 7508R3-FM

- With scaled Acl configurations, SandAcl agent restart may result in additional agent restarts. (408260)

- With "ip pim spt-threshold infinity" configuration and having 10K or more Ipv4 (*,G) multicast routes, in addition to what hardware can accommodate, SandMcast agent may restart repeatedly. Work-around is to reduce the scale of (*,G) multicast routes to a number that can fit in hardware. (415164)

- Configuring 4094 different Egress IPv4 RACL on 4094 SVI will cause repeated SandAcl Agent crash.
  No workaround is available except for scaling down Egress IPv4 RACL. (415577)

- 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)
- Sometimes on reload MACsec sessions on MACsec proxy ports stay down when FIPS restriction is enabled. (430728)
- 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)
- With VXLAN configuration, if the route to a remote VTEP churns while the tunnel hardware resource usage is closer to the hardware limits, then the SandL3Unicast agent might restart unexpectedly. (445834)
- V6 unknown multicast packets are not headend replicated to remote VTEPs. (448950)
- 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)
- Sflow samples of VXLAN packets terminated by the switch, will not have their VXLAN headers stripped. (457096)
- "match nexthop-group" filter is not supported in PBR rules when the corresponding nexthop-group is referenced by ECMP routes. (458189)
- On DCS-When a port-based connector is used for a pseudowire or local cross-connect patch, dot1q-tagged packets will have the outer tag incorrectly removed on ingress/encap if any one of several other features are configured on the switch (on any interface) including: L2 subinterfaces, Dot1q Tunnel, Vlan Mapping, and Vxlan. The only workaround is to unconfigure these other features. (474093)
  Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- SflowAccel does not send flow samples to collectors routed via BGP routes with additional backup paths. (476553)
- If the number of BGP Flowspec matches exceed the hardware TCAM banks and entries available, and then Flowspec is enabled on additional interfaces while in the overflow state, the matches may fail to be programmed on the additional interfaces, even if the published Flowspec matches later fall back under the hardware limits. (477150)

- 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)
- The SandAcl agent may crash if PBR is applied on SVIs during a linecard hotswap (486402)
- When BGP neighbors advertise Flowspec rules containing the police action and the system's policer resources are exhausted, any further change to Flowspec rules may cause the SandAcl agent to restart unexpectedly. Upon a SandAcl agent restart, existing Flowspec rules are removed from hardware and the system will attempt to program the new set of advertised rules, which will succeed if there are enough hardware resources.

  As a workaround, restart the SandAcl agent after freeing enough policer resources, e.g., remove rules from features such as QOS policies that use policer resources. (486688)
- Configuring 'port qualifier size N bits' in 'feature flow-spec port ipv4|ipv6' causes BGP Flowspec to stop working on non-Algomatch switches. (487271)
- In a 6PE configuration, packet ingressing on an L3 subinterface with explicit null label and destined to this switch is not supported. (504582)
- Removing a field-set configuration (i.e., via "no field-set ipv4 prefix <field-set-name>") does not cause traffic-policies that match on the field-set to be updated in the switch. Packets will still match on the IP prefixes or L4-ports of the removed field-set configuration. The removal of the field-set from hardware will take effect in a subsequent change to the policy configuration. (512073)
- TCAM entries may become corrupted when two CLI commands which read TCAM entries from hardware are executed simultaneously. The commands which cause a TCAM hardware read are 'show platform fap acl tcam diff', and 'show platform fap acl tcam hw hits`. (521356)
- SandFap agent restart or a "config replace" operation may cause some subinterfaces to transition to "notconnect" state.

  "shut/no shut" the parent interface or the affected subinterface to bring the subinterface back up. (535787)
- Ipv6 multicast routes with IIF as port-channel subinterface may remain unprogrammed after reloading the peer switch to which the IIF is connected.

  Workaround is to clear the affected routes or flap the affected IIF. (537653)

- 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

- While in Tap Aggregation mode, changing the composition of a tool group may trigger a cascade of events which could lead to the group not able to replicate and transmit packets.

  To restore functionality, exit and then re-enable Tap Aggregation mode. (546239)

- In configurations involving IP/IPv6/MPLS routes resolved multiple IS-IS-SR tunnels (including indirectly via other tunnels such as BGP-LU) an IS-IS-SR tunnel failure on a remote node (such as Loopback shut, or device reset) may cause extended loss on other nodes as reconvergence occurs as the multi-path set is updated. The loss duration will vary depending on the number of resolution BGP-LU tunnels and on the number of internal VLANs (run "show ip int brief | wc -l" to get an estimate of this number) and could be on the order of 30 milliseconds per tunnel, if the VLAN count is on the order of 2000, and if both IPv4 and IPv6 routes are in use.

  A workaround may be to try to reduce the number of BGP-LU tunnels, or to manually adjust the network configuration to divert traffic from the IS-IS-SR tunnel towards the device to be reset. (546432)

- '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)

- SandL3Unicast may not program routes after reloading a switch that is fully populated with linecards. (558148)
  SKUs Affected: DCS-7516N

- For interfaces with both a primary and fallback PBR policy applied, a policy update containing multiple ACL updates could cause packets to miss both the primary and fallback policy during the update. (562198)

- A config-session that applies primary policy to an interface and updates ACLs/classes used by the primary policy 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. (578146)

- 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)
- A tcam profile change in flex-route, cbf, or egressIpv6Racl features may lead to forwarding issue. Restarting the SandL3Unicast agent would recover from the problem. (617588)
- With the configuration "ip hardware fib next-hop rfc5549 dedicated", IPv6 egress ACL may be skipped if IPv6 route is changed from drop to forwarding. (619336)
- 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)
- sand-dma causes a kernel panic with the preceded log "Assertion failed! skb,/usr/include/CpuFabric/strata-cq.h,cq_rx,line=1083". The device will reboot. (637021)
- 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
- The SandL3Unicast agent may restart if there is an excessive amount of route or next-hop churn. (678131)
- 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)
- 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)
- 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)
- The LoopProtect agent may fail to detect the loops in the network after MLAG switch reload.

  The workaround is to remove and re-apply the loop protection configuration. (732430)
- EVPN MPLS ESI filtering may not work as expected on device restart or linecard insertion.

  The workaround is to restart SandAcl agent. (754114)
- Quickly changing TCAM profiles that repeatedly enables and disables the traffic-policy feature may cause the SandTcam agent to continuously crash.

  Remove all traffic-policies applied to interfaces recover from this problem state. (759765)
- When using a large scale of PBR rules, removal and addition of member links used as nexthop groups destinations of PBR policies may cause the SandAcl agent to restart. (759846)
- Constant link flaps on a lag member of MLAG peer links may result in STP looping with VXLAN MLAG (761572)
  Series Affected: 7500R3, DCS-7020, DCS-7280CR2, DCS-7280E, DCS-7280PR3, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280SR3, DCS-7280TR, DCS-7500E, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 Series
- Performing many updates to a traffic policy in a short amount of time may cause the SandAcl and/or SandAegis agent to restart unexpectedly. (770478)
- In an EVPN MPLS configuration, flooded packets may not be sent out to peer PE if all of the below happens:
  1. Multiple peers are reachable via the same immediate nexthop
  2. The end to end tunnel connectivity goes down and comes back up.

  The workaround is to restart the SandL3Unicast agent. (772158)
- BGP Flowspec rules may fail to program after a linecard forwarding agent is restarted. (790662)
- If a linecard is removed where an interface traffic-policy is applied, it is possible for all future traffic-policy changes and/or policy applications to timeout. The workaround for this is to restart the SandAegis agent. (798182)
- Simultaneous read and write TCAM operations may update the TCAM hardware tables affecting the features using TCAM.

"show platform fap tcam hw" command is run as part of "show tech" that does TCAM read/write operations periodically. So it's recommended to disable this command from show-tech using the below commands: management tech-support
policy show tech-support
exclude command show platform fap acl tcam hw (807981)

DCS-7500E Series

- On the DCS-7280E and DCS-7500E series, disabling autonegotiation on a port with a 1G fiber transceiver can cause occasional link flaps.

  The workaround is to always enable autonegotiation on these ports. (70667)
    SKUs Affected: 7500E-72S-LC

- IPv4 egress RACLs will not filter multicast traffic in shared mode. In unshared mode, multicast traffic is filtered only if ingress replication is enabled. (131660)

- An IPv6 Egress ACL applied to a Port Channel is not applied to a newly added member interface, if the member interface already has the same ACL applied to it. The workaround is to remove the ACL from the interface before adding it to the PortChannel. (141736)

- SAND_POSSIBLE_STUCK_PORT may be logged on an Ethernet interface after switch initialization, which can cause an interface to not transmit packets despite being up.

  A workaround is to perform a full reset of the corresponding forwarding chip using the "reset platform arad <forwarding chip name> full" command. All Ethernet interfaces on that forwarding chip will flap; to determine all the Ethernet interfaces on the forwarding chip, run "show platform arad <forwarding chip name> map | grep Ethernet". (201096)

- In multi-agent mode, configuring multiple VRFs and Vxlan routing may lead to continuous OOM failures. (399674)

DCS-7280R and DCS-7500R Series

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

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

  Inserting transceiver 10 seconds after removal from the same slot is the workaround. (293594)
    SKUs Affected: 7500R-8CFPX-LC

- VRF-mapped IP over MPLS packets with TTL=1 in the IP headers may be dropped, if the installed system TCAM profile uses packet types 'mpls ipv4 forwarding routed decap' or 'mpls ipv6 forwarding routed decap'. This could be a predefined TCAM profile such as pbr-match-nexthop-group or a user-defined profile with these packet types.

  Create a new TCAM profile based on current profile and remove these packet types from all applicable features. Note that this will remove support for these packet types in egress ACLs, and 'mpls tunnel termination model ttl uniform dscp pipe' configuration will no longer work. (412811)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

- Configuring more than 6 MPLS labels for the pseudowire path (including VC label and flow label), will cause the SandTunnel Agent to crash and not recover. The only workaround is to reduce the number of transport labels and/or to remove flow label such that the total is 6 or fewer. (443140)

- 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)

- Optimized routes may not be programmed properly if TCAM profile is changed from <profile-1> to <profile-2> with both profiles having 'feature flex-route' in it.

  The workaround is to change the TCAM profile in two steps - first change from TCAM profile <profile-1> to default and then from default to <profile-2>. (483629)

- 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 the congested traffic exceeds the DRAM, we switch to using OCB. In OCB mode due to excessive copp queues even small amounts of traffic gets dropped. Thus disabling dramRecovery via "no platform sand enqueue performance-monitor " helps. (486924)

- 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)

- After hotswap of a linecard the SandAcl agent may restart. All functionality will recover but transient disruption may occur. (554932)
  SKUs Affected: 7500R-36CQ-LC, 7500R-36Q-LC, 7500R-48S2CQ-LC, 7500R-8CFPX-LC, 7500R2-18CQ-LC, 7500R2-36CQ-LC, 7500R2A-36CQ-LC, 7500R2AK-36CQ-LC, 7500R2AK-48YCQ-LC, 7500R2AM-36CQ-LC, 7500R2M-36CQ-LC, 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-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-7280SR2-48YC6, DCS-7280SR2-48YC6-F, DCS-7280SR2-48YC6-M-F, 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 and DCS-7280SR2K-48C6-M-R

- 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)

- BASE-T ports with auto-negotiation enabled with the command 'speed auto', that auto-negotiate to a speed less than 10G can flap when the "speed" CLI command is run on another port or the transceiver is removed on another port.

  The workaround is to not configure the speed with 'no speed' or force the speed of the port to the auto-negotiated speed using the 'speed forced <speed>' CLI command. The 'no speed' configuration is equivalent to the 'speed auto' configuration for the fixed BASE-T ports. (740294)
  SKUs Affected: 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 and DCS-7280TRA-48C6-R

DCS-7280R2 and DCS-7500R2 Series

- Multicast traffic on non-default VRF will be lost for as long as 15 minutes when default VRF is multicast enabled without any L3 interfaces.

Workaround is to disable mulitcast on default VRF or configure a dummy L3 interface. (619693)

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

- During restart of SandAcl agent, packets may not be subjected to ACL rules for a short period of time. (372971)
- Neighboring interfaces may flap if speeds or speed groups are configured on an interface. (394764)
- A statically configured MAC entry may be replaced by a dynamically learned MAC entry leading to transient mis forwarding of packets destined to that MAC address. This transient condition gets corrected in the next poller cycle. (428580)
- SandAegis agent may crash when the last IPv6 rule is deleted from a traffic-policy which has both IPv4 and IPv6 rules (435733)
- VXLAN IPv6 overlay configuration along with VARP and Virtual VTEP configuration is not supported. (442914)
- On Switch with Jericho2 FAP(s), if Pimsm spt threshold is configured as infinity, Pimsm multicast traffic forwarding may not work properly. (447366)
- Speed change on a port with running traffic may lead to stuck queues for packets egressing on that port. If this occurs, shut the port and restart SandFapNi agent to recover. To avoid this, shut the port before changing speeds. (479144)
- Outer DSCP is not derived from Traffic Class in VXLAN encapsulated bridged packet. (487274)
- ECN bits may not be correctly copied from outer header when packets are routed after VXLAN decapsulation. (489998)
- Forwarding of IPv4 and IPv6 multicast packets over MPLS Pseudowires is not supported. (493142)
- Configuring 513 different Egress IPv6 SubInterface ACLs will cause repeated SandAcl Agent restarts. No workaround is available except for scaling down Egress IPv6 SubInterface ACLs. (502185)
- In interface traffic-policies, when configuring a destination port after configuring source port field-set, the ports are treated as field-set names (506962)
- When two or more of the following features: Port ACL, PBR, QoS, Flowspec, Aegis, and CPU Policy are applied on the same interface with a non-default TCAM profile, one or more of the features may not work (526913)

- If the "acl subintf ip/ipv6/mac" feature is added to a TCAM profile without the corresponding "acl port ip/ipv6/mac" feature being added to the same profile, subinterface ACLs will not work. (532393)
- MTU violating VXLAN routed packets with either IPv6 overlay or an Ethernet tag on the underlay ethernet header can cause the egress port to lock up. When in this state, no packets will egress the port and traffic will get deleted on the ingress VOQ. Shutting and unshutting the port will restore the traffic. (537323)
- A VXLAN packet that would be bridged, after decapsulation, to the same interface as it ingressed will be dropped instead of being bridged. (549785)
- MPLS tunnel terminated GRE-over-IPv6-over-MPLS packets are forwarded incorrectly. (550842)
- When VXLAN is enabled, L4 port match in ACLs on VXLAN decap packets does not work. (567246)
- 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)
- 1. A traffic-policy with both v4 and v6 rules applied on an interface may be deleted when
  either all the v4 rules are removed or all the v6 rules are removed.
  2. When traffic-policies are applied and deleted on multiple interfaces in the same config
  session commit, some changes to the traffic-policy configuration on the interface may
  not be applied.

  Thw workaround is to restart the SandAegis agent (578658)
- 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)

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-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

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

- 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)

- Configuring multiple DirectFlow flows of different priorities at once may restart the SandAcl agent continuously. A workaround would be to configure the DirectFlow flows one after the other. (543622)

- The flexroute feature may override the forwarding decisions from PBR policies and BGP Flowspec matches. (544872)

- SandAcl agent may restart when programming flows at many different priorities. Workaround is to limit the number of different priorities to 1024. (588420)

- In an EVPN MPLS multihomed configuration, packets egressing towards non designated forwarder ports may not be dropped. (592102)

- When two prefix lengths are configured with 'ipv6 hardware fib optimize prefix-length' command, the packet matching the smaller prefix length will be mis-

forwarded if there is a matching less specific IPv6 route in the routing table for the same packet. (635484)

- If the forwarding agent restarts, subsequent FEC configuration may not take effect.

  If MACsec is configured, workaround is to power cycle the linecard to avoid key mismatch. Otherwise workaround is to run the "agent Enigma-LinecardX terminate" CLI command, with X being the linecard slot number, which can flap all interfaces on the linecard. (666830)
  SKUs Affected: 7500R2AM-36CQ-LC and 7500R2M-36CQ-LC

- 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)

- Certain features may stop working after a reload on switches that meet all of the criteria listed below:
  - at least one group of SFP28 interfaces is present on the switch
  - at least one interface in the group does not have a transceiver

  The following features may stop working:
  - L2 or L3 Subinterface
  - VLAN Mapping
  - EVPN MPLS
  - VXLAN
  - Private VLAN
  - TAP Aggregation
  - L3 interface

  The workaround is to flap the packet's ingress interface by doing "shut" followed by "no shut". (702086)

- In an EVPN network, SandMact agent might restart due to a SIGQUIT while processing the MAC addresses polled from the hardware MAC table if the number of MAC addresses is close to the hardware MAC table limit. (718204)

- 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)

- When the "packet IPv6 forwarding routed decap" is missing from the operational TCAM profile, tunneled traffic might be misforwarded after decapsulation if the decapsulated packet is IPv6 and CPU-bound.

  The workaround is to apply a TCAM profile that includes "packet IPv6 forwarding routed decap" in any TCAM feature. (729017)

- Parity error in EGQ_VSI_MEMBERSHIP table can lead to packet loss. (737517)

- Configuring tail drop thresholds on an interface, before it's up and running state, will result in assertion of SandFap agent.
  Workaround is to configure the tail drop thresholds once the interface is in up and running state. (753376)

- If an egress IPv6 ACL is applied on a port-channel interface, packets that should be forwarded through that port-channel may be sent to the kernel after the device reload or inserting a linecard.

  The workaround is to shut and no-shut the port-channel interface. (802730)

- DirectFlow cannot enforce a new destination with its actions 'output nexthop' and 'output interface' for traffic destined to the CPU. The flow's counters will increase but the traffic will not be redirected. (809100)

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

- On large systems with more than 4 linecards, the SandCounters agent may restart continuously. (536744)
  Series Affected: 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)

- 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)

- Egress IPv6 ACL logging do not work on packets subject to layer 3 MPLS termination. (640491)
- 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)
- For EVPN VXLAN Multicast, the IIF underlay switchover does not work. The workaround is to clear the mroute with "clear ip mroute *" command. (647945)
- 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)
- IPv6 linklocal packet received on front panel ports or sourced by CPU do not get headend replicated to remote VTEPs when IPv4 or IPv6 multicast routing is enabled on the VXLAN VLAN. (690105)
- 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)
- 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)
- In an EVPN VXLAN setup, the shared router MAC might not be programmed in hardware when the local VTEP IP is removed and re-added. This could result in SW forwarding of packets that are received from the VXLAN tunnel. The workaround is to restart the SandL3Unicast agent. (745442)
- When the agent SandL3Unicast restarts, some previously configured monitor sessions mirroring traffic to an IPv6 GRE endpoint may stop mirroring, even if the session is displayed as programmed by CLI and a nexthop is known for the endpoint.

  The workaround is to delete and add the monitor session again. (787948)
- When an SVI is configured on the VXLAN VLAN, unknown unicast and multicast packets are unnecessarily sent to the CPU after VXLAN decapsulation. (788956)

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

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

DCS-7280R and DCS-7280R2 Series

- When ACLs are applied on a large number of interfaces (more than 1000 SVI interfaces), the SandHalo agent may restart with a SOCKET_CLOSE_LOCAL syslog message. (253015)
- Traffic egressing out of a LAG interface may not be subject to egress ACL on reload (473227)
- RACLs are not supported for VXLAN traffic ingressing on a MLAG peer-link when AlgoMatch is enabled. (485677)
- For an ACL applied to a LAG subinterface when all members of the LAG subinterface on a given chip are down, subsequent edits to any ACL on that chip may cause the default configured action to be applied to all interfaces on the chip where an ACL is present.

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

  To recover when this bug:
  - Remove the "ip access-group" configuration from LAG subinterfaces that are down
  - Remove the LAG members that are down from the LAG that has subinterfaces with an ACL configuration (598255)
- When AlgoMatch is enabled and MPLS is a configured feature in the TCAM profile, mpls-in-ip packets (ip-protocol=137) may bypass ACLs configured on interfaces. (679923)
- When using AlgoMatch ACLs, certain sequences of configuring a routed port with ACL applied that is later configured to be part of a LAG with an ACL applied may cause traffic to be dropped on all ports with ACLs applied:

  1. configure routed port with an ACL
  2. configure routed LAG with no ACL applied
  3. add routed port as member of routed LAG

4. configure routed LAG with an ACL

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

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

- In multi-chip systems with L2 MTU greater than 9236, packets that are larger than 9236 bytes that flow between the switch ASICs will be dropped. (102837)
  Series Affected: 7300X3, DCS-7250X, DCS-7260CX and DCS-7300X 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)

- BUM traffic may temporarily loop into the MLAG pair when a member interface in the peer-link port channel is brought up. (259213)

- When IPv4 uRPF exceptions are already programmed, configuring IPv6 uRPF exceptions imposes a heavy processing load on the Strata Slice agent to do a transition from IPv4 to IPv4-v6 exception mode in the hardware. Therefore, under a scaled config scenario, the Strata Slice agent may restart repeatedly.

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

- IS-IS configured over LAGs may see neighbor adjacency flapping during a hitless restart. (276232)

- A PCI error can break counter collection for an ASIC.

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

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

  Workaround is to perform hitful restart of forwarding agent. (283175)
  SKUs Affected: DCS-7050CX3-32S, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050SX3-48YC12, DCS-7050SX3-48YC12-F, DCS-7260CX3-64, DCS-7260CX3-64-F and DCS-7260CX3-64-R

- Egress ACL on SVI for BUM traffic will not match ACL rules for packets being sent across fabric interfaces. (284543)
  SKUs Affected: DCS-7304 and DCS-7308

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

- 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

- DirectFlow entries that redirect traffic out of a port-channel may incorrectly egress packets on none or multiple members of the port-channel. (378467)
  SKUs Affected: 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-7304, DCS-7308 and DCS-7316

- 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-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)

- On ECMP shrinkage due to link down events, VXLAN tunnels reachable via ECMP may experience extended outage until routes reconverge. (419753)

- 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 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)

- Special filter is installed on MLAG port channel member interfaces to block packets ingressed from MLAG peer via peer link at egress to avoid loop. When the filter is not installed on a member interface before it is added to MLAG port channel, there might be a momentary loop when interfaces belonging to an
  MLAG port channel link up. (471268)

- 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)

- Slice agent might restart repeatedly in scale QoS policy-map configuration with policers, which causes hardware policers to exhaust. The workaround is to reduce number of policy-maps with policers attached to them. (489846)

- 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)

- It is possible that storm-control configuration will still be applied in hardware after removing the configuration. Workaround is to restart the Strata agent. (504901)

- Enabling per VLAN routing with an L3 interface that only has an IPV6 address will result in no routing of packets received on this L3 interface. Workaround is to flap the interface. (528490)

- During an SSO failover, interfaces with VXLAN configured may result in an increased traffic outage on the associated interface (528749)
  Series Affected: DCS-7300X Series

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

  On 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 and CCS-720XP-48Y6

- Attaching a policer to class-default of an egress policy-map consisting of both IPV4 and IPV6 ACLs, is not supported and may lead to forwarding agent restart (534152)

- When an MLAG peer is reloaded with VXLAN dynamic VLAN to VNI configuration, some VNI entries may not get programmed in hardware resulting in traffic outage.

The workaround is to restart the StrataL2 agent. (549749)

- 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)

- 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)

- Ingress VNI rate-limiting feature is not working for VXLAN routed traffic. (587489)
  SKUs Affected: DCS-7060CX-32S, DCS-7060CX2-32S, DCS-7260CX-64, DCS-7260CX3-64, DCS-7260CX3-64E and DCS-7260QX-64

- 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)

- A SSU reload may result in the SCD watchdog causing an extended traffic outage. (605183)

- 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

- MAC addresses might not age out on port-channels spanning across multiple chips (616314)
  SKUs Affected: 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-7304 and DCS-7308

- When the "hardware counter feature vlan-interface/subinterface out" counter feature is enabled with MPLS nexthops, MPLS routes may get programmed incorrectly resulting in traffic black-holing. (625005)
  Series Affected: 7300X3, 7320X, 7368, CCS-720XP, DCS-7050CX3, DCS-7050QX, DCS-7050QX2, DCS-7050SX, DCS-7050SX2, DCS-7050TX, DCS-7050TX2, DCS-7050TX3, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060SX2, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7250X, DCS-7260CX, DCS-7260CX3, DCS-7260QX, DCS-7260X and DCS-7260X3 Series

- DirectFlow flow with port channel defined as output interface may flap after port member status changes. (635036)
  Series Affected: 7320X, DCS-7050QX, DCS-7050SX, DCS-7050TX, DCS-7060SX2, DCS-7260CX, DCS-7260CX3, DCS-7260QX and DCS-7300X Series

- StrataL3 agent may restart unexpectedly when StrataL3 forwarding agent is restarts or when underlay core interface flaps with scaled VXLAN VTEPs. Workaround is to reduce the VTEP scale. (637532)
  Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050TX3 and DCS-7060 Series

- Performing StrataL3 agent restart may result in excessive traffic outage for VXLAN routed traffic. (638886)

- With VXLAN VLANs configured, hitless reload may take longer or eventually time out, causing a hitful reload. (641005)
  Series Affected: 7320X, DCS-7050X, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260CX3, DCS-7260X and DCS-7300X 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)

- 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)

- 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. (649075)
  Series Affected: DCS-7050TX3 Series

- 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 and DCS-7010T-48-R

- 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: DCS-7050TX3 Series
  SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-24ZY4-F, CCS-720XP-48ZC2 and CCS-720XP-48ZC2-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

- 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-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7300X 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)
- 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. (716820)
  SKUs Affected: DCS-7060DX4-32-F and DCS-7060PX4-32-F

- 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

- 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. (716822)
  Series Affected: DCS-7260CX3 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. (716825)
  SKUs Affected: DCS-7050SX3-48YC8

- 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. (716827)
  SKUs Affected: DCS-7050CX3-32S

- 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. (716829)
  SKUs Affected: DCS-7050SX3-48YC12

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

- StrataLag agent will continuously restart unexpectedly if there are more LAGs configured (including LAGs with no members) than the hardware can support.
  The LAG hardware limits can be seen by executing CLI 'show port-channel limits', field 'Max port-channels per group'.

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

- On some 1000BASE-T ports, PTP in E2E one-step transparent clock mode may be inaccurate after speed changes.

  To restore correct PTP behavior for this mode, configure the desired speed without autonegotiation. If changing the speed to 1g, a chip reset will also be necessary after configuring the ports. (737105)

- An interface configured to be a member of a LAG may fail to be programmed in the hardware.
  The CLI 'show port-channel detailed' will show the status of the interface as "Waiting for hardware to program port to RX...".

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

- Configuring the storm-control feature when TCAM is fully utilized causes the forwarding agent to crash. (748934)

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

- A short AC input power loss glitch can cause the system to enter a stuck power state where it must be manually power cycled. (772073)
  SKUs Affected: 7368X4-SC

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

  Restarting the StrataVlanTopo agent fixes the issue. (780110)
  SKUs Affected: DCS-7060CX2-32S, DCS-7060SX2-48YC6, DCS-7260CX-64, DCS-7260CX3-64 and DCS-7260QX-64

- IP PACL configured on interface receiving VXLAN traffic encapsulated to VARP VTEP IP may result in drops.
  Workaround is to configure port range rule in the ACL or apply the following command: `platform trident tcam port-acl vfp disabled`. (784384)

- Looping or double delivery of L2 protocol packets received on the peer-link may be observed in an active-active MLAG setup when L2 Protocol Forwarding is configured on the peer-link. (784857)

- When an interface having "switchport trunk private-vlan secondary" configuration is removed from a port-channel we may observe incorrect tagging on packets egressing the interface.
  A workaround is to remove and add the""switchport trunk private-vlan secondary"" configuration. (786324)

- The strata-dma kernel driver may trigger a kernel panic during bootup if it module insertion coincides with outgoing IPv6 DAD traffic. (788206)

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

  On EOS versions prior to 4.29.0F, on fixed systems, the workaround is to run `platform trident reset`; on modular systems, the workaround is to run `platform trident <linecard> reset` on the affected linecard. On EOS versions 4.29.0F and onwards, the workaround is to run `reset platform trident phy interface INTF pcs` on one of the affected interfaces. (792495)
    SKUs Affected: CCS-720XP-24Y6 and CCS-720XP-48Y6

- IPv6 Underlay is not supported. Receiving EVPN routes from IPV6 VTEPs may result in continuous restart of L3 forwarding agent. (794373)
    Series Affected: 7320X, DCS-7050QX, DCS-7050QX2, DCS-7050SX, DCS-7050SX2, DCS-7050TX, DCS-7050TX2, DCS-7050X2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series

- Too much churn in CoA messages from RADIUS server can cause forwarding agents to continuously restart unexpectedly. (797268)

- Double tagged VXLAN packets sent to CPU, via sFlow, mirroring to CPU, or the CopyToCpu TCAM action, might result in excessive /var/log/messages output. (799760)

- If secure MAC moves are disabled and the port-security configuration is removed from a port-security protect mode port, some MACs might not be cleared or age out.

  A workaround is to enable secure MAC moves using 'switchport port-security mac-address movable'. The problematic MACs can be removed by clearing them

individually using 'clear mac address-table dynamic vlan <vlan> address <mac>' (799899)

- VLAN translation may stop working on a port when:
  - the port was removed from a LAG
  - that LAG has the same VLAN translation configured as the port

  Workaround:
  - restart StrataL2 agent (809637)

- BessMgr agent CPU usage may be high when underlay multicast and sFlow are configured. (811844)

- If a MAC flaps between a port-security protect mode interface and another interface rapidly, it may be displayed on the wrong interface in 'show mac address-table' and might not be counted in number of allowed addresses on the port-security interface (814589)

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-7060X and DCS-7260X 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)
- 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. (625317)
- If hitless reload is performed multiple times with port-security, VXLAN, and 'switchport port-security mac-address movable' configured, some MACs might be duplicated in the hardware and not age out even in absence of traffic. (797299)

7368 and DCS-7060X4 Series

- 802.3X pause is not supported (348758)
- L3 flooded packets do not preserve VLAN priority values. (365998)
- 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
- 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)
- If test patterns are configured on certain interfaces in startup-config, or in a configuration used with "configure replace" which also changes interface speeds, Strata forwarding agent may restart causing all links to flap upon booting or upon replacing configuration, respectively. (435141)
  Series Affected: DCS-7060X4 Series
- Speed change on a port-channel member can cause a flap in BFD and BGP sessions for the entire port-channel (520638)
- Changes in interface L1 configuration via CLI such as speed, forward error correction, or shutdown while traffic is running may cause an unexpected forwarding agent restart. This will lead to a brief traffic outage on all

ports. (595680)
  Series Affected: DCS-7060X4 Series

- Packets with internal VlanId on subinterface will not be dropped. (747009)
- With no QoS configuration, if the packet sizes vary for different queues then significant differences for traffic latency among different queues may be observed. (792903)

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)
- (A1 silicon only) MPLS traffic gets incorrectly prioritized, causing other traffic on the same port to get dropped. (160269)
  SKUs Affected: DCS-7050QX-32
- 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)
- In the presence of large amount BUM traffic across several interfaces, discards may be observed on some CPU queues. (665547)
- If multiple dot1q tunnel VLAN translations are configured with the same outer VLAN ID (S-TAG), the StrataL2 agent may restart when translations are deleted. (729031)

CCS-720XP Series

- Setup with only Storm-control and scaled IP Locking configuration may cause slice agent to restart continuously. (599935)

DCS-7050X3 Series

- Interfaces may be unable to transmit or receive packets after the hitless reload or hitless restart of the forwarding agent, followed by interface speed change, or transceiver removal or insertion. An indication of the problem is the "inDiscards" in the output of "show interfaces counters discards". To recover, execute "platform trident reset". This command will cause all ports to flap and a traffic outage to occur. (570820)
- 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-7050SX3-48YC8

7300X3 and DCS-7050X3 Series

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

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

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)

- 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)

- For QSFP28 optical modules, if the module does not enable Tx and Rx clock data-recovery by default, EOS may not enable it when operating at 100G or 4x25G. This may result in data corruption or link flaps over longer fiber links. (667187)

- 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)

- QSFP-100G-CWDM4 and QSFP-100G-XCWDM4 transceivers with serials beginning with XHT may experience high pre-FEC BER and uncorrectable FEC errors on the peer or fail to link up.

  A workaround is to adjust the tx input equalization of the transceiver with the command 'transceiver electrical lane 1-4 tx-input-equalization module-default' in config-if mode. The command is to be applied on the /1 of the QSFP interface. (739931)
  SKUs Affected: 100GBASE-CWDM4

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

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

7300X3, CCS-720XP and DCS-7050X3 Series

- When performing SSU, there may be failure in table programing resulting in fatal error cold reboot with extended traffic outage. (399774)

- 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)

- If a hitless forwarding agent restart is attempted while the hardware configuration does not match the running configuration, then it may cause an additional forwarding agent restart and loss of traffic. This may also happen if the forwarding agent is restarted while some interfaces have been disabled due to insufficient logical ports. (417835)
  SKUs Affected: DCS-7050SX3-48YC8-F and DCS-7050SX3-48YC8-R

- When "platform trident forwarding-table flexible" is configured, L3 unicast forwarding to hosts / L3 multicast forwarding may not work resulting in unexpected behavior. The workaround is to remove the CLI and use "platform trident forwarding-table partition" instead. (421084)

- Performing SSU may result in extended traffic outage, leading to total traffic loss being slightly over 1 second. (477139)

- 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 and DCS-7300X Series

- Egress ACL rules can not be used to match ip protocol 50. (535104)

- 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)

- 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 Series

- VXLAN encapsulated packets with inner destination MAC as the switch MAC may get dropped when received on a L2 only VXLAN VLAN. (582895)

- Removing SVI configuration can result in Segment Security being disabled on the given VLAN.

  Workaround is to delete and re-create the VLAN. (588149)

- 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)

- In a EPVN VXLAN Multihoming deployment, the flapping of ES-link can lead to duplicate BUM traffic to be forwarded by the non-DF device. (615803)

- When a NAT port is admin shutdown, some dynamic NAT connections may not get cleaned up in hardware.

  The workaround is to restart StrataNat agent. (619186)

- 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)

- Prefixes used exclusively in match-lists for MSS-G feature and not provided by FIB can result in LPM match causing traffic to be forwarded through CPU.

  There is no workaround for this issue. (634355)

- When IGMP Snooping is enabled, unknown multicast traffic may have a brief traffic outage on all interfaces when a speed change or transceiver gets inserted or removed on a QSFP port. (638649)

- A single-event upset (SEU) in the IFP_TCAM can lead to forwarding agent restart, after which the switch forwards traffic normally. (641105)

- 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)

- During config replace with NAT configuration, TCAM churn may result in forwarding agent restart. (676117)

- 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)

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

  Workaround if to restart the slice agent after stopping the traffic. (753304)
  Series Affected: 7300X3, DCS-7050CX3 and DCS-7050TX3 Series

DCS-7160 Series

- In some cases, when PFC (Priority Flow Control) is configured on an interface and if the user performs a "default interface" on a range of interfaces, the XpQos agent is restarted. (245559)
- In the event of the XpAcl agent restarting, the hardware programming of ACL may not be correct.

  In order to recover, remove the ACL configuration on ports/SVIs, and reattach them. (251114)
- On XP platforms, if XpMact agent restarts, and we have a large number of learned mac addresses and a large number of host routes, some of the host routes may get removed from hardware.

  To recover, clear all learned mac addresses and restart XpMact. (255279)
- VXLAN routing & bridging traffic may fail to converge after Sysdb agent is restarted.

  A workaround is to reboot the switch. (255957)
- In cases of high TCAM utilisation by Policy-based Routing, a 'configure replace' can result in Incorrect TCAM programming.
  In such cases, remove and re-apply the same policy-map which will result in programming the correct TCAM rules. (260906)
  Series Affected: DCS-7160 Series
- XPL3Unicast process may unexpectedly restart when the switch is reloaded if PBR policies with nexthop VRF actions are applied on an interface in the non-default VRF. (262713)
- IGMP IPv4 multicast control packets may create loop over a Vxlan Tunnel , if IGMP snooping is turned off in a VLAN. Workaround is to enable IGMP snooping in a VLAN (which is enabled by default in EOS) (429933)
- Unexpected restart of the XpMact forwarding agent while there are unprogrammed MAC entries can cause some IPv4 and IPv6 unicast host routes to become unprogrammed. The fact that the routes are not programmed is only visible via "show platform xp ip(v6) route".

A workaround for the inaccurate bookkeeping is to restart the XpL3Unicast forwarding agent. This will not reprogram the missing host routes.

A workaround for the missing host routes is to reduce MAC scale and then restart the XpL3Unicast forwarding agent.

To prevent this, configure "platform xp host-table dedicated" so that host entries cannot be overwritten by MAC entries. (552433)

- If Dynamic VLAN to VNI map changes while forwarding agent XpMact is restarting, forwarding agent might get killed repeatedly due to stale entries.

  As a workaround, reload this switch. (619014)
  SKUs Affected: DCS-7160-32CQ, DCS-7160-32CQ-F, DCS-7160-32CQ-R, DCS-7160-48TC6, DCS-7160-48TC6-F, DCS-7160-48TC6-R, DCS-7160-48YC6, DCS-7160-48YC6-F and DCS-7160-48YC6-R

- Hardware installation failure of ECMP host routes may cause the XpL3Unicast forwarding agent to restart. (702552)

Limitations and Restrictions in 4.23.14M

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)

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)
- 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)

Containerized EOS

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

vEOS Router / CloudEOS

- The /mnt/flash/cloud_ha_config.json based configuration for the CloudHa agent is not supported from this release. If you are upgrading from an older image (< 4.20.5) please reconfigure HA using EOS CLI.
  Please see information on converting json file based config to CLI commands in the vEOS Router Configuration Guide. (264660)
- While vEOS instances in an AWS cloud can be created using either a Bring Your Own License (BYOL) or Pay As You Go (PAYG) format, there currently is no show command in vEOS to easily help the user or support determine which mode this instance uses. (272649)
- Tx/Rx ring parameters of an "Elastic Network Adapter (ENA)" interface cannot be set through config.json in vEOS. (276382)

- vEOS supports up a maximum of 8 physical cores. With hyperthreading enabled it is 16 vCPUs (278286)
- In dpdk based vEOS(MODE=sfe in /mnt/flash/veos-config), 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)
- Site-to-site path monitored under "router path-selection" may experience flapping, when a lot of packets on the path (>10%) are dropped due to too much traffic injection into the virtual router. Workaround is to use QoS service policy to assign Inband Telemtry control packets (default UDP destination port is 4789) with a higher priority. (434087)
- CloudEOS does not support more than 10k /30 routes. The SFE agent will restart if more than 10k /30 routes are added to the system. The workaround is to reduce the number of /30 prefixes. (540204)
- NIC i40en driver version 2.2.4.0 is required on ESXi 7.0 host for SRIOV virtual function trust mode on/off functionality. (680174)
- NAT interface configuration mode is not supported in SFE. This needs to be kept in mind when migrating to SFE mode from Sfa.

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

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)

- The Arista OpenStack plugin has been enhanced to create a default route for the routers created in OpenStack, when VRF creation and vrf_default_route option is enabled in the networking-arista plugin configuration (295274)

- 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)
- 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)

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)
- A range of 0.01 sec to 255 seconds is allowed for the YANG VRRP advertisement-interval value (range is 1..25500 as advertisement-interval is specified in hundredths of a second), but EOS only accepts values in the range 1 sec to 255 secs. (426391)
- Egress MAC ACL logging is not supported (433169)
- In the 'show interfaces interactions' CLI, the list of speeds an interface is limited to does not account for the limitations that result from speed-group configuration. (434363)
- SSH may fail when sFlow and Management SSH have the same QoS DSCP value configured, and sample rate for sFlow is set to dangerous. (434763)
- 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)
- Installed licenses without corresponding licensing features configured might not be visible from the CLI (504793)
- The YANG path '/network-instances/network-instance/mpls/global/config/null-label' is unsupported, but it is incorrectly listed as supported, and accepted by the OpenConfig agent. (514041)
- The following YANG paths are unsupported, but are incorrectly listed as supported and accepted by the OpenConfig agent.

  /network-instances/network-instance/policy-forwarding/policies/policy/rules/rule/action/config/decapsulate-gre
  /network-instances/network-instance/policy-forwarding/policies/policy/rules/rule/action/state/decapsulate-gre
  /network-instances/network-instance/policy-forwarding/policies/policy/rules/rule/action/config/discard
  /network-instances/network-instance/policy-forwarding/policies/policy/rules/rule/action/state/discard (515061)
- If Octa, OpenConfig or TerminAttr are started under low memory conditions, they can cause a spike in CPU usage before unexpectedly restarting. (534841)
- 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)
- When running CLI command `show processes top`, or the "top" command via the bash shell, statistics about %CPU time are underestimated for most processes.
  To workaround this: use "top -b -n2" and ignore the first output (562245)
- Sampled flow tracking doesn't export proper bridging information (ingress VLAN, egress interface, egress VLAN, etc.) in IPFIX flow records for flows on layer 2 subinterfaces when the bridging VLAN is different than the dot1q VLAN.
  The workaround is to use the same dot1q VLAN and bridging VLAN for these layer 2 subinterfaces. (572688)
- 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)

- Sampled flow tracking doesn't export proper routing information (VRF, next-hop, egress interface, BGP information etc.) in IPFIX flow records for flows on sub-interface with QinQ configuration. (575271)
- The 'logging level all' CLI expands to individual commands in the configuration, preventing dynamic updates of the logging list as part of EOS upgrades. (579047)
- TACACS+ authentication and authorization through SSH does not work with Pulse Secure server because EOS sends different port names in authentication ("ssh") and authorization (actual vty) requests which are rejected by the server. (583678)
- Sampled flow tracking exports IPFIX flow records with egress interface as 'discard' for flows on pseudowire patched interfaces. (603916)
- A CPU internal error (BIT30 SyncFlood BIT11) may cause an unexpected reboot (623392)
  SKUs Affected: DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-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-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-96-F, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F and DCS-7280PR3K-24-F
- 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, DCS-7300X, DCS-7300X3, DCS-7500E, 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: DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F and DCS-7280CR3MK-32P4-R

- When Path MTU discovery is enabled under static tunnel interfaces using "tunnel path-mtu-discovery", the discovered MTU is only used for software forwarding in the slow path. When MTU is configured under static tunnel interfaces using "mtu" config, the configured MTU value is used for hardware forwarding. (681101)

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

- If 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)

- 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-7050SX3-48YC8-F, DCS-7050SX3-48YC8-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-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-96-F, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F and DCS-7280PR3K-24-F

- 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)

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

  To workaround this issue, if a rule in middle of ACL needs to removed then,
  1. Copy the rules which appear after the rule of interest,
  2. Remove all rules from the rule of interest till the end of ACL,
  3. Re-add rules copied in step 1. (743473)

- JSON does not specify the precision of a Number in the specification. eAPI can generate numbers up to 6-bit, which can be truncated by some JSON clients with lower limit (e.g., some implementations have a max value of 2^53-1). The only workaround is to use a different JSON deserializer (769194)

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

Layer 1

- The MACSec outPktsEncrypted and outOctetsEncrypted counters are not supported with the MACSEC Proxy feature (281715)

- On Vxlan MACsec proxy, disable learning is not supported on proxy patch vlan. (284348)
  SKUs Affected: 7500R2M-36CQ-LC, DCS-7280CR2M-30-F, DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R

- MACsec proxy for front panel ports requires replay protection to be disabled on the MACsec profile, otherwise some packets will be dropped due to sequence error. (339533)
  SKUs Affected: 7500R2M-36CQ-LC, DCS-7280CR2M-30-F, DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R

- 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-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 and DCS-7050CX3M-32S

- MACsec frames are dropped when the sci setting in the "mac security profile" does not match the peer's sci setting. (664596)
  SKUs Affected: DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F and DCS-7280CR3MK-32P4-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 and DCS-7050CX3M-32S-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-720XP 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)
- Removes the VRF option from the `show ip hardware ale l2-adj` command. (588134)
- 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)
- Loop protection is not intended to be used along with Spanning Tree Protocol (STP) (745522)

- A clock source with a lower quality level than the internal oscillator is incorrectly shown as standby in "show sync-e selection" when it is unusable. (761975)
- LAG interface counters will have constant counts discrepancy compared to the aggregate interface counts of its member interfaces when there is a link flap event on LAG interface. (785073)

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)
- SSO is not supported when BFD is configured for ISIS IPv6 sessions. (239704)
- 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)
- IPv6 Egress ACLs may not be applied to traffic if 'ip hardware fib next-hop sharing' is also configured (446000)
- PBR policy maps which match nexthop-groups will not be applied on platforms that have hierarchical FEC support, but have disabled it with the "ip hardware fib hierarchical next-hop disabled" or "ip hardware fib hierarchical next-hop max-level 1" commands. (471880)
- 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)

- 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)

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)

- 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)

- 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)
- BMP, SNMP and EosSdk will not report multiple BGP peers configured with the same BGP peering address. (602657)
- When default-originate is configured, default route will not be advertised to a neighbor if a prefix-list configured directly for this neighbor denies the default route. (648587)
- When multiple single-hop EBGP sessions are configured using point-to-point physical interface IP addresses between the same pair of switches, some of the BGP neighborship may get established using the other peer's IP address. When any BGP neighborship comes up using such wrong IP address from cross-link, the BGP neighbor configured on the other link will have connection establishment failure with Notification error: 'connection collision resolution'.

  To workaround this issue, either use 'update-source' option for the neighbor or explicitly enable 'no neighbor <addr> route-to-peer' for all the single-hop EBGP peers using point-to-point physical interface IP address. (686390)
- BGP TCP-AO is not supported on cEOS. (762218)

Rib

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

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

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

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

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

  To deny FIB programming of routes, the RIB route control configuration policy in the form of "rib ipv4|ipv6 fib policy" can be used.
  To prevent RIB redistribution of routes, the routes can be filtered out as part of the "redistribution" policy into the target protocol. (663662)

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

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

MLAG

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

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

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

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)
- MPLS vias in the LFIB are limited to 128 even when the maximum allowed ECMP or UCMP values exceed this number (602879)
- 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)
- With "hardware next-hop fast-failover" configuration, Ti-LFA MPLS tunnel switching over from a path without a MPLS label to a path with a MPLS label isn't supported. (631560)

Multicast

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

  Use "no ip igmp snooping VLAN <vlanId> immediate-leave". Otherwise, some desired multicast traffic might be lost. (21367)
- After ASU reload, if the first PIM hello packet received from the neighbor gets lost and If the neighbor has a high hello query-interval, it will result in traffic loss. Workaround is to use default PIM hello query-interval and increase the PIM hello query-count. (341447)
- With scale greater than 5000 multicast IPV6 routes, show platform sfe mroute ipv6 <vrf> commands might not display the right number of routes when performing multiple addition/deletion operations across VRFs. This is applicable only if user level multicast forwarding agent is configured to be used instead of kernel. (378841)
- Bessd might unexpectedly restart on config replace with scaled setup over 30 VRFs and large number of PIM enabled interfaces. Work around is to reduce the scale. (398185)

- Pim "non-dr install-oifs" feature becomes disabled if the non-dr needs to route multicast traffic for any reason to the interface on which the feature was originally configured. This is how the feature is expected to work. The workaround is to setup the network in a way that the DR always becomes the assert-winner. However, if the non-DR becomes assert winner for some routes on an interface and feature becomes disabled, flap the pim sparse mode configuration on non-DR on that interface to attempt to get out assert winner state. (401672)

- Multicast routes are not created when a source is transmitting only 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)

- When using SwitchApp with routing and MLAG, if link failure causes more than 64 ARP entries to be needed on the peer link then traffic may be dropped.
  Design the network with care to ensure that this limitation does not occur for the relevant failure scenarios. (765647)

DCS-7150 Series

- Under extreme circumstances where a high volume of broadcast or multicast traffic is sent from 10G ports to a mix of multiple 1G ports and one or more 10G ports, continuous over-subscription of the 1G interfaces could result in consumption of all system queue descriptors.

  If this occurs, the 10G ports may be limited to 1G transmission rates. While this is an extreme case, CLI commands can be used to tune the maximum queue length available to 1G interfaces to mitigate this risk. (32294)

- NAT translation does not happen for a packet if the ARP for the next-hop is not resolved or is aged out. Translation will start happening again once ARP resolution completes.

  The workaround is to install a static ARP entry for the next-hop that the NAT translated packet will take. (40132)

- Source port loopback suppression at high traffic rate can reduce the egress bandwidth of a port. This will happen when a port forwards incoming multicast traffic at a high rate to a multicast group which also has this port as a member. This port will receive back copies of the replicated

traffic and drop them due to loopback suppression if there are no receivers on it at the egress stage of the port. (40294)

- NAT and VXLAN cannot be configured on the switch at the same time. (47258)
- When agile ports are configured for 40G with clause 73 autoneg (speed auto 40gfull) while using CR cables, the ports cannot be used in a port channel.

  The workaround is to force the speed to '40gfull' on the ports. (55517)
- Systems with A0 silicon, packets are not forwarded to the egress port when egress truncation and timestamping are enabled together. Removing the configurations resumes packet forwarding. Truncation and timestamping are not supported together. (72756)
- The VXLAN flood VTEP list cannot be more than 128k VLAN-VTEPs, where VLAN-VTEPs is the number of VLANs multiplied by the number of VTEPs. (73228)
- Ingress port access control list (PACL) on a VXLAN-enabled interface is not supported. If ingress PACL is configured on a VXLAN-enabled interface, in-hardware VXLAN encapsulation of incoming packets is not expected to work correctly. VXLAN encapsulation of packets in software is not affected. Removing ingress PACL configuration on the affected interface restores correct VXLAN hardware encapsulation behavior on that interface. (76085)
- Multicast-routed frames get tx-mirrored even when they do not egress the tx-mirrored multicast destination due to VLAN suppression. This is a limitation of the switch ASIC. (87933)
- Enabling counters for Static or Twice NAT connection
  can cause FocalPointV2 agent restarts, when the number of connections exceed 275. (109048)
- On a DCS-7150 series switch, VXLAN encapsulating a packet post routing may result in the outgoing packet being corrupted if the packet was permitted by an ACL with "statistics per-entry" enabled. (294588)
- IPV6 hop-limit based filter in access-list is not supported even if accepted by CLI. (383400)
  Series Affected: DCS-7150 Series
- TX Mirroring is not supported for VXLAN decapsulated and encapsulated traffic and enabling TX mirroring of such traffic may lead to unexpected behavior of mirrored streams. (615841)
- Unidirectional link mode is not supported on 40G ports. (635053)

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)
- VXLAN encapsulated ieee reserved mac addresses are not trapped to CPU and hence
  any control protocol that is VXLAN encapsulated will not work. (423963)
- When sFlow and monitor session are both enabled in the system, the mirrored traffic may be incorrectly throttled to the sFlow copp policy rate. (604224)

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-720XP Series

- "show environment power" does not show insertion status or telemetry information for PWR-1021-AC-RED or PWR-621-AC-RED when there is no input power. (369368)
  SKUs Affected: PWR-1021-AC-RED and PWR-621-AC-RED

- 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)
- When authentication caching is enabled and the link on a switch port is brought down
  due to a device being unplugged, it may take up to 2 minutes for the static mac
  address associated with the port to be removed from the mac address table. This is a
  grace period provided to give time for the device to be re-inserted into the same
  port. (489930)
- 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)
- PTP will not synchronize over interfaces configured with half-duplex link speeds. (740727)
- Phone sending LLDP packets before getting authenticated (by sending non-LLDP packets like DHCP) may not get classified as phones.

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

DCS-7160, 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)

7368, DCS-7300X, DCS-7500R and DCS-7800 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)
- When 1000 VRFs are configured with 4000 SVIs, the Ira agent may experience a SIGQUIT during SSO switchover. There is no workaround for this issue and switchover is ultimately successful. (495529)
- On modular switches configured in SSO redundancy mode, the Tpm agent does not start on the standby supervisor. (677335)

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

- The L3 MTU on interfaces is not enforced on SVIs. (11659)
  Series Affected: DCS-7020, DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series

- Disabling IPv6 routing on an L3 interface is not supported when IPv4 routing is also enabled on that L3 interface. (43093)

- On DCS-7500 series, stateful switchover is not hitless. (54305)

- Linecards with HW revision < 02.00, when configured for 100G and a physical layer problem is present, a one-way link may result. The link experiencing the physical layer problem will be down. However, the peer link may still be up.

  The workaround is to detect this situation, by, for example, looking at the link status on both sides of the 100G link, and perform remedial actions; for example, replacing the transceiver or cable. Running a keepalive protocol like LACP on these links is advised. (67325)
  SKUs Affected: 7500E-12CM-LC and 7500E-72S-LC

- show interfaces counters rates' for port channels may show a slightly inaccurate data rate (including above 100%). (67367)

- Linecards with HW revision < 02.00, under extreme scenarios where we receive linerate small packets on all interfaces connected to the same forwarding ASIC and ingress mirror or sflow dangerous is turned on many of these interfaces, some packets (both the original and the mirror/sFlow copy) can be dropped. (67942)
  SKUs Affected: 7500E-12CM-LC, 7500E-36Q-LC, 7500E-48S-LC and 7500E-72S-LC

- 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)

- VARP and VRRP cannot be enabled at the same time. If VRRP is configured, there can be up to 128 combinations of interface and Virtual Router ID (VRID), but there can only be at most 8 different VRIDs and the VRIDs must differ only in the 3 least significant bits. VRRP for IPv4 and IPv6 cannot be configured at the same time. (75752)
  Series Affected: DCS-7280E and DCS-7500E Series

- 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)
- On the DCS-7280E and DCS-7500E series, configuring QinQ or SelectiveQinQ along with configurable TPID on the same interface is not supported (131757)
- IPv6 Egress ACL deny logging is not supported on subinterfaces. (146487)
- 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)

- Egress IPv4/IPv6 RACL does not work for VXLAN encap-route case. (154551)
- Congestion on an egress interface used by a GRE tunnel that is a mirroring destination can cause drops on forwarding ASICs where the mirrored traffic ingresses the switch. (158001)
- On the DCS-7500E and DCS-7280 series, traffic outage is expected during transition from a single member LAG to two member LAG and vice versa. Use "platform sand lag hardware-only" to avoid this outage. (160439)
- When MLAG peer mac routing feature is enabled on both MLAG peers, the IPv6 neighbor may not be resolved when source ipv6 address in solicited neighbor is link local address. (165067)
- When running PTP with 4X10G break-out cables, a variable delay is observed with the PTP timestamps that can cause instability to the PTP time calculations. (186891)
  SKUs Affected: 7500E-36Q-LC
- When the 'vxlan-routing' hardware TCAM profile is configured, VXLAN traffic flows in overlay may not work alongside IPv6 traffic flows in underlay. The workaround is to use the CLI command 'no platform sand ipv6 host-route exact-match'. (190017)
- L2-ECMP over VXLAN is not supported for multi-homed hosts under EVPN. (193475)
- In Tap Aggregation mode, Tap Aggregation features may not work for 802.1BR/VN tagged packets when 802.1BR/VN tag stripping is not configured. (198798)
- Links are not compatible with systems running EOS 4.18.1F

  The workaround is to upgrade both sides of the link. (199152)
  SKUs Affected: 7500R-8CFPX-LC
- Traceroute to a IP route where the IP route is getting forwarded through a MPLS nexthop-group may incorrectly display the first hop as being unreachable. The actual first hop will appear in the second position. (209273)
- When ingress packet truncation is enabled in Tap Aggregation mode, truncated packets are not counted as packets received on the ingress interface. (215346)
- Frames less than 64 bytes are dropped (as expected),
  but not counted.

  Command "hardware phy cmx42550 <intf > diag read64 mac line 0x138"
  can be used to dump the counter for the purpose of debugging". (218429)
  SKUs Affected: 7500R2M-36CQ-LC, 7800R3-48CQM-LC, DCS-7280CR2M-30-F,

DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R

- Native VLAN setting on Tap interfaces does not take effect. (227770)
- While in Tap Aggregation mode, ingress VLAN membership filtering does not work when packets are destined to a tool interface configured with egress packet truncation. (227841)
- While in Tap Aggregation mode, interfaces configured for egress truncation will continue to consume hardware resources even when they are down. (228384)
- LLDP functionality is not supported on interfaces with ingress truncation enabled. (229983)
- Sflow agent CPU utilization is around 10% even when hardware acceleration is enabled and the agent is not processing any flow samples. (251674)
- With MTU increased beyond 9100B, there can be some multicast packets of size greater than 9100B that are dropped as a port gets close to linerate. (261446)
- Changing PMF profile on one linecard may cause a brief disruption of PMF-based features on other linecards which should not be affected by the same change. (270849)
- A QoS policy, with a class map which has either set-tc or set-dscp on the ingress port and which has DSCP rewrite map on the egress port, results in indeterministic behavior with respect to DSCP remarking of the outgoing packet. (273320)
- If a QOS and PDP Policy both having police action are applied on an interface and the traffic hits both the policies then QOS Policer Counter is messed up. The "show policy-map type qos pmap-name" will give the PDP Policer Counter, not the QOS Policer Counter. (275994)
- Rate values for subinterfaces shown by "show interface counter rates" may deviate up to 15% from corresponding rate on parent interface. (278700)
- With TCP MSS Clamping feature configured, TCP SYN packets might get dropped under high CPU usage conditions if the incoming TCP SYN packet rate is high. (279370)
- Packets larger than 10196 bytes will be discarded by the forwarding ASICs. (280459)
- PFC on DSCP trusted ports is not supported. (281214)
- Loop protect feature is not supported when AlgoMatch HW is enabled. (290706)

- L3 forwarding through SVIs is not support on L2 sub-interfaces. (293410)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- When a port becomes a LAG member, there may be a very brief moment where a few packets may be duplicated on the LAG. (295260)
- Fragmented IP packets will not match on flow-spec rules that include a match on TCP/UDP port. (297309)
- When all the L2 subinterfaces belonging to a parent interface are shutdown, packets are bridged and mac addresses are learnt on parent interface (300202)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- VXLAN routing is not supported on a modular system on which one of the linecards is configured for TAP aggregation. (306180)
- Show interfaces counters ip only show IPv4/Ipv6/MPLS counters per physical interface and not per port-channel (306209)
- When a BGP peer session is cleared, traffic which hits existing flow-spec rules may increment the wrong counter for a brief period of time. (339523)
- If non shaped sub interfaces are configured under same interface having shaped sub interfaces, traffic will egress out of the non shaped sub interfaces in a round robin fashion. It is recommended to not have shaped and unshaped sub interfaces under the same parent. (341851)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Hardware accelerated sflow with profile sflowaccel-v2 doesn't work with 7500E linecards. (342130)
- Stale telemetry flow-spec statistics may be seen until the next "show flow-spec" command is issued. (345195)
- The sFlow sample output interface of IPv6 packet subject to an egress ACL, is set to unknown output interface. (354746)
- Hardware accelerated Sflow does not support collectors reachable over arbitrary tunnel types (354749)
- IPv4 packets with options in the IP header will not be subjected to Unicast Reverse Path Forwarding (uRPF) checks,
  and hence will not be dropped even if the source IP is unknown. (355542)
- unicast RPF is not enforced on IP packets when the destination IP needs ARP resolution (356347)
- Software forwarding is not supported for GRE or IP-in-IP tunnel terminated packets using decap-groups. (358369)

- The egress DSCP value in the Outer IP header for packets that undergo headend
  replication after routing is derived from the Traffic Class (and not copied from the ingress DSCP value). (366189)

- While in Tap Aggregation mode, if the identity tag configuration of a Tap port is changed, traffic received on this port may be sent to an invalid multicast destination momentarily. This will cause the DchUnreachables counter to increment. (372757)

- TX mirroring shows VXLAN encapsulated packets that are actually dropped by split horizon. (375795)

- If more than maximum supported Ipsec tunnels are configured, it is random as to which ones are assigned flow entries within on-chip encrypt/decrypt engine. The set of Ipsec tunnel interfaces assigned flow entries may not be the same as the set of tunnel interfaces which have IKE connections established. As a result, some tunnel interfaces with IKE connection established may not get link up. (376431)

- Configuring more than 250 BFD sessions on a physical interface may result in many of those sessions flapping because of hardware drops on the interface. (389387)

- Drop Precedence based tail drop thresholds can be ignored when multiple queues are congested and buffering resources are low. (395693)

- Queue build-up doesn't show up for 64 byte packets in case of congestion via following commands:-
  1) platform fap diag diag cosq non_empty_queues
  2) show interfaces queue length

  Following command should be used instead:-
  show queue-monitor length (398937)

- PIM sparse mode feature which allows installation of outgoing interfaces on the non-DR for a faster failover does not work on L3 sub-interfaces. This feature relies on an egress ACL to stop multicast traffic from going out of the interface. Multicast egress ACLs do not work on L3 sub-interfaces due to a platform limitation. If the feature is configured, there will be duplication of traffic. The workaround is not use the feature on such interfaces. (405696)

- When a policy-map with match on inner VLAN is programmed on an interface, if single tagged packet traffic is sent to that interface, the behaviour of that packet is indeterministic as inner VLAN tag is calculated using an offset and vlan-tag-format qualifier doesn't differentiate between single and double tagged packets. (415587)

- When MPLS label is pushed on the traffic ingressing CoS trust port then EXP bits are not derived from TC based on packet's COS. (417108)
- Byte counts are not supported for IPSec tunnels. (419031)
- SflowAccel supports up to 200 VRFs. (423892)
- PAUSE frames with correct DMAC=01:80:C2:00:00:01, Ethertype=0x8808, opcode=0x0001 are never forwarded and are always consumed by the hardware. (426047)
- Header truncation for mirrored packets can not be used when the destination of a monitor session is a GRE tunnel. (428547)
- A port transmitting packets smaller than 100B that were tunnel terminated on the switch might be limited to ~90% linerate (429797)
- When the VXLAN routes are programmed using 2 hierarchical FEC levels in hardware, and the outgoing interface of a VXLAN underlay route changes from one
  forwarding chip to another, then it can negatively impact VXLAN overlay route
  convergence. (440233)
- Priority Flow Control pause watchdog actions "errdisable" and "notify-only" actions are not supported. (450857)
- Interface-level Priority Flow Control (PFC) watchdog commands are not supported. (450860)
- Dynamic port-channels coonfigured using LACP are not supported as destinations for monitor sessions (463186)
- 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-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, 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

- Egress MPLS tunnel counters are not supported for implicit null tunnels. (496375)

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

- All asymmetric IRB traffic flows forwarded over EVPN VXLAN toward to a single multi-homed destination are not load balanced across the VTEPs that form the L2
ECMP group. However, note that the traffic destined to different hosts are statically load balanced across the VTEPs. (501343)

- With 'hardware next-hop fast-failover' configuration, the maximum ECMP size supported is reduced to 1023 in R, R2 series and 127 in R3 series swtiches respectively. (502940)

- 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)

- 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)

- `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-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, 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
- "shape rate" configuration is not supported on LAG subinterfaces used for MPLS pseudowires. (555799)
- 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)
- Different streams being rate-limited by the same group policer might exhibit bias towards some streams and fair policing across all the streams is not guaranteed. (578041)
  Series Affected: DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- When all members of a port channel have "forwarding transmit disabled", traffic will not be forwarded to other VLAN members. (587648)
- While in Tap Aggregation mode, traffic steering service policy rules that match on fields from the MPLS header may result in false matches on packets without an MPLS header.

Remove the MPLS key-fields from feature 'tapagg mac' in the active TCAM profile. (592184)

- While in Tap Aggregation mode, traffic steering service 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)

- 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)

- 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)

- IPv4 and IPv6 strict mode Unicast Reverse Path Forwarding (uRPF) does not work if the route matching source IP address resolves via an ECMP nexthop. Such packets will be dropped. (757793)

- With scale QoS scale policy-map config with ACL rules, configure-replace might CLI time-out (766811)

- For tunnel destinations, the decision to fragment is based on unencapsulated packet size, while MTU enforcement is based on the encapsulated packet size. This may cause unencapsulated packets near the MTU packet size to be dropped instead of fragmented. (768410)

- At high combined scale of VLANs and pseudowires in VPLS or L2 EVPN, exhaustion of "managedTt inLif" resources may occur. This may prevent hardware programming of any configured patch panel patches, resulting in "Unprogrammed" status for one or more patchs. It may also result in VPLS packets drops in the decapsulation direction. This situation can be detected by "show hardware capability".

  The workaround is to reduce the scale of these or other features which use

the managedTt inLif resource and then perform a shut/no shut on the affected patch panel or subinterface. (781993)

- "ip hardware fib next-hop sharing" is not supported with VXLAN. (795334)

DCS-7500E Series

- With a delay ECN threshold configured, some packets (under 10 %) experiencing valid
measurable delays greater than the threshold may not be ECN marked. (141587)
- With a delay ECN threshold configured, packets experiencing delays of up to 2 microseconds more than the configured delay may not be ECN marked. (141588)

DCS-7280R and DCS-7500R Series

- When a linecard is inserted in a chassis there can be a temporary marginal drop in fabric throughput for other cards. (205880)
- On the system with DCS-7500R linecards, if DCS-7500R2 linecards are inserted, routes may consume more resources in the hardware routing table. (215393)
- Jericho chip buffering DRAM test at agent startup does not catch bad parts while it should fail to force RMA of bad parts. (218971)
- IPv6 multicast routes are stored in two separate TCAM resources. The group IP is stored in one TCAM while the source IP is stored in another. Since one TCAM contains the source IP it will contain an entry for every multicast route. However, the group IP TCAM will reuse entries for duplicate group IPs. For example, if you have routes (S1,G1) and(S2,G1), one TCAM entry is required for the group IP and two TCAM entries are required for the source IP. This means it is possible that more TCAM resources are allocated to the source IP lookup than the group IP lookup. The scaling achieved during the initial distribution of multicast routes will be preserved. However, if the distribution later changes, optimal scaling for the new distribution may not be reached. (257364)
- The CPU traffic policy feature does not support filtering on
IPv6 extension headers. (364996)
- Ingress ACLs and PBR are not supported for MPLS tunnel terminated traffic using VRF label. (366695)
- VRF terminated traffic is not sampled with sFlow. (369010)
- Drop-precedence thresholds are not supported on mirror session ports. (388932)

- We do not support incoming multicast LDP packets with a destination multicast MAC address.
  RFC 5332 doesn't mandate the use of sending labeled multicast traffic in a multicast ethernet frame. Multicast MAC destination address could be used in upstream label assignment. (393406)
  Series Affected: DCS-7020, DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series

- MTU is not enforced for MPLS packets forwarded by a MPLS swap route, which have a length one byte more than MTU. (464742)

- SandTcam crashes while changing the system TCAM profile to a profile that has feature 'tcp-mss-ceiling ip' on switches that has some linecards using Algomatch as the access-list mechanism and some other linecards using TCAM as access-list mechanism.

  Workaround : Configure 'hardware access-list mechanism tcam' (474595)

- If the CPU is included in monitor sessions with multiple destinations, it will only work if the device mirror0 is selected. If multiple mirroring to CPU sessions with multiple destinations are configured, all packets will be sent to mirror0. (525919)

- With RFC5549 ACLs configurations, the Egress IPv6 ACL deny logging on port channel subinterfaces stops working on performing agent SandL3Unicast restart. (549181)

- MPLS explicit NULL termination isn't supported with flex-route configuration containing non-nibble aligned prefix length. (576801)

- If MPLS tunnels are allocated in uncountable egress banks due to transiently high scale, they can be left in an uncountable bank which will prevent traffic counters from being aggregated for these tunnels.
  As a workaround, a hitfull restart of the forwarding plane agent will reallocate the tunnels in countable egress banks 1 & 2. (635152)

DCS-7280R2 and DCS-7500R2 Series

- On platforms with Jericho+ switches and with large number of pseudowire patches (2500 or more) configured, disabling hierarchical FECs (HFECs) using CLI command "ip hardware fib hierarchical next-hop disabled" may leave some patches unprogrammed and not operational.

  Workaround is to save config with HFEC disabled to startup config and reload the switch. (569868)

- "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)
- IP egress ACLs are not supported for MPLS terminated packets. (383263)
- 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)
- Applying a Qos policy on Vlan interfaces may result in SandAcl restart, since this functionality isn't supported. (403280)
- Applying a Qos policy on sub-interfaces may result in SandAcl restart, since this functionality isn't supported. (403281)
- Traffic drop may happen when we sent traffic to all 8 traffic-classses (aggegrate rate nearing linerate) to port which has been configued with Strict-Priority scheduling. (406550)
- 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)
- "ip tunnel <ip> nexthop-group" configuration for nexthop-groups with type mpls consisting of direct IP nexthops without label encapsulation is not supported. (458065)
- Enabling Tx mirroring may cause the counter VOQ_SRAM_ENQ_RJCT_PKTS to increment, with reason QnumNotValid, as shown in the output from "show platform fap counters pipeline". There is no impact to mirrored or forwarded traffic. (459862)
- Unidirectional links are not supported at 40G speed on QSFP100 interfaces using the CRT50216 PHY. (467672)
  SKUs Affected: 7368-16C, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQM-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)
- 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 and DCS-7280CR3MK-32P4
- Encapsulating PTP packets over VXLAN tunnel is not supported. (526461)
- 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)
- UDP packets with ports 4754 or 4755 may not be forwarded correctly. (648240)

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

- L4 ops and fragments rule cannot be programmed together. (94197)

- While in Tap aggregation mode enabling either VN or BR tag stripping will disable VXLAN header stripping functionality. Traffic steering on the inner IP header of VXLAN packets will also be disabled.

  Removing the VN/BR tag stripping configuration will restore VXLAN stripping and traffic steering functionality. (175829)

- An IPv6oMPLS packet terminated by L3 EVPN MPLS configuration will not be forwarded. The workaround is to disable IPv6 exact match host routes using 'no platform sand ipv6 host-route exact-match' and disable IPv6 route optimizations for prefix length 128. (201414)
  Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series

- Vxlan decapped packets will not match rules that match on 'tcp established' in egress acls (226582)

- VXLAN and GRE tunnels cannot be configured at the same time. (248239)

- If a PTP packet is mirrored to a GRE destination which has time stamping enabled, the timestamps on such packets are not reliable. (280362)

- If a packet is mirrored to a GRE tunnel, but the forward copy is trapped to the CPU, then the ingress VLAN ID carried in the GRE key will be set to 0 when the mirrored copy is encapsulated in the GRE envelop. (434871)

- Mirroring to GRE destinations will not work if VxLAN encapsulation is enabled (548429)

- The "VLAN Editing Extended-64" and "ACL Port IP Egress" TCAM features are mutually exclusive. (554456)

- MPLS label popped packets will not match rules that match on 'tcp established' in egress IP acls (565015)

- L4 destination port matching in egress IP ACLs does not work on untagged bridged traffic. (609461)
- 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)
- 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)
- L2 EVPN MPLS configuration will not work when the same trunk port is used to carry traffic towards both the Provider Edge (PE) and the Customer Edge (CE). (727163)
- When IP PIM sparse mode is enabled on the VLAN interface, egress port ACLs applied to switch ports that are a member of that VLAN match on random fields for bridged multicast traffic. (728267)
- Egress IPv4 ACLs do not apply to unknown multicast traffic (728269)
- QoS policy-maps are not supported on dot1ad, unmatched or untagged subinterfaces. (733077)

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

- Priority tagged packets (802.1Q tagging with the VLAN set to 0) being reflected by an Ethernet reflector interface configured in the egress direction with MAC address swap action will have the priority tag stripped from the packet before they are transmitted. (354190)
- Only port-channel interfaces are supported as peer links. Single member port-channel interfaces are supported. (409958)
- If a packet that is decapsulated with a decap group is mirrored to a GRE tunnel, the packet's ingress VLAN will be set to 0 in the upper 2 bytes of the GRE key. (430986)

- QoS policy-map is not supported on Pseudowire decapsulated traffic flows. (574600)
- Minor variance between configured and observed bandwidth on transmit queue may be seen on 400G ports configured with Weighted Round Robin scheduling. (580294)
- 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)
- IP UDFs in ACLs may not work correctly on IPv6 packets with a routing extension header when matching on bits beyond the IPv6 header. (650286)
- VXLAN is not supported with greater than 256-way ECMP in the underlay. (652266)
- While not in Tap Aggregation mode, packets mirrored to multiple destination ports may be actually forwarded out of the mirroring destinations, or may be malformed.

  Do not use mirroring to multiple destinations outside of Tap Aggregation mode. (670046)
- The MPLS no-php mode is not supported with L3VPN and VPWS services. (720213)
- Traffic throughput on an internal recirculation interface is capped at 100Gbps. (738474)

DCS-7280R and DCS-7280R2 Series

- Flex-route is not supported with security ACLs (Port ACLs, Router ACLs & MAC ACLs) when AlgoMatch is enabled. The workaround is to configure 'hardware access-list mechanism tcam'. (348149)
- CBF does not work in AlgoMatch mode on some custom TCAM profiles (603206)

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

- Untagged packets received on a port with PFC enabled may be assigned the wrong priority. In case of congestion, PFC frames may not be sent or may be malformed.

  The issue is observed when a port uses a default CoS value that is mapped to a traffic class with a different value. For example, if the default CoS is 0 and CoS 0 is mapped to traffic class 1. The workaround is to change the default CoS value on the port or the global QoS mapping such that the default CoS and the traffic class match. (23922)

- System does not stop transmitting traffic for the exact amount of time quanta received in PAUSE or PFC frame. This can lead to packet loss during congestion if the peer switch does not have enough buffers. The workaround is to use XON/XOFF instead of time based flow control, or to increase the time to account for one MTU sized packet. (27190)

- QoS shaping and guaranteed bandwidth will only work in store-and-forward mode. They are not supported in cut-through mode. This includes 'shape rate X' and the 'bandwidth guaranteed X' commands. (56790)

- Accelerated Software Upgrade (ASU) from 4.13.x (x < 6) to 4.13.6 and above requires a special upgrade procedure due to a software image format change. (90097)

- 10 Gbps packet replication on all SFP+ ports at the same time may lead to packet discard. (91149)

- Egress L2 MTU violations in cut-through mode may result in unexpected discards of other frames. (95577)

- VXLAN encapsulated packets cannot be TX mirrored at the egress properly. The mirrored packets will have an incorrect MAC header. (97272)

- 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-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, DCS-7050CX3, DCS-7050TX3, DCS-7050X3, 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)
- Cut-through mode is not supported on SFP+ interfaces. (252731)
  Series Affected: DCS-7050X3, DCS-7060X4 and DCS-7260X3 Series
- A switch may not be able to bridge or route VXLAN traffic if it hits one or all of following conditions:
  - Nexthop resources are exhausted as indicated by syslog VXLAN_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 and DCS-7050SX3-48YC8-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
- Routed unicast traffic egressing a SVI for which destination MAC is not learnt are sent to CPU and then flooded on the egress VLAN. (368228)
  Series Affected: 7368 and DCS-7060X4 Series
- SSU is not supported with Vxlan and MLAG config. (388669)
- 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 an interface undergoes a hitless speed change, sFlow sample numbers and sequence counts for that interface may be reset. (452949)
  Series Affected: 7300X3, 7368, DCS-7050X3, DCS-7060X4 and DCS-7260CX3 Series
- BUM traffic is not sampled by egress sFlow. (459902)
- 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 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)
- 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)
- Packets sent from the CPU cannot be egress mirrored unless they are being sent out of an SVI. (495146)
- 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)
- 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)
- 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-7060SX2, DCS-7060X and DCS-7060X2 Series
- The traffic load across ports comprising a port channel(s) may be uneven for certain traffic patterns and/or port channel configuration and state.

  Workaround:
  - change the seed in the hashing algorithm until satisfactory result is achieved by executing configuration CLI command:
    "arista(config)# port-channel load-balance trident <seed>" (555833)
- Implicit ICMPv6 permit rule(s) in TCAM is not seen in output of "show platform trident tcam acl detail" command.

  Workaround is to check the entry and packet hits for the entry using output from "show platform trident tcam detail". This command shows all ACL entries in TCAM and packet hits per entry. (563228)
- Enabling L2 protocol forwarding of LLDP on a given port will implicitly enable forwarding of EAPOL and MKA packets on the same. (569478)
- BGP packets are ignored by IP counters if hardware IP counters are not enabled. (599275)
- 7304X3-FM and 7304X3-FM2 are not guaranteed to interoperate within the same chassis. 7308X3-FM and 7308X3-FM2 are not guaranteed to interoperate within the same chassis. The system will not emit any warnings. (621072)
  SKUs Affected: 7304X3-FM, 7304X3-FM2, 7308X3-FM and 7308X3-FM2

- Reaching a VLAN's configured limit for locally learned dynamic MACs will not generate a syslog message. (631035)
- The "show tech-support" command does not contain output from the "show mac address-table dynamic local limit" command. (631039)
- 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)
- Mirroring packets to CPU can cause congestion on CPU queues and disrupt control plane traffic. It is not recommended to leave this feature on for extended periods of time. (675773)
- 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-720XP, DCS-7050CX3 and DCS-7260CX3 Series
- Support for configuring MTU over 9214 in a config session. (713459)
- 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)
- If the same ACL is applied on control plane and data plane, the byte counter for the ACL will be incremented for only data plane packet hits while the packet counter will be incremented for both data plane and control plane packet hits.

A workaround is to create different ACLs to apply on control plane and data plane. (788296)

DCS-7010 Series

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

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

- When in cut-through forwarding mode, mirror sessions with both ingress mirror sources and egress mirror sources configured to mirror packets to the same mirror destination port will not be programmed in hardware.

  The workaround for this is to switch to store-and-forward mode where the limitation is not present. (154721)

- Cut-through mode is not supported on the SFP+ interfaces. (182728)

- Cut-through mode is not supported at 1G speed. (219427)

- Mirror To GRE destinations with 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 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)

- show running-config incorrectly displays interfaces for non-inserted cards. Interface statuses for missing or powered down cards are displayed as inactive. (306105)
  Series Affected: 7368 Series

- Layer2 source MAC learning may occur on packets received with CRC errors. (325507)

- Per vlan routing and sub interface are not supported together. (486439)
- Grpc buffer-usage counters are not supported on Tomahawk3 platforms, and can cause forwarding agent crashes if configured (593626)
- 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)

CCS-720XP Series

- On interfaces with speeds less than 10G, PFC cannot be transmitted. (689087)

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-O-2Q-400G and CAB-O-4Q-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)
- SFP 10GBASE-T transceivers which do not support Soft RX_LOS or Soft TX_DISABLE are not supported on DSFP ports. (702025)
  SKUs Affected: SFP-10G-T

7300X3, CCS-720XP and DCS-7050X3 Series

- The DirectFlow "set vlan action" command is not supported. (237046)
- When configuring a DirectFlow action, specifying "interface hardware-loopback" as the destination of "action output" is not supported. (243031)
- IPv6 packets with WESP payload (next header 141) bypass the IPv6 security ACLs and service policies. (243387)
- Directflow configuration with "set vlan", but no "action set output interface" configured is not supported. (345852)
- NAT is not supported on subinterfaces (353236)
  SKUs Affected: 7300X3-32C-LC, 7300X3-48YC4-LC, DCS-7050CX3-32S and DCS-7050SX3-48YC12
- 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 number of MPLS next-hop groups configured exceeds 512, 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 MPLS next-hop groups configured to stop the agent from continuously restarting. (428162)
- 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
- 50g port speed is not supported in this release (445999)
  SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- 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)
- "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-7050TX3 and DCS-7050X3 Series
- DirectFlow actions on interfaces configured as dot1q-tunnel are not supported. (686339)
- Ingress mirroring double tagged packets to CPU will have the outer VLAN tag stripped.
  Workaround is to use a port mirroring destination instead. (722891)
- Packets with Broadcast source MAC may be copied to CPU instead of being dropped. (745001)
- Maximum number of replications is 2047 per packet. (775632)

DCS-7160 Series

- If there is a port ACL (PACL) or router ACL (RACL) and if the incoming packet has
  IP options with L4 TCP/UDP/ICMP header, then those packets are dropped. (172932)

- No support for matching on DSCP field on IP packets for security ACLs. (174114)
- If the MAC address of the next-hop to reach a remote VTEP is not learned, and if the remote VTEP is configured for Head End Replication (HER) for a VXLAN VLAN, then, the the VTEP is not included in the list of HER VTEPs and no VXLAN BUM packets will be sent to that VTEP. The VTEP is included in the list of HER VTEPs once the underlay MAC address of the next-hop to reach that VTEP is learned.

  The workaround is to ensure that the underlay MAC of the next-hop to reach remote VTEPs are always present (which is the normal operation). (178395)
- Control plane Policy Map counters are not supported. (180303)
- A maximum of 3840 remote VTEPS are supported when the switch is configured as a Vxlan Vtep. (188553)
- Any IPv6 ACL containing a rule with the first 96 bits set as 0 cannot be configured when the algomatch profile is being used. (218453)
- Customer may witness around 100 milliseconds of packet loss when VxLAN ECMP nexthop changes. (219888)
- There may be traffic loss of up to 100 milliseconds when a VxLAN nexthop (as part of
  ECMP)undergoes modifications. (268600)
- The DSCP value of BGP control plane packets set by the command 'bgp transport qos
  dscp <value>' is overwritten when 'qos rewrite dscp' is enabled and the outgoing
  routed interface is a SVI. The rewritten value is based on traffic-class to DSCP
  map configuration. (288617)
- If a different XP profile than the one that is currently running is configured, and ACL
  configuration/ACL viewing commands without running a chip reset/reboot (via platform xp xp0 reset),
  the command can cause XpAcl to restart. (418083)
- VLAN ID match in IP ACLs is not supported on DCS-7160 platforms. (514366)
- If an incoming non-routable VXLAN-encapsulated packet is decap-eligible, but the encapsulation does not conform to VARP VTEP - VARP MAC binding, the packet will not be dropped and will continue to be switched based on the inner packet header. The impact is limited to extra or duplicate packets in the destination bridging domain, and does not affect packet forwarding

correctness. However, it indicates an issue with the sender switch which is potentially generating illegal VXLAN-encapsulated packets. (562335)

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.