

Arista Networks EOS 4.22.8.3M Release Notes

Version: 0.2

09 August 2022

Table of Contents

EOS 4.22.8.3M Release Highlights

New Platforms and Hardware

SNMP MIBs

SNMP Traps

Supported Hardware

Reference to MLAG ISSU Upgrade/Downgrade Table

Resolved Caveats in 4.22.8.3M

Accelerated System Update

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

Known Software Caveats in 4.22.8.3M

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

Containerized EOS

vEOS Router / CloudEOS

CVP

Network Provisioning

CVX

General

IPSEC

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

MPLS

Multicast

NAT

vEOS Router / CloudEOS

DCS-7150 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

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

DCS-7500E Series

DCS-7280R and DCS-7500R Series

DCS-7280R2 and DCS-7500R2 Series

DCS-7280R3 and DCS-7500R3 Series

DCS-7020 Series

DCS-7500R3 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-7060, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260X, DCS-7260X3 and
DCS-7300X Series

DCS-7010 Series

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

DCS-7260X3 Series

7368 Series

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

DCS-7050X2 Series

CCS-720XP Series

DCS-7050X3 Series

Transceiver Series

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

DCS-7160 Series

Limitations and Restrictions in 4.22.8.3M

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

vEOS Router / CloudEOS

CVP

Telemetry

CVX

General

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

MPLS

Multicast

SwitchApp

vEOS Router / CloudEOS

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

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

DCS-7500E Series

DCS-7280R and DCS-7500R Series

DCS-7280R2 and DCS-7500R2 Series

DCS-7280R3 and DCS-7500R3 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-7060, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260X, DCS-7260X3 and
DCS-7300X Series

DCS-7010 Series

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

DCS-7260X3 Series

7368 Series

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

DCS-7050X2 Series

CCS-720XP Series

Transceiver Series

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

DCS-7160 Series

Conditions and Impacts

EOS 4.22.8.3M Release Highlights

Note : The release type was incorrectly set as "F" instead of "M". No new features were added in this release.

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

Supported Hardware

Fixed System

- | | | |
|-------------------|----------------------|-------------------|
| • DCS-7150S-24-CL | • DCS-7280QR-C36-M-R | • DCS-7050SX-64-F |
| • DCS-7150S-24 | • DCS-7280QR-C36-R | • DCS-7050SX-64-R |

• DCS-7150S-24-CL-F	• DCS-7280SR-48C6	• DCS-7050SX-64-SSD-F
• DCS-7150S-24-CL-R	• DCS-7280SR-48C6-F	• DCS-7050SX-64-SSD-R
• DCS-7150S-24-CL-SSD-F	• DCS-7280SR-48C6-M-F	• DCS-7050TX-128-F
• DCS-7150S-24-F	• DCS-7280SR-48C6-M-R	• DCS-7050TX-128-R
• DCS-7150S-24-R	• DCS-7280SR-48C6-R	• DCS-7050TX-48
• DCS-7150S-24-CL-SSD-R	• DCS-7280TR-48C6-R	• DCS-7050TX-48-F
• DCS-7150S-52-CL	• DCS-7280TR-48C6-F	• DCS-7050TX-48-R
• DCS-7150S-52-CL-F	• DCS-7280TR-48C6-M-F	• DCS-7050TX-48-SSD-F
• DCS-7150S-52-CL-R	• DCS-7280TR-48C6-M-R	• DCS-7050TX-48-SSD-R
• DCS-7150S-52-CL-SSD-F	• DCS-7280QR-C72	• DCS-7050TX-64
• DCS-7150S-52-CL-SSD-R	• DCS-7280QR-C72-F	• DCS-7050TX-64-F
• DCS-7150S-64-CL	• DCS-7280QR-C72-M-F	• DCS-7050TX-64-R
• DCS-7150S-64-CL-F	• DCS-7280QR-C72-M-R	• DCS-7050TX-64-SSD-F
• DCS-7150S-64-CL-R	• DCS-7280QR-C72-R	• DCS-7050TX-64-SSD-R
• DCS-7150S-64-CL-SSD-F	• DCS-7280SR2-48YC6	• DCS-7050TX-72-SSD-F
• DCS-7150S-64-CL-SSD-R	• DCS-7280SR2-48YC6-F	• DCS-7050TX-72-SSD-R
• DCS-7150SC-24-CLD	• DCS-7280SR2-48YC6-M-F	• DCS-7050TX-96-SSD-F
• DCS-7150SC-24-CLD-F	• DCS-7280SR2-48YC6-M-R	• DCS-7050TX-96-SSD-R
• DCS-7150SC-24-CLD-R	• DCS-7280SR2-48YC6-R	• DCS-7010T-48
• DCS-7150SC-64-CLD	• DCS-7020TR-48	• DCS-7010T-48-F
• DCS-7150SC-64-CLD-F	• DCS-7020TR-48-F	• DCS-7010T-48-R
• DCS-7150SC-64-CLD-R	• DCS-7020TR-48-R	• DCS-7050SX-72
• DCS-7170-64C	• DCS-7280QRA-C36S-F	• DCS-7050SX-72-F
• DCS-7170-64C-F	• DCS-7280QRA-C36S-R	• DCS-7050SX-72-R
• DCS-7170-64C-R	• DCS-7280SR2A-48YC6	• DCS-7050SX-72-SSD-F
• DCS-7170-32C	• DCS-7280SR2A-48YC6-F	• DCS-7050SX-72-SSD-R
• DCS-7170-32C-F	• DCS-7280SR2A-48YC6-M-F	• DCS-7050SX-96
• DCS-7170-32C-M-F	• DCS-7280SR2A-48YC6-M-R	• DCS-7050SX-96-F
• DCS-7170-32C-M-R	• DCS-7280SR2A-48YC6-R	• DCS-7050SX-96-R
• DCS-7170-32C-R	• DCS-7280SRA-48C6	• DCS-7050SX-96-SSD-F
• DCS-7170-64C-M-F	• DCS-7280SRA-48C6-F	• DCS-7050SX-96-SSD-R
• DCS-7170-64C-M-R	• DCS-7280SRA-48C6-M-F	• DCS-7050TX-128-SSD-F
• DCS-7170-32CD	• DCS-7280SRA-48C6-M-R	• DCS-7050TX-128-SSD-R
• DCS-7170-32CD-F	• DCS-7280SRA-48C6-R	• DCS-7050TX-72-F
• DCS-7170-32CD-R	• DCS-7280TRA-48C6	• DCS-7050TX-72-R
• FAN-7001C	• DCS-7280TRA-48C6-F	• DCS-7050TX-96-F
• PWR-1100AC	• DCS-7280TRA-48C6-M-F	• DCS-7050TX-96-R
• PWR-1100AC-F	• DCS-7280TRA-48C6-M-R	• DCS-7010T-48-DC
• PWR-1100AC-R	• DCS-7280TRA-48C6-R	• DCS-7010T-48-DC-F
• PWR-750AC	• DCS-7020TRA-48	• DCS-7010T-48-DC-R
• PWR-750AC-F	• DCS-7020TRA-48-F	• DCS-7260QX-64
• PWR-750AC-R	• DCS-7020TRA-48-R	• DCS-7260QX-64-F
• PWR-500-DC	• DCS-7280CR2A-60	• DCS-7260QX-64-R
• PWR-500-DC-F	• DCS-7280CR2A-60-F	• DCS-7260QX-64-SSD-F
• PWR-500-DC-R	• DCS-7280CR2K-60	• DCS-7260QX-64-SSD-R
• PWR-500AC	• DCS-7280CR2K-60-F	• DCS-7050SX-72Q
• PWR-500AC-F	• DCS-7280SRAM-48C6	• DCS-7050SX-72Q-F

- PWR-500AC-R
- PWR-745AC
- PWR-745AC-F
- PWR-745AC-R
- FAN-7010
- PWR-1900-DC
- PWR-1900-DC-F
- PWR-1900-DC-R
- PWR-1900AC
- PWR-1900AC-F
- PWR-1900AC-R
- PWR-747AC
- PWR-747AC-F
- PWR-747AC-R
- FAN-7002H
- FAN-7002H-R
- FAN-7000H-R
- PWR-1600AC
- PWR-1600AC-F
- FAN-7001D
- FAN-7001D-F
- PWR-400-DC-BLUE
- PWR-400-DC-RED
- FAN-7001C-F
- PWR-400AC-BLUE
- FAN-7011M
- FAN-7011M-F
- FAN-7011M-R
- PWR-1021-AC-RED
- PWR-511-AC-BLUE
- PWR-511-AC-RED
- PWR-621-AC-RED
- FAN-7001C-R
- PWR-400AC-RED
- FAN-7011H
- FAN-7011H-F
- FAN-7011H-R
- PWR-511-DC-BLUE
- PWR-511-DC-RED
- PWR-1511-AC-BLUE
- PWR-1511-AC-RED
- PWR-1511-DC-BLUE
- PWR-1511-DC-RED
- PWR-501AC
- DCS-7280SRAM-48C6-F
- DCS-7280SRAM-48C6-R
- DCS-7280CR2-60
- DCS-7280CR2-60-F
- DCS-7280CR2A-30
- DCS-7280CR2A-30-F
- DCS-7280CR2K-30
- DCS-7280CR2K-30-F
- DCS-7280QRA-C36S-M-F
- DCS-7280QRA-C36S-M-R
- DCS-7280SR2K-48C6
- DCS-7280SR2K-48C6-M-F
- DCS-7280SR2K-48C6-M-R
- DCS-7020SR-24C2
- DCS-7020SR-24C2-F
- DCS-7020SR-24C2-R
- DCS-7280CR2M-30
- DCS-7280CR2M-30-F
- DCS-7280SRM-40CX2
- DCS-7280SRM-40CX2-F
- DCS-7280SRM-40CX2-R
- DCS-7020SR-32C2-F
- DCS-7020SR-32C2-R
- DCS-7020SRG-24C2
- DCS-7020SRG-24C2-F
- DCS-7020SRG-24C2-R
- DCS-7280CR3-32P4
- DCS-7280CR3-32P4-F
- DCS-7280CR3-32P4-M-F
- DCS-7280CR3-32P4-M-R
- DCS-7280CR3-32P4-R
- DCS-7280CR3K-32P4
- DCS-7280CR3K-32P4-F
- DCS-7280CR3K-32P4-R
- DCS-7280CR3-32D4
- DCS-7280CR3-32D4-F
- DCS-7280CR3-32D4-M-F
- DCS-7280CR3-32D4-M-R
- DCS-7280CR3-32D4-R
- DCS-7280CR3K-32D4
- DCS-7280CR3K-32D4-F
- DCS-7280CR3K-32D4-R
- DCS-7280PR3-24
- DCS-7280PR3-24-F
- DCS-7050SX-72Q-R
- DCS-7050TX-72Q
- DCS-7050TX-72Q-F
- DCS-7050TX-72Q-R
- DCS-7050QX2-32S
- DCS-7050QX2-32S-F
- DCS-7050QX2-32S-R
- DCS-7050SX2-128
- DCS-7050SX2-128-F
- DCS-7050SX2-128-R
- DCS-7050TX2-128
- DCS-7050TX2-128-F
- DCS-7050TX2-128-R
- DCS-7050SX2-72Q
- DCS-7050SX2-72Q-F
- DCS-7050SX2-72Q-R
- DCS-7060CX-32S
- DCS-7060CX-32S-F
- DCS-7060CX-32S-R
- DCS-7060CX2-32S
- DCS-7060CX2-32S-F
- DCS-7060CX2-32S-R
- DCS-7260CX-64
- DCS-7260CX-64-F
- DCS-7260CX-64-R
- DCS-7260CX-64-SSD-F
- DCS-7260CX-64-SSD-R
- DCS-7060SX2-48YC6
- DCS-7060SX2-48YC6-F
- DCS-7060SX2-48YC6-R
- DCS-7260CX3-64
- DCS-7260CX3-64-F
- DCS-7260CX3-64-R
- DCS-7050CX3-32S-F
- DCS-7050CX3-32S-R
- DCS-7050SX3-48YC12
- DCS-7050SX3-48YC12-F
- DCS-7060CX-32S (HwRev: 13.20)
- DCS-7060CX-32S-F (HwRev: 13.20)
- DCS-7060CX-32S-R (HwRev: 13.20)
- CCS-720XP-24ZY4
- CCS-720XP-24ZY4-F
- CCS-720XP-48ZC2
- CCS-720XP-48ZC2-F

- PWR-501AC-F
- PWR-501AC-R
- FAN-7000
- FAN-7000-F
- FAN-7000-R
- PWR-460AC
- PWR-460AC-F
- PWR-460AC-R
- PWR-460DC
- PWR-460DC-F
- PWR-460DC-R
- DCS-7280TR-48C6
- DCS-7280QRA-C36S
- DCS-7020SR-32C2
- DCS-7280SE-64
- DCS-7280SE-64-F
- DCS-7280SE-64-R
- DCS-7280SE-72
- DCS-7280SE-72-F
- DCS-7280SE-72-R
- DCS-7280SE-68
- DCS-7280SE-68-F
- DCS-7280SE-68-R
- DCS-7280CR-48
- DCS-7280CR-48-F
- DCS-7280QR-C36
- DCS-7280QR-C36-F
- DCS-7280QR-C36-M-F
- DCS-7280PR3-24-M-F
- DCS-7280PR3K-24
- DCS-7280PR3K-24-F
- DCS-7050TX-128
- DCS-7050TX-72
- DCS-7050TX-96
- DCS-7050CX3-32S
- DCS-7050SX-128
- DCS-7050SX-128-F
- DCS-7050SX-128-R
- DCS-7050SX-128-SSD-F
- DCS-7050SX-128-SSD-R
- DCS-7250QX-64
- DCS-7250QX-64-F
- DCS-7250QX-64-R
- DCS-7250QX-64-SSD-F
- DCS-7250QX-64-SSD-R
- DCS-7050QX-32
- DCS-7050QX-32-F
- DCS-7050QX-32-R
- 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-7050SX-64
- DCS-7050SX3-48YC8
- DCS-7050SX3-48YC8-F
- DCS-7050SX3-48YC8-R
- DCS-7260CX3-64E
- DCS-7260CX3-64E-F
- DCS-7260CX3-64E-R
- CCS-720XP-24Y6
- CCS-720XP-24Y6-F
- CCS-720XP-48Y6
- CCS-720XP-48Y6-F
- DCS-7050CX3-32S-SSD-F
- DCS-7050CX3-32S-SSD-R
- DCS-7160-32CQ
- DCS-7160-32CQ-F
- DCS-7160-32CQ-R
- DCS-7160-48YC6
- DCS-7160-48YC6-F
- DCS-7160-48YC6-R
- DCS-7160-48TC6
- DCS-7160-48TC6-F
- DCS-7160-48TC6-R
- FAN-7002
- FAN-7002-F
- FAN-7002-R
- FAN-7002H-F
- FAN-7000H
- FAN-7000H-F

Modular

- FAN-7002
- FAN-7002-F
- FAN-7002-R
- FAN-7002H-F
- FAN-7000H
- FAN-7000H-F
- PWR-3K-AC
- PWR-3K-AC-F
- PWR-3K-AC-R
- PWR-2700-DC
- PWR-2700-DC-F
- PWR-2700-DC-R
- PWR-3K-DC-BLUE
- PWR-3K-DC-RED
- PWR-3KT-AC-BLUE
- PWR-3KT-AC-RED
- DCS-7508N
- 7500E-6CFPX-LC
- DCS-7500-SUP2
- DCS-7500-SUP2-D
- 7500R-36CQ-LC
- 7500R-36Q-LC
- 7500R-48S2CQ-LC
- 7504R-FM
- 7508R-FM
- 7512R-FM
- DCS-7512N
- 7500RM-36CQ-LC
- 7500R-8CFPX-LC
- 7500R2-18CQ-LC
- 7500R2-36CQ-LC
- 7516R-FM
- 7500R3-24P-LC
- DCS-7808-CH
- DCS-7804-CH
- DCS-7508
- DCS-7504
- 7300-SUP
- 7300-SUP-D
- 7300X-32Q-LC
- 7300X-64S-LC
- 7300X-64T-LC
- 7304X-FM
- 7308X-FM
- DCS-7304
- DCS-7308
- 7316X-FM
- DCS-7316

- PWR-2900AC
- 7500E-36Q-LC
- 7500E-48S-LC
- 7500E-72S-LC
- 7504E-FM
- 7508E-FM
- DCS-7500E-SUP
- DCS-7500E-SUP-D
- 7500E-12CM-LC
- 7500E-6C2-LC
- DCS-7504N
- 7500E-48T-LC
- 7500E-12CQ-LC
- DCS-7516-SUP2
- DCS-7516-SUP2-D
- DCS-7516N
- 7500R2A-36CQ-LC
- 7500R2AK-36CQ-LC
- 7500R2AK-48YCQ-LC
- 7500R2M-36CQ-LC
- 7500R2AM-36CQ-LC
- 7500R3-36CQ-LC
- 7504R3-FM
- 7508R3-FM
- 7500R3K-36CQ-LC
- 7320X-32C-LC
- 7324X-FM
- 7328X-FM
- 7300X3-32C-LC
- 7304X3-FM
- 7308X3-FM
- 7300X3-48YC4-LC
- 7368
- 7368-16C
- 7368-SUP
- 7368-SUP-D
- 7368X4-SC

Transceiver

- 1000BASE-BX10
- 1000BASE-BX10-D
- 1000BASE-BX10-U
- 40GBASE-AOC
- 10GBASE-LRL
- 10GBASE-ZR
- 40GBASE-PLRL4
- 40GBASE-SR4
- 40GBASE-XSR4
- 40GBASE-LR4
- 40GBASE-LRL4
- 40GBASE-PLR4
- CFP2-100GBASE-LR4
- 10GBASE-AOC
- 40GBASE-UNIV
- 100GBASE-CR4
- 100GBASE-LR4
- 100GBASE-LRL4
- 40GBASE-ER4
- 10GBASE-DWDM
- 10GBASE-DWDM-T
- 10GBASE-DWDM-ZR
- CFP2-100GBASE-ER4
- 100GBASE-PSM4
- 100GBASE-SR4
- 40GBASE-SRBD
- 100GBASE-AOC
- CFP2-100GBASE-XSR10
- 100GBASE-CWDM4
- 25GBASE-AOC
- 25GBASE-CR
- 25GBASE-LR
- 25GBASE-SR
- 100GBASE-SRBD
- 100GBASE-SWDM4
- 100GBASE-ERL4
- 100GE-DWDM2
- CFP2-100G-DWDM-DCO
- CFPX-100G-DWDM
- CFPX-200G-DWDM
- 1000BASE-LX
- 1000BASE-SX
- 1000BASE-T
- 10GBASE-CR
- 10GBASE-LR
- 10GBASE-SR
- 10GBASE-SRL
- 100GBASE-XSR4
- 100GBASE-DR
- 10GBASE-ERBD
- 10GBASE-ERBD-D
- 10GBASE-ERBD-U
- 10GBASE-ERLBD
- 10GBASE-ERLBD-D
- 10GBASE-ERLBD-U
- ADPT-O-Q-100G
- 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
- OSFP-400G-2FR4
- OSFP-400G-LR8
- OSFP-400G-SR8
- QDD-400G-LR8
- 100GBASE-FR
- CAB-Q-2Q-100G
- OSFP-400G-DR4
- AOC-D-D-400G
- 10GBASE-ER
- 40GBASE-CR4

Reference to MLAG ISSU Upgrade/Downgrade Table

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

Resolved Caveats in 4.22.8.3M

Accelerated System Update

- Traffic flows on trunk ports in VNI mapped VLANs may experience greater than 1 second outage during SSU. (706641)

7300X3, 7320X, 7368, CCS-720XP, DCS-7010, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260X, DCS-7260X3 and DCS-7300X 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 4.22.8.3M

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)
- Performing SSU with an MTU value less than 1280 on an IPv6 interface may result in a fatal error with extended traffic outage. (400408)
- 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
- During Leaf-SSU, interfaces configured with auto-negotiation may flap. This will lead to a brief traffic disruption on the affected interfaces. (619772)
Series Affected: DCS-7050CX3, DCS-7050SX3 and DCS-7260CX3 Series

BGP EVPN L2/L3 VXLAN/MPLS

- DF election may not work correctly after changing VLAN - VNI binding. (484811)
- 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 the the MLAG primary may trigger the lost of some dynamic VLAN allocated for certain VRF to VNI mappings on the secondary.

The workaround is to restart the VXLAN agent. (548097)

- In an EVPN network, if a MAC address is advertised by a remote device with a non-zero move count and is subsequently removed as a result of a mass withdraw event (EVPN type 1 auto discovery route withdraw), if that MAC address moves to the local device before the remote device withdraws the corresponding EVPN type 2 MAC/IP route, the local device will fail to advertise its own EVPN MAC/IP route. As a workaround, once the relevant MAC/IP routes are withdrawn from the remote device, clearing and re-learning the affected MAC addresses will cause them to be advertised correctly. (549976)
- In a MLAG EVPN IRB deployment, reloading the MLAG secondary or putting both MLAG peers into and out of Maintenance Mode may result in an IP-VRF not have an associated VLAN.

The workaround is to restart the VXLAN agent. (569663)

- When performing an operation that can cause configuration churn on interfaces with Ethernet Segment configuration (such as a config replace) at high scale, the BGP agent may restart unexpectedly. The issue doesn't repeat after the BGP agent has restarted (unless there is another large configuration churn event at high scale). (576033)
- 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)

- When moving a VLAN from one VLAN-aware-bundle to another, MACs associated with the VLAN may fail to program if the VLAN was added to the new bundle before it was removed from the old bundle.

The workaround is to delete and re-add the VLAN, so the VLAN is added while it is not present on any other bundle. (594211)

- In an EVPN deployment, the VLAN flood list may be deleted when moving a VLAN from one VLAN aware bundle to another MAC-VRF.

The workaround to correct the flood list is to remove the VLAN from the MAC-VRF and then add it back. (603591)

- In an EVPN VXLAN All-Active Multihoming deployment, the ARP entry of a multi-homed host learned through MAC IP route advertisement can be lost. This issue does not apply to IPv6 neighbor entry for a multi-homed host. The issue does not apply to hosts connected via L2 only (no SVI for the VLAN).

A workaround is to clear the missing ARP entry on the multi-homed peer that learn it over the data plane and advertised the routes. (608813)

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

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

- Installing an EVPN MAC-IP route with MAC address value of zero may cause BGP process to restart. (696815)
- In an EVPN IRB setup with VXLAN, if a host moves from being local to being remote relative to a switch, traffic from behind the switch towards the now-remote host may not get resolved.

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

Containerized EOS

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

vEOS Router / CloudEOS

- With DPS enabled, there is a small chance that SFE agent may restart unexpectedly after a DPS path flap event. (527588)
- When an IPsec connection is established with a primary IP address of an interface and a secondary IP address is configured on an interface and that secondary IP address is used to establish an IPsec connection, the IPsec connection with the secondary IP address doesn't come up.
The workaround is to restart the IPsec agent. (572808)
- 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)

CVP

Network Provisioning

- OpenConfig configurations involving route targets may be incorrectly ignored. (528374)
- ISIS auth config (at global and interface levels) is lost when any ISIS or interface related config is pushed via GNMI or NetCONF. (543354)

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 an 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 the Macro-Segmentation Service (MSS) is enabled on an MLAG VTEP and the VTEP receives the ARP entries from the VXLAN Control Service (VCS), a reload of the secondary MLAG peer can result in ARP entries getting deleted on the peer. The workaround is to disable MLAG and re-enable it on the secondary peer. (411505)
- In a Macro-Segmentation Service (MSS) deployment, a firewall REST API providing a malformed response may cause enforced policy to be withdrawn and re-applied. (512505)
- 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)

- After a CVX failover the NetworkTopologyAggregator agent may restart.

A workaround is to trigger another CVX failover. (545008)

- When using Macro Segmentation Service (MSS) with Palo Alto Networks firewall, presence of non-ASCII characters in any field (e.g. policy description) on the firewall might not allow the CVX to retrieve any configuration from it.

Workaround is to replace non-ASCII characters with their ASCII equivalent if possible. (558655)

- 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)

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

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)
- When one or more neighbors is moved to a new update-group as a result of neighbor outbound policy change, we might transiently send withdrawals/incorrect delta advertisements.

This is eventually corrected when policy evaluation of new update-group eventually completes. (226201)

- 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 VM Tracer agent may restart when connected to a VMware vCenter instance with a VMKernel not assigned with an IPv4 address. (258986)
- The MacroSegmentation Service (MSS) feature doesn't work in conjunction with the VARP VTEP IP configuration required for VXLAN routing where a subset of VTEPs are L2 VTEPs. (263532)
- Support for VRRPs in OpenConfig not fully implemented (268736)
- An ACL specified in a YANG update to configure an SSM range is not correctly resolved and the configuration is not updated. An error is returned to the user. (278056)
- 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
- OpenConfig agent crashes when there is an unknown keyword in the ODL RPC request (369699)
- 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)

- A gNMI get on /interfaces/interface/routed-vlan/ipv4/config/enabled returns many interfaces which are not enabled. (378280)
- 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)
- When performing a gNMI get on /interfaces/interface[name=Ethernet7/1]/ethernet, the gNMI server returns an rpc errors as - 'skipped' (396967)
 - If a YANG client is used to update VLAN configuration, OpenConfig may add an empty name to VLAN configuration. (399481)
 - The EOS DHCP IPv4 Server can enter a deadlock state when show commands and configuration commands are run simultaneously.

To clear the deadlock, execute `sudo kill -9 <pid>` where <pid> is the process ID of the Kea DHCP Server process. (399832)

- gNMI Get on nodes that failed their when statements return rpc skipped errors instead of returning nothing. (400274)
- gNMI get displays nodes whose ancestor nodes failed when statement(s); These nodes should not be displayed (400283)
- In the CLI BFD-Router Configuration mode, if the CLI command "local-address" is used to modify a BFD local-address the previous address and the new address may be present at the YANG /bfd/config/local-address path. In the CLI Interface Configuration mode, if the command "bfd neighbor" is used to modify a neighbor's IP address the previous address and the new address may be present at the YANG /bfd/interfaces/interface/config/remote-address path. (400528)
- If a YANG client deletes values at /bfd/config/local-address, the values may remain in yang tree even though the YANG delete is successfully applied on the switch (400568)
- Routing protocols such as OSPF, IS-IS are disabled when VLAN or interface config is changed via gNMI or NETCONF. (401590)
- When the forwarding agent hitfully starts up, PTP packets sent downstream may have invalid timestamp and the switch may incorrectly advertise itself

as the GM. Valid PTP packets and correct GM advertisement will converge post forwarding agent startup. (402786)

- In NETCONF output, YANG leafrefs to identityrefs are encoded incorrectly, with <Module> and <Name> elements. (402952)
- gNMI 'get' will not show YANG data defined in arista-bfd-augments.yang even though the configuration exists on the switch. (406427)
- 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 may restart unexpectedly. (407002)
- If the number of PTP vlan's enabled on a switch using 'ptp vlan' command exceeds 65536, PTP agent will become unresponsive and eventually restart. Reducing the VLANs below the limit will avoid getting in this state. (408199)
- When in transparent two-step mode PTP packet sequenceId collision may cause transient jump in offsetFromMaster. Workaround is to use one-step transparent clock mode. (408202)
- When more than 255 subnets are configured for DHCP Server, the agent will continuously restart. To correct the DHCP Server error state, the configured number of subnets must be changed to 255 or less and the DHCP Server ipv4SubnetLastUsedId state in Sysdb must be manually cleared. (411213)
- DirectFlow 'action output nexthop <ip addr>' ignores TTL of incoming packet and forwards the packet even if the TTL has expired. (417994)
- In OpenConfig, some VRF configuration can be missing from the YANG tree if hundreds of VRFs are configured. (426960)
- OpenConfig may restart if IPv6 static routes are configured on the switch. (450019)
- OpenConfig is affected by CVE-2020-7919 in 32-bit EOS. (450879)
- When an OpenConfig NETCONF client is used to add new classes to policy map, the new classes are added before any previous classes. (451025)
- Licenses bound to a device, such as MACsec licenses, would not be found after boot until the license store is refreshed. (451099)
- gNMI, NETCONF, and RESTCONF access to the device may be disrupted if client certificate authentication is configured and a client connects with a malicious certificate. (464752)
- 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
- Octa sends more data than requested for a gNMI Subscription with origin "eos_native" and mode SAMPLE. (482142)
- Octa may restart unexpectedly when handling a gNMI Subscription with origin "eos_native" and mode SAMPLE. (482143)
- When EOS reaches the system-wide open-files limit, agents can restart unexpectedly with out-of-memory log messages, and may cause the system to restart unexpectedly. (488465)
- 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)
- The CLI session can close unexpectedly when running a command that internally encounters an IO error. This can happen when running "no shutdown" under daemon mode while the exe is executable but unreadable (e.g., with a permission of 0711). If the command is loaded via a configuration file, partial configuration may be committed. (501541)
- On a device with a high rate of updates to MPLS LFIB counter the frequent counter update rate may cause OpenConfig to restart unexpectedly. (506329)
- Value for mlag leaf might not be mapped in certain circumstances. (518968)
- A customer-provided agent using EosSdk may restart when configuring or modifying an ACL. (524324)

- Applying a Policy-map created through OpenConfig may result in unexpected restart of the Strata-FixedSystem agent (525595)
- A mis-configuration of the event-handler feature when it uses the on-logging trigger with no regex command will result in a continuous triggering of the handler's action. (528458)
- Octa can incorrectly report collection entries as deleted under /Smash eos native paths.

Restarting Octa will make the incorrectly reported deleted entries appear again, however they or other entries can subsequently be incorrectly reported as deleted again. (537209)

- gNMI clients subscribed to paths with origin eos_native may stop receiving notifications.
To work around the issue restart Octa or start a new subscription. (539454)
- When a member interface of a PTP-enabled port channel is removed from the port channel, a subsequent speed change on the former member interface may cause the Strata agent to unexpectedly restart. A workaround would be to disable PTP on the port channel before removing the member interface. (540175)

Series Affected: DCS-7050CX3 and DCS-7050SX3 Series

- The "arp reply relay" feature does not support a VXLAN domain with L2 VTEPs. (540369)
- If multiple clients subscribe to overlapping /Smash eos native paths provided by Octa, Octa can unexpectedly restart. (540472)
- If an OpenConfig RESTCONF server is configured with an access control list, then the ACL may not take effect. (548311)
- A NETCONF edit-config request for classes in a policy map may cause a change in the order of classes in the same policy map or another one. (550060)
- OpenConfig mixed schema configuration request can fail if the CLI modes send are capable of accepting a "commit" CLI. A workaround is to add the "commit" as part of the request. (554372)
- 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)
- The python2 and python3 versions in EOS are vulnerable to CVE-2021-3177. (560484)
- BmpAgent, Sflow or BgpCliHelper may abort and restart with the agent log message "DirectRegionHugeMap::remove() i:42 is not region:XXX" (563870)

- In PTP transparent clock mode, a PTP packet updated by the switch may have its VLAN tag stripped when it is forwarded. This can happen if the VLAN-tagged PTP packet ingresses a PTP-enabled VLAN trunk port, egresses a PTP-enabled VLAN trunk port, and there are other PTP-enabled ports configured with a switchport access VLAN. (566037)
- If PTP is configured, performing a configure replace that disables PTP and performs speed changes on interfaces may not disable PTP on those interfaces. A workaround would be to toggle PTP globally or toggle PTP on PTP disabled ports. (569723)
- Adding or modifying class-map in a policy-map via OpenConfig may not result in applying the new policy-map in hardware. (573971)
- Adding or modifying Route-map via OpenConfig may not result in applying the new Route-map in hardware. (575325)
- In a VXLAN routing setup using MLAG configuration, there may be excessive UDP (port 51023) traffic observed over the MLAG peer link. This occurs when there is lots of VXLAN ARPs received from remote VTEPs with sender MAC set to virtual MAC. The excessive traffic may result in self-IP queue drops on the receiving MLAG peer, and high CPU utilization by VxlanSwFwd agent on the sending MLAG peer. (583243)
- In a VM Tracer deployment, connecting to a vCenter instance managing an unresponsive ESX host may cause the VM Tracer agent to restart and prevent the processing of CDP and LLDP updates from other ESX hosts. (591593)
- Under a high system load, the PTP agent may intermittently fail to generate packets to downstream PTP devices. There is no known workaround for this issue. (592524)
- In PTP two-step transparent clock mode, delay response packets traversing through the switch may have their correction field incorrectly updated.

A workaround would be to use PTP one-step transparent clock mode. (592527)

- PORT_ID allocation will fail on systems with more than 4095 interfaces. This includes peer interfaces in an MLAG configuration (593759)
- 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)

- Resolve issues with incorrect use of AAA APIs. (601875)

- eAPI may skip re-evaluating user credentials when certificate based authentication is used. (606686)
- 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)
- Configuring an ACL on a VLAN interface may lead to a leafref integrity violation in the OpenConfig YANG tree precluding any successful SET operation on the YANG data.
A workaround is to remove the ACL on the VLAN interface. (614069)
- 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: DCS-7020, DCS-7280CR2, DCS-7280CR3, DCS-7280E, DCS-7280PR3, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R, DCS-7500R2 and DCS-7500R3 Series

- 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)

- A PTP Management message with a certain TLV size close to 65535 will cause the PTP agent to restart unexpectedly. Agent will recover after the restart.
Workaround is to block such PTP Management messages from entering the PTP network. (638107)

- 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)
- 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-7050SX3, DCS-7050TX, DCS-7050TX2, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260CX, DCS-7260CX3, DCS-7260QX, DCS-7260X, DCS-7260X3 and DCS-7300X 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)

- Certain ill-crafted certificates might cause infinite loop during verification. (674519)
- 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)

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 Series
- Ports configured for 100G without FEC with the CRT50216 PHY may remain down after an AIS or PRBS signal is received.
A work around is to enable and then disable FEC with 'error-correction encoding reed-solomon' followed by 'no error-correction encoding' (498631)

- MACsec frames shorter than 88-bytes long without SCI and 96-byte MACsec frames with SCI are dropped upon reception. (514526)

SKUs Affected: 7500R-8CFPX-LC and 7500RM-36CQ-LC

- Interfaces drop unencrypted pause or PFC frames while MACsec is enabled. (538882)

SKUs Affected: 7500R2AM-36CQ-LC, 7500R2M-36CQ-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

- 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-negotiation, EthernetX/1 linking down can cause EthernetX/3 to link down. The above is true for EthernetX/5 and EthernetX/7 as well.

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

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

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

SKUs Affected: 7500R3-24P-LC, DCS-7280CR3K-32P4-R, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F and DCS-7280PR3K-24-F

- Interfaces with 1000BASE-T SFP transceivers may report a false link up when configured with default, "no speed", "speed auto", or "speed auto 1000full". The workaround is to configure affected interfaces with "speed forced 1000full" instead. (614753)

Series Affected: 7320X, DCS-7060CX, DCS-7060CX2, DCS-7060SX2 and DCS-7260CX Series

- 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, 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-7280CR3K-32D4, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4-F and DCS-7280CR3K-32P4-R

- 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

- Inserting a 1Gbps SFP transceiver or 10Gbps BASE-T SFP transceiver configured at 1G into a QSFP port using a QSFP-SFP Adapter (QSA) causes 1 logical port to be allocated. If the interface is configured with "speed forced 1000full" or "speed auto 1000full", transceiver removal causes 3 additional logical ports to be allocated. These logical ports may be reallocated to other active interfaces and cause them to become inactive.

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

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

- 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-7050SX3 and DCS-7260X3 Series

Layer 2

- If Ebra agent restarts on a VTEP that is running EVPN control plane with route type 5, then it can result in VxLAN traffic loss. (276063)
- MACsec delay protection feature is not working. (282186)
- When running rapid-PVST at high scales of interfaces and vlans, the device may occasionally not transmit a BPDU, which can lead to a BPDU timeout on the neighbor. (442663)
- Performing an SSO switchover on an MLAG primary peer running Multiple Spanning Tree Protocol may result in STP topology reconverging. (470962)
- Restarting STP agent with a high number of interfaces and vlans in PVST mode may result in STP topology reconverging. (472675)
- In VXLAN deployment, remapping the VNI from one VLAN to another VLAN may cause the destination VTEP to drop ARP requests received over VXLAN tunnel in that VNI. (482403)
- Re-configuring a member interface from a port-channel configured with "port-channel lacp fallback individual" and currently in fallback individual mode to another port-channel that does not have the fallback configuration can cause the member interface to remain in Bridged forwarding state instead of Data Link. This can potentially cause traffic loop.

To workaround, un-configure the affected interface from the port-channel and reconfigure it. To avoid this issue while moving an interface from one port-channel to another, un-configure the interface from current port-channel before configuring it into the new port-channel. (525342)

- On a device that's an MLAG or had MLAG configured if egress sFlow is enabled, interfaces can go to blocking with syslog SPANTREE-4-INTERFACE_SELF_LOOPED.

To workaround, disable egress sFlow. (546134)

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

- LACP port-channel member interface(s) may unexpectedly leave the port-channel

due to "partner not in sync" syslog reason without member interface link status change. The member interface will rejoin the port-channel immediately after. However, the event may cause port-channel interface to flap if there is only one member interface configured in it. (546649)

- On a device that's an MLAG or had MLAG configured in the past and configured with rapid-pvst STP mode, if more VLANs are configured than platform limit, then the StrataVlanTopo agent may continuously restart. Device won't log SPANTREE_TOPOLOGY_LIMIT_EXCEEDED to indicate platform topology limit exceeded.
To workaround, reduce number of VLANs to platform limit specified in Arista product documentation. (548254)
- 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)
- 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)

Layer 3

- If a loopback interface is only configured with a link-local IPv6 address, OSPFv3 will not advertise the address. (131369)
- 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 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 a large number of BFD sessions are active, configuration changes may cause sessions to become stuck in incorrect states. Restart the BFD agent to work around the problem. (377125)
- Setting `ipv6 icmp rate-limit-unreachable 0` and shutting down all the next hop interfaces for an IPv6 default route resilient ECMP causes the L3 Agent to restart. (395322)
- When an ARP entry is added while there is a matching, existing virtual IP with a non-/32 subnet, traffic loss may occur because traffic may be flooded instead of forwarded to the CPU.

Workaround is to either use multiagent mode or only configure virtual IPs with a /32 subnet. (402802)

- Large numbers of RSVP node protection enabled RSVP tunnels sharing the same outgoing interface can cause the layer 3 forwarding agent to transiently crash after a link down occurs on that interface if hardware fast failover is enabled as well. (409973)
- 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)
- When BFD is configured on a link and the link is flapping, routes resolved over this link may stay unprogrammed if BFD status is down. Removing BFD configuration on the affected link should resolve this issue. (415032)

- Receiving a trigger Resv message for an RSVP-TE bypass tunnel that is in active use may cause a brief drop in traffic forwarding over the bypass (415287)
- 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)
- Cspf agent may restart when TI-LFA is configured and the backup path is via a neighbor to which the point of local repair has ecmp paths and atleast one path is not directly connected to the neighbor (427437)
- When processing flow-spec rules with match components that carry <match operator, value> (for example: match component type 3, type 4, etc), if the match operator specifies a length that is greater than what is needed to encode the actual value, then "show flow-spec ..." may incorrectly report the installed state of the rule. In addition, the SandAcl agent may crash continually if multiple flow-spec rules were received that differ only in the match operator's specified length.
For example: The two rules noted below are expected to match an IP protocol value of 10. However, they encode the number of bytes needed to represent the value differently.
Rule1: Match Component Type 3, <operator: len = 0, value: 10>
Rule2: Match Component Type 3, <operator: len = 1, value: 10> (433033)
- 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 static ARP entry is added via CLI with a MAC of 0.0.0, Arp will continuously crash. The workaround is to either remove the static ARP entry or add a non zero MAC address to the static ARP entry. (444677)
- With FIB compression being enabled, the StrataL3 agent and SandL3Unicast agent can restart unexpectedly. (445777)

- 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)
- IS-IS adjacency will get dropped when the IS-IS SHA authentication feature is used in interop scenarios with non Arista devices. (463543)
- 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)
- Running 'show ipv6 hardware ale routes <prefix>' or 'show ipv6 hardware fib routes <prefix>' with a prefix that has trailing bits that exceed the prefix length may cause the platform agent to restart. (473799)
- 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)
- After several subsequent interface shut/no shut, some routes might not forward traffic. The workaround is to toggle the interface once more. (476768)
- When the MAC address changes on a BFD peer, the session may flap. This can happen when the interface fails over from active to standby. (477559)
- 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, particularly when all or most of the entries map to the same route, the Arp agent may encounter a heartbeat timeout, leading to the Arp agent restarting. 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)
- All ECMP static vxlan evpn routes for a prefix get deleted when one of the contributing static vxlan evpn routes is deleted. (493505)
- 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)

- When changing the OSPF Router Id, OSPF will not flush its existing self-originated LSAs from the routing domain before the new Router Id takes effect (499551)
- 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)

- For short-lived TCP connections, NAT may not tear down the address mapping when it should. The workaround is to lower the TCP timeout using the `ip nat translation tcp-timeout` CLI. (526152)
- IP routes and FECs may fail to install in some instances when we recover from a resource overflow scenario. A workaround is to restart the L3 platform agent. (526893)
- On scale setup, with certain network events that cause massive route updates, ARP agent's internal state may go out of sync, leading to the 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 the ARP agent. (532708)

- Stateful switchover can cause a small traffic loss because of host routes without ARP resolution. (533234)
- When there are more than 254 RIP enabled interfaces then it may result in restart of the protocol agent if a route needs to be advertised to more than 254 neighbours. (535893)
- 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 disabling MPLS on an interface in the non-default VRF, it may cause software forwarding or control plane generated IP packets for routes which point to tunnels, nexthop-groups or when routing from one IP address-family to another (RFC 5549) to route the packet incorrectly. This will only affect packets in the default VRF. These packets may end up being routed in the non-default VRF.

Configure an SVI with no autostate configured in the default VRF to work around this issue. (542757)

- 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)
- 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)

- When using OSPFv3 and encryption or OSPFv2 with SHA encryption, the Rib agent may stop responding and restart unexpectedly. (560294)
- 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)
- Rib agent may restart when BGP PIC ("BGP additional-paths install") is configured and a BGP path's next hop resolves over a BGP route with a backup path.
Workaround is to unconfigure the additional-paths configuration. (566312)
- When 64 areas are configured in OSPFv3, router configuration may not be applied after a system reload. To workaround the issue, configure less than 64 OSPFv3 areas. (570346)
- 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)
- Rib agent or the multiagents (Isis/Ospf) can restart with a large prefix-list configuration (578037)
- A DHCPv6 prefix received with valid lifetime of zero may persist forever in the routing database in the presence of 'ipv6 dhcp relay install routes'. A workaround might be to toggle "ipv6 dhcp relay install routes" command, which deletes all DHCPv6 routes, including the stale ones with negative lifetime. (578276)
- 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)

- With BFD VXLAN configured, BFD agent restart after a first supervisor failover may result in continuous BFD agent restart.

There is no workaround once the restart loop starts. Performing another supervisor failover before restarting the BFD agent will avoid the restart loop.
(586673)

- 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)
- The OSPF agent may restart when sending a graceful restart LSA on an interface with SHA authentication configured. (593659)
- 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)
- OSPFv3 fails to install routes from intra-area-prefix(NAP) LSAs if the referenced LS ID in the NAP LSA does not match the Link State ID of the RTR LSA (605376)
- In setups with large number of ARP or ND entries in the same subnet, upon a more specific route in that subnet getting removed, the Arp agent may reevaluate all of the ARP or ND entries in that subnet, resulting in high CPU usage in the platform agent.

Workaround is to split up large connected routes by adding more specific static interface routes to make the subnet that the ARP/ND entries fall under smaller. (606649)

- 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)
- ARP requests for NAT pool addresses might fail due to incorrect initialisation.
As a workaround restart NAT agent to recover from the issue. (610947)
- 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)
- 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)
- 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 reinitiates the instance like: LSP buffer size update, IS level update, address-family update etc (710168)

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)
- There may be some traffic loss on ASU or SSO restart at high scale (i.e with a large number of next hops). (362995)
- Performing a stateful switchover in a switch with a large number of VRFs configured, with BGP config in all VRFs and IGP configured in a large number of VRFs, may cause the system to restart unexpectedly, leading to traffic loss. (376773)
- The "Not best" output in "show ip bgp <prefix> detail" and "show ipv6 bgp <prefix> detail" does not consider the IGP cost step in BGP best path selection. These commands will claim that affected paths are "Not best" due to some comparison performed after the IGP cost step. This is specific to the show command output; best path selection itself handles IGP cost correctly. (401354)
- With BGP configured in VRF, traffic loss may be seen in some VRFs on "reload fast-boot" due to Forwarding state preserved(F) bit not being set in the open message (414457)
Series Affected: DCS-7280SR and DCS-7280TR Series
- Routes may not be updated in the kernel when its next hop gets resolved via a drop route. (444682)
- ISIS/OSPF/OSPFv3 agents can restart unexpectedly if interface static routes over tunnel interfaces are redistributed into them. The workaround is to

configure gateway static routes using nexthop IP instead of interface static routes. (450213)

- IS-IS SR LFIB may be deleted and recreated if the ECMP resolution for IS-IS SR adjacency segment changes (457220)
- 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)
- Bgp agent may restart unexpectedly when 'bgp always-compare-med' option is configured and large number of competing paths for the same prefix from several neighbors with different MED values are present (482275)
- During SSU or SSO, 'update wait-for-convergence' behavior is unnecessarily forced for AFI-SAFIs which do not support SSU or SSO. (483591)
- At large route scale, the restart of either the Bgp agent or the IpRib agent may cause repeated restart of the IpRib agent (508763)
- If a BGP peer (or peer group) description is longer than 255 characters, polling the aristaBgp4V2PeerDescription MIB object may fail. In addition, if BGP traps/notifications are enabled, the Snmp agent will restart whenever it attempts to send the aristaBgp4V2BackwardTransitionNotification trap.

Workaround: Limit the BGP peer(group) description to 255 characters. (528521)

- A switch configured with MPLS EVPN may proxy reply and not flood a gratuitous ARP announcement message, which may prevent cluster failover migration of a virtual IP to a new server. (541463)
- The configuration 'bgp missing-policy direction out action deny' may be ignored briefly when a route map applied on outbound to a BGP peer is deleted.

A workaround for this is to remove or replace the route map from the peer without deleting the route map. Once the update has taken effect the route map can be deleted. (554264)

- 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)
- For 6PE session, route refresh requests are sent out for IPv6 Unicast rather than for IPv6 Labeled Unicast. (569595)
- An IPv6 nexthop may optionally contain an additional link-local IPv6 address. A BGP neighbor sending IPv6 Labeled Unicast NLRI with a nexthop that includes both global and link-local nexthops will be incorrectly marked as malformed, causing IPv6 LU AFI/SAFI to be disabled for that neighbor. (575422)
- In multi-agent mode, when displaying the results of "show ip bgp detail" the output may be corrupted by showing content previously output by the currently running command in place of new content (607255)
- The Bgp agent may restart unexpectedly if two threads simultaneously apply input policy and the result is a rejected path. (652370)
- In some EVPN deployments with large number of 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)

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)
- On Mlag setup with dual primary detection enabled, traffic loss might be seen when recovering from dual primary condition. (297307)
- On Mlag setup with dual primary detection enabled, traffic loss might be seen when recovering from dual primary due L3 convergence. With recovery delay support, MLAG interfaces can be configured with a longer recovery time to

allow

L3 convergence before MLAG interfaces come up. (411211)

- MLAG ISSU breaks on a switch running a version of EOS < 4.20.5 if it receives ARP entries from the VXLAN Control Service (VCS) while the peer has been upgraded to EOS version >= 4.20.5 (497776)
- When the M-LAGs flap due to missing LACP packets from their partner interfaces and end up in active-partial state while a burst of MAC learning events are being exchanged between the two MLAG peers, then it is possible for the L2 agent to skip programming some of the MAC addresses synced between the two peers.

The workaround is to flap the M-LAG members on the peer whose MAC table is out of sync. (543878)

- 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

- The Rsvp agent may restart unexpectedly when a large number of LSPs are using the same node protection bypass tunnel and that tunnel gets deleted. (415301)
- When EOS routers are deployed as both LSR and tail end, an RSVP primary LSP going through the EOS LSR may flap if its associated bypass LSP originating at that LSR flaps even when the bypass LSP is not in use. (417161)
- The 'hardware next-hop fast-failover' configuration may have not function correctly on a router with port-channel interfaces configured. (439078)
Series Affected: DCS-7500E, DCS-7500R and DCS-7500R2 Series
- 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)
- MPLS agent might restart due to route churn in FIB containing large number of routes. (526406)
- 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)

- The "configure replace" command causes LDP sessions to be cleared when they were previously cleared using the "clear mpls ldp neighbor". The cleared sessions will immediately recover. (549437)
- 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)

Multicast

- In a VLAN that only has an IGMP proxy querier(0.0.0.0) present, IGMP snooping does not start flooding traffic in the VLAN when the querier disappears. (258931)
- 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)
- In an environment with a very large number of MSDP peers, during cleanup the MSDP Agent may unexpectedly restart. (354364)
- If large number of IGMP reports are received in a very short interval, some of them might get dropped resulting in the multicast route not being created. The routes will be re-learned on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIgmp", consider increasing the bandwidth for this Copp class. (356675)
- 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)
- Running PIM over SCTP transport and PIM over datagram at the same time could result in periodic PIM join-prune messages, that are sent to non-PORT neighbors to have an incorrect source address. As a result will be dropped by the neighboring PIM device. This will result in multicast traffic loss. (398383)

- When using Igmp host proxy, longer than expected traffic gaps might be seen when proxying source specific joins. (408784)
- In an environment with a very large number of MSDP peers, during cleanup the MSDP Agent may unexpectedly restart. (418190)
- Deletion of PrefixIntfSm may cause unexpected restart of the corresponding Mroute agent (471017)
- 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)
- 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)
- If the SandMcast agent restarts while IPv6 multicast routes are programmed, the routes may become incorrectly programmed. The programming of the IPv6 multicast routes can be corrected by toggling IPv6 multicast routing off and back on. (519865)
Series Affected: DCS-7020, DCS-7280CR2, DCS-7280CR3, DCS-7280PR3, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R, DCS-7500R2 and DCS-7500R3 Series
- Forwarding agent might unexpectedly restart or use incorrect RPF setting during unicast route churn. (539478)
- On devices running PIM, on system boot up, might suffer loss or slower re-convergence of multicast traffic flows. (570930)
- On devices connected to multicast receivers that are part of a MLAG pair, restart of IgmpSnooping/Mlag agent may result to more IgmpSnooping restarts. This state has no impact on forwarding behavior of the box.

Workarounds: Reload the device (574602)

- 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)

- When an interface is a (*,G) assert loser, traffic duplication might be present until the assert winner sends the next periodic assert of around 3 minutes. (587719)
- The PimBidirDf agent could continuously send DF winners when multiple RPs are configured and one RP is deleted. This may cause high CPU in the PimBidirDf agent. Workaround is to restart the PimBidirDf agent. (603791)
- 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)
- Icmp 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 icmp host proxy access-list <acl>` and use `ip icmp 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)

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

NAT

- Changing the configuration of NAT synchronization connected MLAG peers might cause the NAT agent to unexpectedly restart. (680566)
Series Affected: CCS-720XP and DCS-7050SX3 Series
- On MLAG set-up using NAT synchronization between the two peers, TCP connection might reset while the connection is being established.

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

Series Affected: DCS-7170 Series

- %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)

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)
- Adding an app-profile with no rules to a DPS policy causes the Sfe agent to restart. (401675)
- Hyperthreading incorrectly disabled on MSFT Azure. This causes loss of performance due to reduced number of cores. (402981)
- set parent interface's mac-address to system-mac for subintf to work (407368)
- 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)

- If IPv6 or multicast packets are sent to vEOS on two or more interfaces vEOS forwarding engine can crash. Disable IPv6 and mcast configuration on the interfaces of the neighboring routers connected to vEOS. (497553)

- 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)
- 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)

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)
- The system may not power cycle when configured to power cycle upon CRC detection on the system FPGA. (431702)
- Further improvements to reduce the probability of watchdog reboots due to SEUs on the switch chip (501119)

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)
- 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 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* forwarding 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

CCS-720XP Series

- Under heavy CPU load, the Strata agent may unexpectedly restart because of a timeout. (348087)
- 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)
- Inserting and removing a QSFP+ transceiver may result in the forwarding plane agent restart and a brief traffic outage to occur. A workaround is to configure a forced speed on the corresponding interface via CLI before inserting/removing a transceiver.
Changing interface speed interface configuration via CLI may require user confirmation to continue. Upon user confirmation, the forwarding plane agent restarts and a brief traffic outage occurs. (388282)

Series Affected: CCS-720XP Series

- 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 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)

- 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

- If Dynamic authorisation request is sent with incorrect format of Calling-Station-Id then this may lead to ungraceful handling of the request. (548065)

- Link partners connected to the SFP28 ports with optics might see PCS/FEC errors, link flaps or failure to link up. (556805)

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

- For a given tx-queue, traffic on unicast queue may starve the multicast queue. (616128)

Series Affected: CCS-720XP Series

- 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

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)

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

- When many IPv6 routes are configured and the BFD protocol in use, BFD sessions may flap when a linecard is removed or shut down. (279129)
- ManagementActive Agent repeatedly restarts and Management0 interface is not reachable in Modular systems. Doing switchover will address the issue. (369565)
- Hotswapping a card may cause a kernel panic and the switch will reboot unexpectedly. Not running traffic on the card before powering down will make this hit less frequently, but it may still hit. (387173)
- On modular systems with redundant supervisor cards, if the standby supervisor has connected management interfaces, ZeroTouch provisioning may fail to download a configuration from any connected DHCP servers. The workaround is to "shutdown" the management interfaces on the standby supervisor. (507185)
- During SSO, a fabric card's forwarding agent may experience an additional unexpected restart. In such cases, the forwarding agent will emit a "STRATA-3-RXDMASTUCKERROR" syslog message. A brief traffic outage may be observed. (572329)
Series Affected: DCS-7300X Series
- Replacing certain linecards on modular systems with one that has a lesser switch count while SSO is enabled will result in agent restarts on the standby supervisor. This affects replacing a 7300X-32Q-LC with a 7300X-64S-LC and replacing a CCS-750X-48ZXP with either a CCS-750X-48TP or CCS-750X-48ZP.

The workaround is to configure the system for RPR mode, replace out the cards, and then reconfigure for SSO mode. (586499)

SKUs Affected: 7300X-32Q-LC and 7300X-64S-LC

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

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

CCS-720XP Series

- PoE ports that are able to classify but do not power on for an extended period of time may cause Sysdb memory usage to grow, eventually resulting in out-of-memory events, particularly with TerminAttr

A workaround is to configure "poe disabled" on the affected ports (528060)
SKUs Affected: CCS-720XP-24Y6, CCS-720XP-24ZY4, CCS-720XP-48Y6 and CCS-720XP-48ZC2

DCS-7020, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7500E, DCS-7500R, DCS-7500R2 and DCS-7500R3 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)
- When a large number of PBR service policies are configured and the SandAcl agent restarts a number of times, the SandAcl agent may continuously restart.

System restart will remedy the situation. (166081)

- 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)
- Unknown unicast and multicast packets are unnecessarily copied to CPU after VXLAN decapsulation when there is an SVI configured on the VXLAN VLAN. (258990)
- When a 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)
- Egress acl deny logging does not work when the packet is destined to a RFC5549 next-hop. (264045)
- 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 flowspec has both IPv4 and IPv6 rules and the tcam space is sufficient to program only the v4 rule set and not v6 (or vice versa), the v4 rule set may be partially installed. (342060)
- 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)
- With MPLS pop routes and IP egress ACL configuration, restarting the L3 forwarding agent may lead to another unexpected restart. (349534)
- When IPv4 or IPv6 flowspec rules exceed the available TCAM capacity, the rules for the other protocol may not be installed even if there is space in the TCAM for it. (369374)
- While in Tap Aggregation mode with VXLAN tags stripping feature enabled, forwarding of VXLAN tagged packets will be disrupted if the Sand agent or the SandFap agent restarts.

Reconfiguring the VXLAN tag stripping feature on the relevant tap ports will restore functionality. (372603)

- Traffic coming into subinterfaces are subject to the flowspec rules applied on the parent interface's VRF. (380533)
- 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)
- While in TAP Aggregation mode, swapping the transceiver on a configured tool port may cause identity tagging on this port to stop functioning.

Toggling the switchport mode or "tool identity dot1q" configuration on the affected tool port will restore functionality. (391367)

- BGP Flowspec and Cpu Traffic Policy may be unable to recover from programming failure when there is initial TCAM bank exhaustion, but at a later time, TCAM banks do become available and the feature is then applied to another interface on another forwarding ASIC. The workaround is to remove and reapply the feature on all the interfaces. (393500)
- The Type Of Service (TOS) field in MPLS-terminated IP packets may become corrupted when using a user-defined TCAM profile.

A workaround is to add the affected packet type to 'feature ip qos' in the user-defined TCAM profile. (394854)

- 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)

- On PTP enabled switches, CRC errors may occur in internal time synchronization that can cause PTP time to be unstable. PTP time will recover stability on its own after a few seconds. (399179)
- Insertion of a fabric module may cause brief traffic loss on the device (402508)
SKUs Affected: 7504R3-FM and 7508R3-FM
- Packets that will get vxlan routed do not get redirected by PBR policy. (402759)
- Adding new VLAN members to a CPU traffic policy using the police action may result in policers not being allocated properly for those members. The workaround is to deconfigure and reconfigure the policy. (402804)
Series Affected: DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- Applying ipv6 ACL and then bringing up any lag subintf will allow traffic not matching any rule to flow through.
Reapplication of the ACL resolves the issue. (404160)
- When a traffic-policy with non-default actions is modified in a config-session the policy action stops working. (406486)
- With scaled Acl configurations, SandAcl agent restart may result in additional agent restarts. (408260)
- When an ACL with a 65K or more rules is applied to an interface, SandAcl agent will continuously restart. Workaround is to remove the ACL from the configuration. (409638)
- With multiple configuration operations, such as add/modify/delete, on QoS policy maps, a restart of SandAcl agent may result into a core file and hitful operations for a brief period when 'feature qos mac' is in use. (410390)
- 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)

- Rollback from an unsuccessfully-applied CPU traffic-policy will result in no CPU traffic-policy being applied, rather than reverting to any previous traffic-policy configuration. This only affects the 'cpu traffic-policy <> vrf all' command. (421391)
- 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)
- After configuring a Vxlan interface, egress traffic will be tagged with vlan 0 if it enters the switch on an l2 subinterface. (432767)
- Pbr may be unable to recover from programming failure when there is initial TCAM bank exhaustion for a named policy, but at a later time, TCAM banks do become available and the policy is then applied to another interface on another forwarding ASIC. The workaround is to remove and reapply the named policy on all the interfaces. (434114)
- Flapping BGP Flowspec configuration on an interface where the neighbor is publishing large number of Flowspec matches may cause the forwarding agent to crash. (439838)
- When the outgoing interface of a VXLAN underlay route changes from one forwarding chip to another, it can negatively impact VXLAN overlay route convergence. (440233)
- While a line card is being power cycled, corrupt VXLAN encapsulated packets may be sent. (441141)
- Performing many updates to a PBR policy in a short amount of time may cause the SandAcl agent to restart unexpectedly. (441912)
- With PIM enabled on a VXLAN SVI configuration, reload of a switch might lead to incorrect flooding of PIM packets to remote VTEPs, causing PIM neighborship to not form between the MLAG pair, leading to misforwarding of routed multicast traffic in overlay. "shut", "no shut" of SVI would resolve the issue. (445280)
- Arp suppression in a multi-homed L2 EVPN MPLS configuration with more than two peers may cause loops. (445770)
 Series Affected: DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7280TR, DCS-7500R and DCS-7500R2 Series

- 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)
- If an access-list is used in a PBR policy and also in a QOS policy, editing this access-list may result in a failure to apply the QOS policy. The workaround is to use two separate access-lists for the PBR and QOS policies. (450595)
- The SandFap agent unexpectedly restarts when a PRBS type unsupported by the interface is configured; this will lead to traffic loss on all the interfaces. (450978)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- "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)
- SandL3Unicast agent may leak memory when unshared nexthop group entries are modified repeatedly. (478158)
- Changing the nexthop for routes may cause a momentary traffic loss during reprogramming. (478170)
- 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)

- TCAM bank reservation may not work after reload.
The workaround is to remove and reapply the reservation config for any existing feature after all the SandFap agents are up. (488675)
- With flex-route compression optimization configuration, exhaustion of hardware resources lead to continuous SandL3Unicast agent restart. (492870)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- VXLAN routing traffic may get mis-forwarded when there is rapid churn in the reachability of VTEPs (495364)
- With hardware accelerated sflow, if there are more than 14 sflow hardware accelerators in the system, the SandFap agents can restart if sflowAccel is disabled (496646)
- "hardware next-hop fast-failover" is not supported with MPLS EVPN flood adjacency resolved over an MPLS SR tunnel that is in turn resolved over TI-LFA tunnel. (503263)
- In a 6PE configuration, packet ingressing on an L3 subinterface with explicit null label and destined to this switch is not supported. (504582)
- When a DirectFlow flow is modified, the original flow entry might still be programmed in the TCAM along with the modified one. The original TCAM entry will not be visible using CLI. The workaround is to manually restart the SandAcl agent. (508034)
- With hardware accelerated SFlow enabled and more than 20 LAGs in the system and mirroring sessions configured, the Sand agent restarts continuously. (512973)
- 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)
- 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)

- With hardware fast failover configuration, ECMP to non-ECMP FEC transition leads to programming of backup path in hardware even when the active path is up. (540338)
- 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)

- LAGs with VLAN translation configuration may experience brief misforwarding of packets to incorrect VLANs during LAG member flaps.

A workaround is to avoid overlapping customer and bridging VLANs in the VLAN translation configuration. (553890)

- 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)
- If a QoS policy with rules of different ACL types fails to program in hardware for one of the ACL types, the subset of rules that was programmed successfully may not match incoming traffic correctly if other QoS policies are configured and applied in the system. (571599)
- 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)

- On a shaped subinterface tx-queue, configuring the bandwidth doesn't take effect unless no-priority is also configured. (588133)
- Adding a new member into a LAG with a high-scale VLAN translation configuration may cause some packets to be misforwarded to incorrect VLANs.

A workaround is to keep the new LAG member administratively shut before configuring it into a LAG. After a few seconds, unshut the member. (588883)

- When a PBR policy is applied to an SVI and new chips are added to the system, PBR may not work for the interfaces associated with the new chip if a policy is ever applied. (607599)
- While in Tap Aggregation mode, traffic forwarding may be disrupted on two-way (tap-tool) ports following a stateful switchover or a Mirroring agent restart. (614967)
- 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)
- When the EVPN VXLAN VLAN aware bundle configuration is changed, the SandL3Unicast agent may restart. (678131)
- When inserting a linecard, the adjacency update may cause momentary mis-forwarding of packets. (679075)
- 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)

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)
- AradLanz agent will restart if an invalid VOQ ID is read from the hardware. (557570)

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

- With CPU traffic policy configured with high VRF scale, the Acl agent may unexpectedly restart. (379101)
- SSU is not supported with port-channels containing members with auto-negotiation enabled. (407035)
Series Affected: DCS-7280SR and DCS-7280TR Series
- PFC pause frames, if received on /3 subport of Sesto2 master ports, will not be honored on the affected platforms. (411352)
SKUs Affected: 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 and DCS-7280QRA-C36S-M-R

- 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)
- The SandAcl agent may restart continuously if a linecard is replaced with another linecard with fewer forwarding chips. (505164)

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

- When an SVI is configured with an MTU value and a new interface is created with breakout cable and configured to be in that vlan, the new interface is not programmed with the correct MTU value. The workaround is to re-apply the MTU configuration on the SVI. (536927)

- 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)
- Reloading a switch in Tap Aggregation mode with multiple large tool groups (35 members or more) and continuous ingress traffic may result in the forwarding chip's multicast buffer being locked up; this will causing some traffic to not be forwarded.

A workaround is to reset the affected forwarding chip with the CLI command "reset platform fap full". (547590)

- If the forwarding agent is restarted after configuring a service profile on an interface, the forwarding agent may continuously restart unexpectedly.

A workaround is is to remove the service-profile and re-add it. (549335)

- When a device with VLAN translation entries configured on LAGs where the VLAN tag matching rules overlap with bridging VLANs on the same LAG is rebooted, there may be a window of time in which packets are mapped to incorrect bridging VLANs (571307)
- An interface with a link down debounce of less than 1.5 seconds configured may not respond to a local fault deassertion at the physical layer fast enough causing the link to spuriously flap.

Workaround is to configure a longer link down debounce time. (571797)

- 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 than the configured transmit power. (707195)

SKUs Affected: 7500R-8CFPX-LC

DCS-7280R2 and DCS-7500R2 Series

- Under high ambient temperature conditions, the switch may reboot due to a critical system temperature. (544313)
SKUs Affected: DCS-7280CR2A-60-F and DCS-7280CR2K-60-F
- Multicast traffic on non-default VRF will be lost for as long as 15 minutes when default VRF is multicast enabled without any L3 interfaces.

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

DCS-7280R3 and DCS-7500R3 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)
- Configuring IPv4 and IPv6 rules or configuring mpls match rules on a PBR policy may cause the SandAcl agent to restart. (407070)
- The traffic class in IPv6 packets egressing the switch after MPLS label termination might be incorrect. (420345)
- IPv6 over MPLS packets received with IPv4 Explicit NULL (i.e. Label=0) are misforwarded. (425904)
- 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)
- Even with a MAC ACL to deny IS-IS packets, MAC address of the IS-IS neighbor is learnt. (431856)
- After a restart of the forwarding agent packets originating on the CPU are dropped. (458206)
- MPLS packets with more than 10 labels, are not forwarded correctly. (467327)
- 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)
- When 'set dscp' action is applied to ECN capable packet, and there is no ECN congestion experienced by the packet, the ECN bits of the packet get reset to 0. (482029)
- Linecard hotswap or a switch reboot when the traffic is running may cause all packets to drop at ingress of one or more forwarding chips. Restart of the SandFapNi agent instance for the affected chips will resolve the issue. (483181)
- Forwarding of IPv4 and IPv6 multicast packets over MPLS Pseudowires is not supported. (493142)

- With large Nexthop group updates, in rare instances Packet can egress with wrong header data. (499845)
- High volume of mirrored traffic may be throttled aggressively; moreover, it may cause congestion and packet loss for control-plane traffic injected from the controller. (521159)
- 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)
- Packets forwarded by MPLS swap routes that are optimized as forward routes may get corrupted. (542196)
- 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)
- Mirrored packets may be dropped inadvertently due to L3 MTU check (562671)
- When VXLAN is enabled, L4 port match in ACLs on VXLAN decap packets does not work. (567246)
- Upon recovery from a PFC pause storm, Normal transmission of PFC frames on the interface and traffic class where we received the pause storm will not resume.

Workaround: To restore PFC functionality, disable and re-enable PFC on that interface. (571168)

- On a device with VXLAN configured, a large number of split horizon drops may result in some packets being discarded on neighboring devices due to CRC errors (574067)
- 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)

DCS-7020 Series

- When multiple Ipsec tunnel interfaces are shutdown, then "no shut" in a short period of time, routes associated with some tunnel interfaces may not be installed after "no shut". The workaround is to do shut/no-shut again on the individual tunnels where routes not installed. (377984)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

- When traffic is flowing from a 10G interface to a 1G interface and PFC is enabled on the ingress 10G interface, if the administrator changes the COS to TC mapping all the incoming traffic on the 10G interface is dropped. The only way to recover is by restarting SandFap. (400653)

DCS-7500R3 Series

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

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

- Ethernet over MPLS packets being terminated and decapped on a LDP pseudowire will be hashed as if the inner header is IPv4 rather than Ethernet (382763)
- The flexroute feature may override the forwarding decisions from PBR policies and BGP Flowspec matches. (544872)
- When a line card is swapped with a different model line card, and there was previously some LAG configuration for ports that do not exist on the new line card, the SandFap slice agent corresponding to that line card may restart repeatedly.

A workaround is to re-insert the old line card and wait approximately two minutes. Then, re-insert the new line card. (545221)

Series Affected: DCS-7500E, DCS-7500R and DCS-7500R2 Series

- Under some conditions, a forwarding chip's ingress buffer system can get in a corrupted state which can cause packets coming in that forwarding chip to be dropped. (548943)
- The Arp agent may continually restart when hardware accelerated sFlow is enabled.

Toggling sFlow, i.e. "[no] sflow run", may fix the issue. (571099)

- 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)
- With a port without any explicit speed configuration, many removals and insertions of a 1G transceiver in an SFP+ or SFP25 port, or of a 40G transceiver in a QSFP100 port, may cause the forwarding agent to restart and all links to flap.

The workaround is to apply a speed configuration on the port matching the supported speed of the inserted transceiver. (600618)

- 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

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

- When a line card is swapped with a different model line card, and there was previously some LAG configuration for ports that do not exist on the new line card, the SandFapNi slice agent corresponding to that line card may restart repeatedly.

A workaround is to re-insert the old line card and wait approximately two minutes. Then, re-insert the new line card. (536443)

Series Affected: DCS-7500R3 Series

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

Series Affected: DCS-7500R3 Series

- Under sustained heavy load, when more than one interface is receiving traffic at line rate 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)
- SandAcl can restart when line card is hot swap with higher line rate traffic in progress. (556308)
- 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)

- After a config replace LAG members may start dropping all untagged ingress packets due to dropVoqInPortNotVlanMember.

To work around, remove the interface from the LAG and then add the interface back in the LAG. (578380)

- 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)
- 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)

- 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)

DCS-7280R and DCS-7280R2 Series

- VXLAN packets that are decapped and bridged are incorrectly subject to the IP ACL applied on the ingress port. (221457)
- When ACLs are applied on a large number of interfaces (more than 1000 SVI interfaces), the SandHalo agent may restart with a SOCKET_CLOSE_LOCAL syslog message. (253015)
- Using AlgoMatch ACLs on routed ports, subinterfaces or SVIs may not work if QoS policies are applied on the same interface. (387932)
- IPv6 routing with virtual VTEP configuration will not work on platforms that run in AlgoMatch mode. Workaround is to change the access-list mechanism configuration to tcam (472142)
- 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)

7300X3, 7320X, 7368, CCS-720XP, DCS-7010, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260X, DCS-7260X3 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_HWHOP_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 either the peer-link port channel or in an MLAG 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

- Applying a large configuration containing many interface changes may result in Strata forwarding agent restart, which could cause traffic loss during the reconfiguration. (330840)
- Configuring more than five different speeds simultaneously on the system will result into continuous forwarding agent restarts. (339513)
Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX3, DCS-7050X3, DCS-7260X and DCS-7260X3 Series
- 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
- Enabling EVPN MPLS may result in continuous StrataL2 forwarding agent restarts as it is not supported. (363491)
- 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)
- 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
- Executing "platform trident forwarding-table partition" in quick succession may result in unrecoverable traffic loss. The work around is to do remove and reissue the command "platform trident forwarding-table partition". (368260)
- 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
- When PTP is configured on a both a primary lane and a secondary lane on the same front panel interface, a speed change on the primary lane can cause an unexpected Strata agent restart, causing all lanes for that interface to

become inactive until an ASIC reset is performed. This reset would result in all interfaces flapping. (390862)

- DHCP unicast packets which are not destined to switch are copied to CPU, which may forward the packet to destination along with one copy forwarded in hardware. This results in duplication of unicast DHCP packets. (393520)
- When the forwarding agent hitfully starts up, PTP packets sent downstream may have invalid timestamp and the switch may incorrectly advertise itself as the GM. Valid PTP packets and correct GM advertisement will converge post forwarding agent startup. (395916)
- 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)
- Switch may drop ingress UDP traffic destined to the switch MAC address when destination UDP port of the ingress packet matches VXLAN protocol default UDP port even if the destination IP address of the packet does not match any of the configured local VTEP IP addresses. (404138)

Series Affected: 7300X3, CCS-720XP, DCS-7050X2 and DCS-7050X3 Series

- DirectFlow flows with mirroring action may not be reprogrammed after a supervisor switchover or StrataMirror agent restart.

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

- Restart of StrataL2 may lead to continuous StrataL2 restarts. Workaround is to reload the system. (408138)
- DirectFlow flows do not take precedence over conflicting router ACLs. (408734)
Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX2, DCS-7050SX3, DCS-7050X2 and DCS-7050X3 Series
- DirectFlow flows can be uninstalled from a linecard on reaching the idle timeout for the flow even though matching traffic continues to match the flow on another linecard. The flow can be restored by re-programming it or disabling the idle timeout.

Series Affected: 7300X3 (411851)

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

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

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

- With IPv6 multicast configured, changing the speed of a port will lead to disruption of IPv6 multicast traffic (418250)

Series Affected: DCS-7260X3 Series

SKUs Affected: DCS-7050SX3-48YC8-F and DCS-7050SX3-48YC8-R

- 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)
- Hitless Strata agent restart may cause a Port Channel interface to flap if it contains members with auto-negotiation speed configuration (419369)
- 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
- 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

- When establishing a tunnel route when the hardware resources are exhausted but later recovered, the tunnel route may not be programmed in hardware. The workaround is to restart the StrataL3 agent. (433945)
- 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 ingress 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)

- LANS fabric thresholds may not be applied if the Strata slice forwarding agent restarts. Workaround is to remove and reapply the thresholds. (480422)
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

- 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)

- After a failed attempt to apply an ACL to an SVI, rollback to previous working config may leave the ACL in a state where it does not match expected traffic.

Reapplying the previous working ACL will correct the issue. (498355)

- 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)
- When the control plane is down during SSU, ARP packets from the remote VTEPs may not be forwarded to edge members of the VLAN. There is no workaround to this issue. (501085)
- 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)
- If a multi-ASIC linecard is hotswapped with a single-ASIC linecard in the same slot, the Snmp agent may continuously restart. The only workaround is to reload the system. (515667)

Series Affected: 7300X3 and DCS-7300X Series

- Restarting Strata with an ACL config that doesn't fit in TCAM can leave system in an incorrect state where subsequent attempts to program ACLs can fail. (522674)
- 'deny any any log' MAC ACL rule will not drop DHCP traffic and DHCP traffic will still be snooped. (523037)
- 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.

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

- A mirror to GRE packet may incorrectly be tagged with a VLAN tag corresponding to the native VLAN if the next hop for the GRE tunnel is through a trunk port that is in an SVI. A workaround to this issue is making the GRE tunnel reachable through an access port if in an SVI or a routed port. (533903)
- 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)
- VXLAN encapsulated ARP request and reply packets that have more than one dot1q tag will not be properly forwarded. (535958)
- 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. There is no workaround to this. (536312)
Series Affected: CCS-720XP, DCS-7050CX3, DCS-7050SX3 and DCS-7050X3 Series
- Card may shutdown at temperatures below critical temperature threshold. (536948)
SKUs Affected: 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 and DCS-7170-64C-R
- 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)

- Deleting a PTP-enabled port channel while PTP is in Boundary mode and then disabling PTP with the 'ptp mode disabled' command may cause the member interfaces of the deleted port channel to fail to send some PTP packets if PTP is set back to Boundary mode with 'ptp mode boundary'. Disabling and re-enabling PTP on the affected interfaces or port channel containing the affected interfaces with 'no ptp enable' and 'ptp enable' will restore normal operation. (558065)
- Loop protect fails to detect loops on MLAG port channels. (559665)
Series Affected: 7300X3, 7320X, CCS-720XP, DCS-7050CX3, DCS-7050SX3,

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

- 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)
- 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
- VXLAN BUM traffic received on a MLAG peer link may violate split horizon checks and result in sending traffic back to the VXLAN fabric resulting in MAC flaps in remote VTEPs.
Workaround is to flap the MLAG peer link port channel member that is affected. (567667)
Series Affected: 7300X3, CCS-720XP, DCS-7050QX2, DCS-7050SX2, DCS-7050SX3, DCS-7050TX2, DCS-7050X2 and DCS-7050X3 Series
- The forwarding agent may incur an additional restart during the Stateful Switchover (SSO). This may lead to a brief traffic outage. Conditions include: any peer fabric interface down (visible through `show platform trident connections`) and fabric modules unplugged. (568323)
Series Affected: 7300X3, 7320X and DCS-7300X Series
- Forwarding agent restarts, changing the forwarding table partition, or cold boots may result in the internal VLAN not being set on an interface.

As a workaround, a manual restart of the StrataLag will fix this. (569584)
- 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)

- Strata forwarding agent may continuously restart when a PBR policy with an access list "permit vxlan" is applied on an interface. The workaround is to remove "permit vxlan" rule from the matching criteria. (576286)
- In a MLAG setup, non routable packets destined to VARP MAC address coming in on a VXLAN VLAN (either on VXLAN edge port or encapsulated and entering a core port) can get looped over the peer-link. (577462)
Series Affected: DCS-7250X and DCS-7300X Series
- If uRPF is configured only in one VRF, deletion of that VRF may disable routing for all VRFs configured on the device. (578952)
- 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)
- Configurations with more than 8 interfaces configured on QSFP ports 53-56 cause continuous forwarding agent restarts leading to all interface to remain down.
The workaround is to avoid configuring more than 8 interfaces on QSFP ports 53-56 in affected releases. Such configurations are limited to:
 - 1) Four 1G (with QSA), 40G or 100G
 - 2) One 4x10G or 4x25G breakout, up to one 2x50G breakout, remaining ports at 1G (with QSA), 40G or 100G
 - 3) Up to three 2x50G breakouts, remaining ports at 1G (with QSA), 40G or 100G
 - 4) One 10G or 25G (with QSA), up to one 2x50G breakout, remaining ports at 1G (with QSA), 40G or 100G (593597)
 Series Affected: DCS-7050X3 Series

- DHCPv6 unicast packets which are not destined to switch are copied to CPU, which may forward the packet to destination along with one copy forwarded in hardware resulting in duplication of unicast DHCPv6 packets. (595070)
- A SSU reload may result in the SCD watchdog causing an extended traffic outage. (605183)
- 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
- In EVPN VXLAN setup with multicast routing enabled, ARP broadcast request and IPV6 neighbour solicitation packets may flood additional copies in the VLAN including a copy out the original source interface. (611692)
Series Affected: CCS-720XP, DCS-7050CX3, DCS-7050QX2 and DCS-7050SX3 Series
- Changing the membership of an interface from port-channel A to port-channel B might result in no control-plane forwarding to/from pre-existing members of channel group B during stateful switchover. (611967)
Series Affected: DCS-7300X Series
- ARP packets destined to VARP MAC on a L2 VLAN may get dropped instead of forwarding the packet on the VLAN. (613277)
- Routing of certain IPv6 prefixes may not work correctly when systems has one or more set of routes having following bit patterns. 1) Upper 96 bits are same and 2) Bits 65 to 96 has non-zero value. (614629)
Series Affected: 7368 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-7050SX3, DCS-7050TX, DCS-7050TX2, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2, 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-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7260CX, DCS-7260CX3, DCS-7260QX and DCS-7300X Series

- The StrataL3 Agent may unexpectedly restart when running the "show platform trident l3 software mgroups" on a switch with a large number of multicast nexthops. (636683)
- 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)

- 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)

- 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)

- A multicast group membership change that results in all receive traffic to CPU being dropped, may result in the resetting of the Strata agent.

There is no known workaround. (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)
- When segment security features are configured on a switch and interfaces flap, some Layer3 routed traffic trying to egress some L3 interfaces may result in L3 MTU check failures and get software forwarded. The workaround is to restart StrataL3 agent or reload the switch. (668718)
- 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-7010, DCS-7050X, DCS-7050X2, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260X and DCS-7300X Series

SKUs Affected: DCS-7304

- 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)

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-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7300X Series

- 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 and DCS-7300X Series
- When the number of remote VTEPs and VXLAN VLANs exceed 2000, executing "show platform trident vxlan next-hop" may result in unexpected restart of StrataL3 forwarding agent. Since the command is part of show tech, StrataL3 forwarding agent may restart periodically if scheduled show tech is enabled. Workaround is to reduce the scale of no of remote VTEPs. (705738)

- 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)

DCS-7010 Series

- 802.1x pause and guaranteed bandwidth cannot be configured at the same time. (91141)
- IP packets may be dropped or punted to CPU for VLAN with routing enabled, if L2 miss policy is also configured on the same VLAN.

Workaround is to restart the StrataVlanTopo agent. (503725)

- The switch will not boot if the real-time clock (RTC) battery is depleted. (675835)

Series Affected: DCS-7010 Series

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

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

A forwarding agent restart is required to forward traffic normally. (260334)

- If RPR or SSO is configured, Strata fabric agent restarts may cause switchover to occur. (368082)

SKUs Affected: DCS-7304 and DCS-7308

- 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. (383160)

- 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

- Ports with 1g transceivers will still count towards the limit of 5 speed classes even when disabled and may lead to continuous restart of the Strata forwarding agent. It can cause traffic outage on other ports.

Workaround is to remove the 1g transceiver. (338741)

- 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)

- A single event upset (SEU) in the TDM_CALENDAR* tables may lead to packet loss.

A forwarding agent restart is required to forward traffic normally. (547328)

- 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)

- Load balancing across ECMP members will not work when DLB is enabled and there are more than 31 members present.

When issue occurs "RX - Port bitmap zero drop condition" counters will increment in below command:

```
#show platform trident counters | nz
```

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

7368 Series

- An interface may fail to link with a MAC Local Fault indication.

The workaround is to run interface configuration command "shutdown" followed by "no shutdown". (341316)

Series Affected: 7368 Series

- 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

- Lossy traffic could be discarded at ingress interface instead of egress interface, if the ingress interface has "priority-flow-control mode on" configured. Because of this, the lossy priorities on the ingress port are vulnerable to head-of-line blocking. Buffer build-up on the lossy priorities is a necessary condition for the problem. (514254)
- Speed change on a port-channel member can cause a flap in BFD and BGP sessions for the entire port-channel (520638)
- System may experience transient traffic drop on some of the existing routes, when new routes having some specific distribution are added to the system. (524510)
- Executing the CLI command "show platform trident l3 shadow l3-iif-profile" causes StrataL3 to unexpectedly restart. (582412)

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)
- A single event upset (SEU) in the ES_PIPE* tables may lead to a forwarding agent restart, after which the switch forwards traffic normally. (547326)
- 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 line rate 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-Fix system agent to restart following which the interface continues to transmit properly. (245501)
- When the same VRRP group is configured on more than one VLAN, the switch may route packets to VRRP MAC addresses that correspond to groups for which it is not master on some VLANs. The workaround is to use different VRRP groups across VLANs. (444826)

- Packets sourced from an authenticated MAC on a Dot1X MBA authenticated port get punted to CPU. Unicast packets destined to authenticated MACs egress out tagged on a Dot1x MBA authenticated access port. Restarting StrataL2 will fix the issue. (532463)
- Vxlan multicast decapsulation will not work for packets that are sourced from unknown source VTEP. (538407)
- Continuous flapping of L3 interfaces could cause the StrataL3 forwarding agent to restart, resulting in a momentary traffic loss. (550872)
- When the per VLAN routing feature is enabled, BFD packets will be dropped in HW resulting in existing sessions going down or preventing new sessions from getting established.
Workaround is to disable per vlan routing feature. (568298)

CCS-720XP Series

- IPv6 BGP peering may not be formed due to HLIM=1 packets not reaching to CPU if
 - there is a speed change on any port,
 - there is a transceiver hot swap
 - IPv4/6 multicast routing is enabled.
 Workarounds:
 - Use HLIM>1 BGP packets by configuring `neighbor <IP Address> ebgp-multihop <hlim>` on the BGP peer.
 - Or if HLIM=1 packets are to be used, then for the first two cases turn off IPv4 and IPv6 unicast routing on all VRFs and then enable back again. There is no workaround for the third case. (563045)

DCS-7050X3 Series

- Hardware port resource (logical port) exhaustion may cause one of the following:
 - One or more interface to discard all packets.
 - Forwarding agent to restart (including continuous restarts).

Hardware port resource exhaustion is identified by STRATA-4-HW_PORT_RESOURCE_FULL syslog. The switch may run out of hardware port resources (logical ports) in following cases:

- Several 4x10G or 4x25G breakout ports are configured.
- 40G-only transceivers inserted with a default speed configuration (e.g., 40GBASE-SRBD).

The workaround is to reduce the number of 4x10G or 4x25G breakout connections to 2 (one breakout per following interface groups Ethernet49-52 and Ethernet53-56) and 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. (546409)

- 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

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)
- Third party QSFP twinax cables may fail to initialize. This will result in them not showing up in 'show inventory'. Susceptible cables will have lower page 0 byte 2 bit 2 set to 0. This can be confirmed with the 'show idprom transceiver ethernet <intf> extended' while running an unaffected release. (556437)
SKUs Affected: 100GBASE-CR4 and 40GBASE-CR4
- Removing a transceiver within approximately 3 seconds of inserting the same transceiver from a QSFP200, QSFP-DD or OSFP transceiver slot may cause XcvrAgent to unexpectedly restart. (569999)
- 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)

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

- Egress ACL on SVI for BUM traffic will not match ACL rules for packets being sent across fabric interfaces. (284543)
- A single-event upset (SEU) on the table `ING_VLAN_TAG_ACTION_PROFILE_1` may not be corrected, which may lead to traffic loss. Restarting the StrataL2 agent is required to clear the error and forward traffic normally. (287239)
- Configurations with more than 8 interfaces configured on QSFP ports 53-56 cause SFP ports 41-48 to become inactive. Upgrades to EOS-4.23.1F, EOS-4.24.0F or newer cause SFP ports to become active, and QSFP ports 55-56 to become inactive instead.

For example, configuring "speed forced 25gfull" on Ethernet53/1-56/4 causes Ethernet41-48 to become inactive. Upgrading to a newer EOS release causes Ethernet41-48 to become active and Ethernet55/1-56/4 to become inactive.

The workaround is to avoid configuring more than 8 interfaces on QSFP ports 53-56 in affected releases. Such configurations are limited to:

- 1) Four 1G (with QSA), 40G or 100G
- 2) One 4x10G or 4x25G breakout, up to one 2x50G breakout, remaining ports at 1G (with QSA), 40G or 100G
- 3) Up to three 2x50G breakouts, remaining ports at 1G (with QSA), 40G or 100G
- 4) One 10G or 25G (with QSA), up to one 2x50G breakout, remaining ports at 1G (with QSA), 40G or 100G (380984)

SKUs Affected: DCS-7050SX3-48YC8

- 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)
- When a DirectFlow flow entry with both an L2 rewrite modification and any other field is both matched on and modified in the same flow, traffic may egress without performing L2 rewrite modifications. (414504)
Series Affected: 7300X3 Series
- 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)

- When Evpn multihoming is configured, BUM traffic to some Vxlan VLANs may get dropped and copied to cpu if the multicast index allocated is greater than 8191. The workaround is to restart L3 forwarding agent. (419207)
- When EVPN multihoming is configured, IPV6 multicast traffic may get looped back to ethernet-segment on non designated forwarder and lead to loss of bandwidth. (443123)
- Performing SSU may result in extended traffic outage, leading to total traffic loss being slightly over 1 second. (477139)
- If EVPN VXLAN multihoming is configured, and link towards ethernet-segment goes down, then some macs that reside behind that link may get stuck in hardware. As a result, traffic destined to those macs may get dropped. The workaround is to reboot the switch. (495546)
- Having any ACLs with UDF qualifiers as well as any of the following LAG/ECMP hashing features present in the configuration: RRoce hashing or TTL/HopLimit hashing, will result in incorrect operation of the LAG/ECMP hashing where the switch incorrectly hashes on IP payload. To avoid this problem the conflicting features should not be configured at the same time. (497016)
Series Affected: DCS-7050CX3, DCS-7050SX3 and DCS-7300X Series
- With GRE tunnels or IPinIP tunnels or Nexthop-group tunnels configuration along with subinterface in the tunnel core interface, tunnel encapsulated packet will not be tagged with dot1q tag as configured in the subinterfaces, instead packet will go out tagged with internal vlanId assigned for the subinterface. There is no workaround for this issue other than removing subinterface configuration in the core interface. (523365)
Series Affected: 7300X3, CCS-720XP and DCS-7050CX3 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)

- A single event upset (SEU) in the TDM_CALENDAR* tables may lead to packet loss.

A forwarding agent restart is required to forward traffic normally. (547329)

- When the switch is reloaded with EVPN All Active config, the BUM traffic from some of the VTEPs may not be pruned correctly resulting in duplicate delivery of packets on some of the shared ethernet segments. The workaround is to restart the StrataL3 forwarding agent. (560034)
- 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

- During bootup or L3 agent restart, one or more linecards may not be programmed with all the routes. This will cause traffic loss or misforwarding. There is no workaround for this issue. (568757)

SKUs Affected: 7300X3-32C-LC

- 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)

- When switching from EVPN configuration to static VXLAN, remote MACs may not be learned even though configuration is switched to dataplane learning.

Workaround: Flap the VXLAN interface or restart StrataL3 agent. (596455)

- Performing SSU (fast reload) from an older release may lead to complete traffic loss. Workaround is to issue chip reset using "platform trident reset" command. (596509)
- 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)
- 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)
- 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)

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)
- If the capacity of prefixes/routes that can be accommodated is reached, under some conditions,
new routes are not programmed even if the the number of prefixes/routes are

brought down within
the capacity.

Workaround is to reduce the amount of learned routes and restart
XpL3Unicast agent. (372207)

- 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)
- When using the tcam-acl profile, ACL removal from an interface may not take
effect in some situations. Workaround would be to restart XpAcl. (431682)
- When using the tcam-acl profile, if there exists an interface with multiple
ACLs applied, the rules are not guaranteed to be programmed in the correct
order. Workaround would be to avoid having interfaces with multiple ACLs
applied. (431684)
- 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

Limitations and Restrictions in 4.22.8.3M

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)

vEOS Router / CloudEOS

- NIC i40en driver version 2.2.4.0 is required on ESXi 7.0 host for SRIOV virtual function trust mode on/off functionality. (680174)

- NAT interface configuration mode is not supported in SFE. This needs to be kept in mind when migrating to SFE mode from Sfa.

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

CVP

Telemetry

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

CVX

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

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

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

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

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

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

- Intercept hosts on trunk ports with native vlan are only supported on 7050X2 platforms (257580)
- The Macro-Segmentation Service (MSS), requires MLAG fast redirection be enabled on the service MLAG TOR pair, when configured to work with a L2 transparent FW. This configuration minimizes traffic loss when any individual interface in the MLAG port channel goes down. (268291)

- The MssClient agent, can restart unexpectedly, when the MacroSegmentation Service, deployed with a L2-transparent firewall and is configured to intercept more than 10,000 end-points on a single switch (a number that's far greater than the total number of intercepts possible on the switch). (283221)
- 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)
- Add support for Port-Channel interfaces for ContainerTracer. (429898)
- 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)
- Do not create an alias named "alias". (65392)
- 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)
- The boot config command takes effect immediately inside a config session and cannot be aborted (346171)
- EOS support alphanumeric VLAN names, however ODL gives an error when reading the leaf /interfaces/interface/routed-vlan/config/vlan if the name contains alphabetic characters.

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

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

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

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

SKUs Affected: DCS-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 updating an interface's description, user needs to specify it for both base interface and subinterface with index=0, this is not desired.

Fix is to treat the 'description' of subinterface (with index=0) as read-only and ignore it when handling OpenConfig update requests. (376474)

- 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)
- If the config script creates a Port Channel interface during ZTP, "ipv6 address auto-config" will get added to the interface's configuration. (380384)
- 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)
- The commands "show (startup|running)-config [all] section <section>" may cause the cli session to exit if the output is big (>~1MB) and paginated. (417111)
- 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)
- 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)
- 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)
- Installed licenses without corresponding licensing features configured might not be visible from the CLI (504793)
- If Octa, OpenConfig or TerminAttr are started under low memory conditions, they can cause a spike in CPU usage before unexpectedly restarting. (534841)
- 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 workaroud 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 workaroud 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 workaroud 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)

- Power Loss wrongly recorded as the reload cause on this platform instead of the correct reload cause. (604561)

SKUs Affected: 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-7280CR3K-32D4, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4-F and DCS-7280CR3K-32P4-R

- "show reload cause" may show "Unknown" with no debug information if the reload reason is an internal CPU error. (623392)

SKUs Affected: 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-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R,



DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, 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-7500E, DCS-7500R, DCS-7500R2 and DCS-7500R3 Series
- Sampled flow tracking (IPFIX) and sFlow export flow records with incorrect nexthop for the flows using tunnel nexthops with hierarchical FECs (HFECs). (677481)
- 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)

## 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
- 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 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-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F and DCS-7280CR3K-32P4-R

- Interfaces with BCM82391 PHY do not support auto-negotiation. (537530)  
SKUs Affected: 7500R-8CFPX-LC and 7500RM-36CQ-LC
- MACsec frames with SL less than 24-bytes are silently dropped on ingress (714597)  
SKUs Affected: 7500R-8CFPX-LC and 7500RM-36CQ-LC

## Layer 2

- LACP PDU fast rate ("lacp rate fast" on an Arista switch) should not be configured on any port in an MLAG pair, or any port connected to an MLAG pair. (13950)
- Static ARP inspection may not work on modular platforms after SSO switchover. The workaround is to disable and re-enable static ARP inspection after switchover. (244477)
- The command "lacp rate fast" is not supported on modular systems with stateful supervisor switchover (SSO) enabled. Neighboring devices should not have "lacp rate fast" configured on ports connected to the said system. (248121)
- In a MACsec profile, when a primary key's CKN is changed, then operating primary continues to be principal actor even when a fallback key is configured till the time the new primary establishes a successful MKA session. There is zero traffic disruption. (275678)
- Port Security Protect does not support trunk ports. If a port is secured to source MAC <A> in VLAN <X>, traffic from source MAC <A> in other VLANs of the trunk port will also be allowed. (344773)
- While performing SSO, a device's peers might not receive LLDP packets within default TTL (120 seconds) and those peer devices may not maintain LLDP status information for the device undergoing SSO. After SSO is complete, the peers will again have LLDP state available.

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

- An MLAG switch cannot be used as an EVPN multi-homing provider edge. When a switch has MLAG enabled, MAC addresses from ports with an ethernet segment configured are advertised through BGP as if they were single-homing, and local ethernet segments are not advertised through BGP. (397867)
- Even with fallback to unprotected traffic configured, traffic loss in the order of milliseconds may be seen during initialization (491232)
- Upon programming a large number (~250) of VTEPs in a flood set at once, the L2Rib agent may restart. Workaround is to program them in batches, waiting a few seconds between each batch. (544051)

- 802.1X is supported on LAG member ports using EAPOL authentication. (544556)
- The SyncE wait-to-restore (WTR) timer is armed when administratively shutting / unshutting an Ethernet interface. This is inconsistent with disabling / enabling SyncE on the interface in which case the WTR timer is not armed.

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

## Layer 3

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

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

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

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

- On Macro-Segmentation Service (MSS) deployment with L3 firewalls, a flow entry that fail to get programmed due to lack of TCAM space, is not

programmed when TCAM space becomes available at a later time. The workaround is to "shut" and "no shut" on the direct flow config. (372232)

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

- The update wait-install feature may not work as expected if there are routes which fail programming due to hardware resource exhaustion. The routes which failed programming may still get advertised out to peers. (372765)
- DHCP packets will not be processed by EOS DHCP Server when DHCP Relay is relaying packet from one SVI where DHCP Relay is configured to another SVI where DHCP Server is configured. (389017)
- When a sampled packet egresses a trunk port and the path is ECMP, sFlow will arbitrarily pick one of the ECMP next hops for the next hop IP of the router extension. (415747)
- DHCP clients will fail to obtain IPv6 leases when "ipv6 dhcp relay destination" is set to IPv6 multicast address. (426796)
- 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)

## 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)
- BGP session over indirectly connected interface that was down due to lack of IGP path to peering address may take around one minute to re-establish once the IGP path is available. (287628)
- 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)
- There may be intermittent traffic loss when the routes resolving BGP nexthops is deleted in RIB, eventhough there is a back up route avaiable to resolve the nexthops. Bgp routes should get resolved through the backup route without any intervention and the traffic should recover automatically. (391399)
- 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 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)

## Rib

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

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

## MLAG

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

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

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

## MPLS

- The MPLS agent may continuously restart on a router with high BGP LU Tunnel scale. (377939)
- BGP LU Tx MPLS routes corresponding to BGP LU advertisements with multiple labels are not supported. (395029)
- The "color-dscp-mappings" CLI configuration sub-mode of "segment-routing" in "router traffic-engineering" has no effect.

The feature that it configures has not been implemented. The configuration sub-mode is now considered deprecated, but existing configurations will not be removed on upgrade. (402515)

- 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)

## 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)
- Traffic loss may happen upon Pimsm/Pimsm6 agent restart. (377375)
- 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)
- The issue happens when PimReg agent restart while one of the smash path hasn't loaded yet.  
This should happen rarely, and agent restart again should fix the issue. (381992)
- In multi-agent mode, MSDP can fail to find an RPF peer for another MSDP speaker, if there is no direct peering and the best available BGP path to the MSDP speaker is ECMP. This is because ECMP paths are not available to MSDP. This can result in failure to accept SA messages from that speaker.

The workaround is to add an default MSDP peer configuration. (385304)

- 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)

- Maximum vif created in kernel is intended to be 512, due to this bug, it is set to 511, one of them is used for pim register interface, leaving 510 vif available to use. CLI allows user to configure up to 511 pim enabled interface, so one of the interface won't get vif from kernel. Due to this issue, none of the multicast traffic for this interface punting to CPU will be dropped in kernel. The workaround is to configure up to 510 pim enabled interfaces per vrf.

Update your release note here. (408111)

- Multicast routes are not created when a source is transmitting only fragmented multicast packets. Switch to using SFE for multicast software forwarding. (605326)

## SwitchApp

- IGMP snooping filters are not supported when SwitchApp over forwards IGMP reports. Reports are over forwarded if the "igmp flooding" CLI is enabled through the FPGA instance CLI, or if the FPGA profile does not support correct forwarding of IGMP reports. (601917)

## vEOS Router / CloudEOS

- The /mnt/flash/cloud\_ha\_config.json based configuration for the CloudHa agent is not supported from this release. If you are upgrading from an older image (< 4.20.5) please reconfigure HA using EOS CLI. Please see information on converting json file based config to CLI commands in the vEOS Router Configuration Guide. (264660)
- While vEOS instances in an AWS cloud can be created using either a Bring Your Own License (BYOL) or Pay As You Go (PAYG) format, there currently is no show command in vEOS to easily help the user or support determine which mode this instance uses. (272649)
- Tx/Rx ring parameters of an "Elastic Network Adapter (ENA)" interface cannot be set through config.json in vEOS. (276382)
- vEOS supports up a maximum of 8 physical cores. With hyperthreading enabled it is 16 vCPUs (278286)
- In dpdk based vEOS(MODE=sfe in /mnt/flash/veos-config), SR-IOV mode doesn't work with Intel x520/82599 on ESXi. As a workaround use mixed mode instead of the SRIOV only mode in ESXi (296718)
- Hot detaching network interface instances is unsupported. Reboot the instance after detaching a network interface. (306132)

- In dpdk based vEOS, the supported scale is 100 dynamic nat flows can be created per second. (407966)
- 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 Telemetry 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)

## 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 PAACL 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 PAACL 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 to 7Kbytes. (258823)
- The tcpdump session CLI command is not able to start the tcpdump session for mirror sessions. The workaround is to execute the tcpdump command directly in bash mode. (273481)

- 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)
- When the mirror destination is a GRE tunnel with Port-Channel as the nexthop interface, the packet mirroring may not happen consistently due to some of the Port-Channel members not programmed in the mirroring hardware table. There is no workaround for this issue. (380328)
- firewall profile on DCS-7170 platform does not support IP multicast (401134)
- While programming mirror destinations in hardware, if the interface corresponding to the mirror destination experiences a link flap, it may not be programmed in the hardware. The workaround is to reconfigure the mirror destination in CLI after removing it. (414534)
- VXLAN encapsulated ieee reserved mac addresses are not trapped to CPU and hence any control protocol that is VXLAN encapsulated will not work. (423963)
- If the members of a mirror destination Port-Channel disappear while programming the mirror destination in hardware, the BfnMirroring agent restarts due to an assertion failure. After the restart, the system recovers and the hardware entries are programmed correctly. (424633)

## DCS-7170 Series

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

## CCS-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

- 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-24ZY4-F, CCS-720XP-48Y6-F and CCS-720XP-48ZC2-F
- 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)

## 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)

## DCS-7020, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7500E, DCS-7500R, DCS-7500R2 and DCS-7500R3 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)
- 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)
- The system can run out of hardware resources to program ingress mac ACLs for a configuration that includes a large number of interfaces. (127383)  
Series Affected: DCS-7500R Series
- If a double tagged frame is received on an untrusted or DSCP-trusted interface, the CoS is retained if the outgoing frame is tagged, unless the frame is routed (131614)
- IPv6 Egress ACL deny logging is not supported on subinterfaces. (146487)
- PBR policy is supported only on interfaces in the default VRF. If a PBR policy is configured on an interface in a user defined VRF, the behavior is undefined. (148895)
- Disabling IPv4 routing on an L3 interface is not supported when IPv6 routing is enabled on that L3 interface. When "no ip routing ipv6 interfaces" is configured to enable the forwarding of IPv4 packets over IPv6 nexthops over interfaces that do not have IPv4 address, but only a IPv6 address, it will also allow routing IPv4 traffic over these interfaces. (151356)
- When 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)
- 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, 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)
- With uRPF configuration, packets sent by the Cpu that match a route forwarding into MPLS, GRE, or a Vxlan tunnel are incorrectly dropped in the hardware due to uRPF violation. (343891)
- Stale telemetry flow-spec statistics may be seen until the next "show flow-spec" command is issued. (345195)

- Filtered mirroring is not supported on linecards or fixed-systems operating in TAP Aggregation mode when traffic steering is enabled. (349883)
- 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)
- 'show platform fap mpls mroute' takes between 3-18 minutes for output to appear on the screen (395867)
- If a security-ACL is configured on a sub-interface beyond the supported limit of ACLs,

SandAcl might restart continuously until the number of configured ACLs is brought back within the supported limits (398815)

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

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

- PIM sparse mode feature which allows installation of outgoing interfaces on the non-DR for a faster failover does not work on L3 sub-interfaces. This feature relies on an egress ACL to stop multicast traffic from going out of the interface. Multicast egress ACLs do not work on L3 sub-interfaces due to a platform limitation. If the feature is configured, there will be duplication of traffic. The workaround is not use the feature on such interfaces. (405696)
- When a policy-map with match on inner VLAN is programmed on an interface, if single tagged packet traffic is sent to that interface, the behaviour of that packet is indeterministic as inner VLAN tag is calculated using an offset and vlan-tag-format qualifier doesn't differentiate between single and double tagged packets. (415587)
- When MPLS label is pushed on the traffic ingressing CoS trust port then EXP bits are not derived from TC based on packet's COS. (417108)
- Byte counts are not supported for IPSec tunnels. (419031)
- SflowAccel supports up to 200 VRFs. (423892)
- PAUSE frames with correct DMAC=01:80:C2:00:00:01, Ethertype=0x8808, opcode=0x0001 are never forwarded and are always consumed by the hardware. (426047)
- A port transmitting packets smaller than 100B that were tunnel terminated on the switch might be limited to ~90% linerate (429797)
- Mirroring to GRE tunnel is not supported on Jericho and Jericho2 mixed systems (438101)
- 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)
- 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
- 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 switches respectively. (502940)
- VRF label termination into non-default VRF is not supported for multi-label MPLS packets. (504763)
- Egress counter features configured via "hardware counter feature ..." do not include packet or byte counts for packets sent from the local CPU. (525380)
- With flex-route compress option enabled or the flex-route internet profile along with a hardware tcam profile including the flex-route feature, if the

default route resolved over tunnel, it may lead to SandL3Unicast agent to restart continuously. (525611)

- 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)

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

- 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)
- 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 match untagged traffic erroneously.

Add the vlan-tag-format key-field to features that filter on VID in the active TCAM profile. (597603)

- VLAN matching in ACLs attached to monitor sessions may result in false matches on untagged traffic. (598186)
- 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 applied than can be configured on any given FAP in the system.

A workaround is to reduce the number of mirroring ACLs applied; or apply smaller subset of mirroring ACLs to different source ports on different FAPs. (628026)

- SflowAccel does not support collectors reachable over cross AF routes. (640292)
- 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)

## 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

- An IPv6 MPLS packet terminated by L3 EVPN MPLS configuration will not be forwarded. The workaround is to disable IPv6 exact match host routes using 'no platform sand ipv6 host-route exact-match' and disable IPv6 route optimizations for prefix length 128. (201414)



Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series

- When a linecard is inserted in a chassis there can be a temporary marginal drop in fabric throughput for other cards. (205880)
- On the system with DCS-7500R linecards, if DCS-7500R2 linecards are inserted, routes may consume more resources in the hardware routing table. (215393)
- Jericho chip buffering DRAM test at agent startup does not catch bad parts while it should fail to force RMA of bad parts. (218971)
- IPv6 multicast routes are stored in two separate TCAM resources. The group IP is stored in one TCAM while the source IP is stored in another. Since one TCAM contains the source IP it will contain an entry for every multicast route. However, the group IP TCAM will reuse entries for duplicate group IPs. For example, if you have routes (S1,G1) and (S2,G1), one TCAM entry is required for the group IP and two TCAM entries are required for the source IP. This means it is possible that more TCAM resources are allocated to the source IP lookup than the group IP lookup. The scaling achieved during the initial distribution of multicast routes will be preserved. However, if the distribution later changes, optimal scaling for the new distribution may not be reached. (257364)
- The CPU traffic policy feature does not support filtering on IPv6 extension headers. (364996)
- Ingress ACLs and PBR are not supported for MPLS tunnel terminated traffic using VRF label. (366695)
- VRF terminated traffic is not sampled with sFlow. (369010)
- Drop-precedence thresholds are not supported on mirror session ports. (388932)
- We do not support incoming multicast LDP packets with a destination multicast MAC address.  
RFC 5332 doesn't mandate the use of sending labeled multicast traffic in a multicast ethernet frame. Multicast MAC destination address could be used in upstream label assignment. (393406)

Series Affected: DCS-7020, DCS-7280CR2, DCS-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)
- 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)
- VXLAN is not supported with greater than 256-way ECMP in the underlay. (652266)

## DCS-7280R2 and DCS-7500R2 Series

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

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

## DCS-7280R3 and DCS-7500R3 Series

- "Drop" MAC entries only drop packets whose destination MAC address matches the "drop" MAC entry. Packets whose source MAC address matches the "drop" MAC entry are not dropped by the "drop" MAC entry. (341123)
- sh policy-map type control-plane copp-system-policy shows wrong hardware programming status. (367000)
- 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)
- Cli to change tail-drop threshold for 200/400G ports is not supported by software. (396029)
- 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-classes (aggregate rate nearing line-rate) to port which has been configured with Strict-Priority scheduling. (406550)
- Per TC based VOQ tail-drop thresholds isn't supported in software. (413344)
- 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)
- 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, 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-7280CR3K-32D4, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4-F and DCS-7280CR3K-32P4-R
- SFlow samples of packets which are dropped by an ACL will have incorrect output interface indication. (470650)
- ECN bits may not be correctly copied from outer header when packets are routed after VXLAN decapsulation. (489998)
- The sFlow flow samples generated for VXLAN decapsulated packets, have the first 6 header bytes of the sampled packet set to zeros. (502957)
- VXLAN VARP VTEP with IPv6 overlay routing requires vxlan-routing TCAM profile to be applied on the switch. (503496)

- 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, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3K-32D4 and DCS-7280CR3K-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)
- 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 Isec tunnel interface for policy-based routing nexthop is not supported. (378861)  
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Isec tunnel interfaces are not supported as part of nexthop-group. (378890)  
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Dynamic NAT, AKA Port Address Translation, is not supported for IPSec tunnels. (384624)

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

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

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

- Vxlan decapped packets will not match rules that match on 'tcp established' in egress acls (226582)
- VXLAN and GRE tunnels cannot be configured at the same time. (248239)
- If a PTP packet is mirrored to a GRE destination which has time stamping enabled, the timestamps on such packets are not reliable. (280362)
- If a packet is mirrored to a GRE tunnel, but the forward copy is trapped to the CPU, then the ingress VLAN ID carried in the GRE key will be set to 0 when the mirrored copy is encapsulated in the GRE envelop. (434871)
- 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)
- 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
- L4 destination port matching in egress IP ACLs does not work on untagged bridged traffic. (609461)
- Support for termination of GRE or Vxlan encapsulated packets with a configuration consisting of a tcam profile with flex-route feature. (628254)
- 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)

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

- QoS policy-map is not supported on Pseudowire decapsulated traffic flows. (574600)
- When there is congestion in Tc6/Tc7 across multiple queues. Tc6/Tc7 Queues can hog up entire DRAM. This can lead to DRAM rejects of lower TC traffic of unrelated stream thus causing some discards. In the current case we had 2 to 3% lower instead of 100% throughput of data traffic (591523)
- 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)

## 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-7060, DCS-7060X, DCS-7060X2, 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-7050SX3, DCS-7050X2 and DCS-7300X Series

- If ingress ACL is applied on an interface, traffic sent to CPU from that interface is falsely counted as InAclDrop in "show int counters ingress acl drop". (203465)
- Clause 37 auto-negotiation on SFP28 and QSFP28 ports is not supported. A workaround is to use 1G forced speed configuration on the affected ports. (206778)

Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX3, DCS-7050X3, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7260CX3 Series

- When "reload fast-boot" is performed from a release prior to 4.17.2, switches with BGP peering connections with the switch undergoing "reload fast-boot" could experience BGP KeepAlive timeouts before they see port flaps. This could result in longer than expected data plane downtime during the upgrade. (214700)

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

- A multi-bit ECC error in forwarding ASIC packet memory cannot be automatically corrected. A hitful restart of the slice agent managing the forwarding ASIC or a reboot of the system is required to clear the condition. (214950)
- If ip-ttl is the only enabled field for port channel hashing, after reload hitless, packets may get hashed to a different link than before the reload. (223568)

Series Affected: DCS-7050TX, DCS-7050X and DCS-7060X Series

- A configured mirroring ACL may not be applied in hardware. This can be confirmed by checking the output of 'show platform trident tcam mirror'. (236908)
- When a VXLAN encapsulated packet is received on a VXLAN VLAN where PBR policy is applied, the switch will not PBR forward the decapsulated packet. (244093)
- Sflow output interface determination will not work when the output interface is a subinterface. (250424)
- A switch may not be able to bridge or route VXLAN traffic if it hits one or all of following conditions:
  - Nexthop resources are exhausted as indicated by syslog `VXLAN_HWHER_NEXTHOP_RESOURCE_FULL`.
  - MAC table overflow as indicated by show platform trident mac-address-table missing.

To recover from this condition reduce the config scale and flap the VXLAN interface. (253601)

- Flexible port-channel hashing may not hash accurately when the inner L4 header option is configured for IPV6 GRE packets. (267487)
- If a packet stream is policed by multiple policers, the policed rate may be lower than any of the configured policer rates. (271046)
- To perform filtered mirroring of a packet based on VLAN from an interface configured to perform VLAN translation, the ACL must be programmed with the translated VLAN. The mirrored packet will be the same as the packet at the source interface. (278599)
- BGP neighbor sessions may flap when URPF is configured on an interface. (279113)
- To perform filtered mirroring of a packet based on VLAN from a subinterface source, the ACL must be programmed with the internal vlan of the subinterface. The mirrored packet will be the same as the packet at the source interface. (280943)
- Matching on inner VLANs is not supported in mirror ACLs. If an entry is specified matching on inner VLAN, it will be ignored and permit all packets. (284371)
- SSU is supported when upgrading from the associated release in each release train: 4.18.8, 4.19.8, 4.20.7, 4.21.0. (289548)
- Issuing "no shutdown" on a member link of port-channel may send duplicate packets for sub-second on that member link (291514)
- DirectFlow (and features such as the MacroSegmentation Service that use DirectFlow) are unsupported on the 7260, 7300 series of switches. (296715)
- When primary vlan and secondary vlan are part of different MST instances, the hardware programming of lag will not happen. (300119)
- Not all interfaces in the groups Ethernet49-Ethernet52 and Ethernet53-Ethernet56 can be broken out in to 4x25G, 4x10G or 2x50G mode due to MAC resource limitation on the system. Interfaces without MAC resource will be marked inactive and 'STRATA-4-HW\_PORT\_RESOURCE\_FULL' will be logged in syslog.

Speed change on interfaces in the groups Ethernet49-Ethernet52 and Ethernet53-Ethernet56, can free up the MAC resources to make inactive interfaces active in that group. (306121)

SKUs Affected: DCS-7050SX3-48YC8-F 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
- 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 fast-boot" is not supported with virtual IP address config. (411323)
- 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)
- Configuring the same IP address for GRE Tunnel Interface source and decap group destination is not supported. (431026)
- 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-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 and DCS-7260CX3 Series
- With VXLAN dataplane learning disabled, locally learned MAC addresses will not age out if VXLAN encapsulated packets from remote VTEPs are received with inner packet's source mac as the locally learned MAC. The locally learned MAC entry needs to be cleared manually by issuing 'clear mac address-table dynamic vlan <VLAN> address <ADDRESS>'. (460607)
- If 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)
- Egress mirroring of multicast packets that are going to be encapsulated on the switch, independent of mirroring, is not supported. (512880)
- 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-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series
- The traffic load across ports comprising a port channel(s) may be uneven for certain traffic patterns and/or port channel configuration and state.

Workaround:

- change the seed in the hashing algorithm until satisfactory result is achieved by executing configuration CLI command:

"arista(config)# port-channel load-balance trident <seed>" (555833)

- Implicit ICMPv6 permit rule(s) in TCAM is not seen in output of "show platform trident tcam acl detail" command.

Workaround is to check the entry and packet hits for the entry using output from "show platform trident tcam detail". This command shows all ACL entries in TCAM and packet hits per entry. (563228)

- Enabling L2 protocol forwarding of LLDP on a given port will implicitly enable forwarding of EAPOL and MKA packets on the same. (569478)
- BGP packets are ignored by IP counters if hardware IP counters are not enabled. (599275)
- 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)
- 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 Series

## DCS-7010 Series

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

## 7320X, DCS-7060, DCS-7060X and DCS-7260X Series

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

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

- Cut-through mode is not supported on the SFP+ interfaces. (182728)
- Cut-through mode is not supported at 1G speed. (219427)
- Mirror To GRE destinations with multiple ECMP nexthops or LAG is not supported. A single link of the LAG or ECMP nexthops will receive all mirror traffic. (639980)

Series Affected: 7368 Series

## DCS-7260X3 Series

- Configuring 100G, 50G, 40G, 25G and 10G simultaneously might result in unexpected packet forwarding behavior. (235075)

## 7368 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)
- 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)
- 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)

## 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 unlearned. 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)

- 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)

## 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)
- Cut-through mode is not supported on SFP+ interfaces. (252731)
- Directflow configuration with "set vlan", but no "action set output interface" configured is not supported. (345852)
- NAT is not supported on subinterfaces (353236)  
SKUs Affected: 7300X3-32C-LC, 7300X3-48YC4-LC, DCS-7050CX3-32S and DCS-7050SX3-48YC12
- NAT is not supported on ECMP interfaces (353704)  
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)
- If the packet ingresses on one line card and egresses on a different line card, then a truncated egress mirrored packets will be 4 bytes shorter than

expected. (445200)

SKUs Affected: DCS-7304 and DCS-7308

- When IPV6 VXLAN underlay core interface is an access port on a SVI, VXLAN encapsulated packets may be sent out with VLAN tag in the outer ethernet header. The work around is to configure the port as a trunk port. (533929)
- "Dont fragment" bit is not set in the outer IP header of the VXLAN encapsulated packets when the payload is non-IP. (605003)
- When Postcard Telemetry is enabled, EVPN Multicast will not be supported. (614072)

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

- DirectFlow actions on interfaces configured as dot1q-tunnel are not supported. (686339)

## 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.