

Arista Networks EOS 4.23.2F Release Notes

Version: 1.1

05 March 2020

Table of Contents

EOS 4.23.2F Release Highlights

New Platforms and Hardware

SNMP MIBs

SNMP Traps

Supported Hardware

Reference to MLAG ISSU Upgrade/Downgrade Table

Enhancements in 4.23.2F

Layer 3

Multi-agent

Rib

Layer 2

Layer 1

MLAG

CVX

MPLS

General

Multicast

Chassis Series

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

DCS-7280R and DCS-7500R Series

DCS-7280R3 and DCS-7500R3 Series

DCS-7050X[2], DCS-7060X[2], DCS-7250X and DCS-7300 Series

DCS-7050X2 and DCS-7300 Series

CCS-720XP Series

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

7368 Series

CCS-720XP Series

Transceiver Series

vEOS Router Series

Resolved Caveats in 4.23.2F

Layer 3

Multi-agent

Rib

Layer 2

Layer 1

MLAG

Accelerated System Update

CVX

MPLS

General

Multicast

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

DCS-7280R2 and DCS-7500R2 Series

DCS-7280R and DCS-7500R Series

DCS-7020 Series

DCS-7280R3 and DCS-7500R3 Series

DCS-7160 Series

DCS-7170 Series

DCS-7050X[2], DCS-7060X[2], DCS-7250X and DCS-7300 Series

DCS-7010T Series

DCS-7060X and DCS-7300 Series

CCS-720XP Series

DCS-7260X3 Series

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

7368 Series

CCS-720XP Series

DCS-7050X and DCS-7300 Series

DCS-7300 and DCS-7500 Series

Transceiver Series

Known Software Caveats in 4.23.2F

Layer 3

Multi-agent

Rib

Layer 2

Layer 1

MLAG

Accelerated System Update

CVX

MPLS

General

Multicast

DCS-7150 Series

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

DCS-7280R and DCS-7500R Series

DCS-7020 Series

DCS-7020RA, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards Series

DCS-7280R3 and DCS-7500R3 Series

DCS-7280E and DCS-7500E Series

DCS-7160 Series

DCS-7170 Series

DCS-7170 Series

DCS-7050X[2], DCS-7060X[2], DCS-7250X and DCS-7300 Series

DCS-7050X2 and DCS-7300 Series

DCS-7010T Series

DCS-7060X and DCS-7300 Series

DCS-7050 and DCS-7300 Series

DCS-7260X3 Series

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

7368 Series

CCS-720XP Series

DCS-7050X and DCS-7300 Series

DCS-7300 and DCS-7500 Series

DCS-7020RA, DCS-7160, DCS-7170, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards Series

Containerized EOS Series

Transceiver Series

vEOS Router Series

BGP EVPN L2/L3 VXLAN/MPLS

Limitations and Restrictions in 4.23.2F

Layer 3

Multi-agent

Rib

Layer 2

Layer 1

MLAG

Accelerated System Update

CVX

MPLS

General

Multicast

DCS-7150 Series

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

DCS-7280R and DCS-7500R Series

DCS-7020 Series

DCS-7280R3 and DCS-7500R3 Series

DCS-7160 Series

DCS-7170 Series

DCS-7050X[2], DCS-7060X[2], DCS-7250X and DCS-7300 Series

DCS-7050X2 and DCS-7300 Series

DCS-7010T Series

DCS-7060X and DCS-7300 Series

DCS-7260X3 Series

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

7368 Series

CCS-720XP Series

DCS-7050X and DCS-7300 Series

DCS-7300 and DCS-7500 Series

DCS-7020RA, DCS-7160, DCS-7170, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards Series

Transceiver Series

vEOS Router Series

BGP EVPN L2/L3 VXLAN/MPLS

Conditions and Impacts

EOS 4.23.2F Release Highlights

New Platforms and Hardware

- 7300-SUP2-D
- 7800R3-36P-LC
- 7800R3-48CQ-LC
- 7800R3-48CQM-LC
- 7800R3K-48CQ-LC
- 7808R3-FM
- DCS-7280CR3-96-F
- DCS-7280CR3K-96-F
- DCS-7800-SUP
- DCS-7808-CH
- FAN-7012H-RED
- OSFP-400G-DR4
- OSFP-400G-XDR4
- PWR-1011-AC-BLUE
- PWR-1011-DC-BLUE
- PWR-1011-DC-RED
- PWR-2411-AC-RED
- PWR-2411-DC-RED
- PWR-721-DC-RED
- PWR-D1-3041-AC-BLUE
- QDD-400G-DR4
- QDD-400G-XDR4
- SFP-10G-T

SNMP MIBs

- SNMPv2, SNMPv3
- RFC 3635 EtherLike-MIB
 - obsoletes RFCs 1650, 2358, 2665
- RFC 3418 SNMPv2-MIB
 - obsoletes RFCs 1450, 1907
- RFC 2863 IF-MIB
 - obsoletes RFCs 1229, 1573, 2233
- RFC 2864 IF-INVERTED-STACK-MIB
- RFC 4292 IP-FORWARD-MIB
 - obsoletes RFC 1354
- ARISTA-SW-IP-FORWARD-MIB
 - IPv4 only
- RFC 4363 Q-BRIDGE-MIB
- RFC 4188 BRIDGE-MIB
- ARISTA-BRIDGE-EXT-MIB

- RFC 4113 UDP-MIB
 - obsoletes RFC 1213
- RFC 4022 TCP-MIB
 - obsoletes RFC 1213
- RFC 4293 IP-MIB
 - obsoletes RFC 1213
- HOST-RESOURCES-MIB
- LLDP-MIB
- LLDP-EXT-DOT1-MIB
- LLDP-EXT-DOT3-MIB
- ENTITY-MIB
- ENTITY-SENSOR-MIB
- ENTITY-STATE-MIB
- RMON-MIB
 - rmonEtherStatsGroup
- RMON2-MIB
 - rmon1EthernetEnhancementGroup
- HC-RMON-MIB
 - etherStatsHighCapacityGroup
- RFC 3636 MAU-MIB
 - ifMauDefaultType and ifMauAutoNegStatus are writeable
- SNMP-TLS-TM-MIB
 - RFC 6353
- SNMP-TSM-MIB
 - RFC 5591
- ARISTA-ACL-MIB
- ARISTA-SNMP-TRANSPORTS-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-SMI-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-PRODUCTS-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-QUEUE-MIB
 - Excluding 7150
- RFC 4273 BGP4-MIB
- RFC 4750 OSPF-MIB
- ARISTA-CONFIG-COPY-MIB
- ARISTA-CONFIG-MAN-MIB
- ARISTA-REDUNDANCY-MIB
 - 7500R, 7500E, 7300X, 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- MSDP-MIB
- PIM-MIB
- IGMP-MIB
- IPMROUTE-STD-MIB
- VRRPV2-MIB
- ARISTA-QOS-MIB
- ARISTA-ENTITY-SENSOR-MIB

- ARISTA-BGP4V2-MIB
- ARISTA-VRF-MIB
- ARISTA-DAEMON-MIB
- ARISTA-IF-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-MAU-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-PFC-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-FIB-STATS-MIB
- ARISTA-GENERAL-MIB
- ARISTA-HARDWARE-UTILIZATION-MIB
- ARISTA-NEXTHOP-GROUP-MIB
- ARISTA-SW-IP-FORWARDING-MIB
- ARISTA-TEST-MIB
- ARISTA-XCVR-DWDM-MIB
- ARISTA-XGS-MIB
- IEEE8021-PFC-MIB
- IEEE8023-LAG-MIB
- MPLS-LDP-STD-MIB
- NOTIFICATION-LOG-MIB
- OSPF-TRAP-MIB
- OSPFV3-MIB
- PTPBASE-MIB

SNMP Traps

- RFC 2863 IF-MIB
 - linkUp, linkDown
- LLDP-MIB
 - lldpRemTablesChange
- RFC 3418 SNMPv2-MIB
 - coldStart
- NET-SNMP-AGENT-MIB
 - nsNotifyRestart
- ENTITY-MIB
 - entConfigChange
- ENTITY-STATE-MIB
 - entStateOperEnabled, entStateOperDisabled
- OSPF-MIB
 - ospfNbrStateChange, ospfIfConfigError, ospfIfAuthFailure, ospfIfStateChange
- BGP4-MIB
 - bgpEstablished, bgpBackwardTransition
- ARISTA-REDUNDANCY-MIB
 - aristaRedundancySwitchOverNotif only for: 7500R, 7500E, 7300X, 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R

- ARISTA-CONFIG-MAN-MIB
 - aristaConfigManEvent
- SNMPv2-MIB
 - authenticationFailure
- VRRPv2-MIB
 - vrrpTrapNewMaster
- MPLS-LDP-STD-MIB
 - mplsLdpSessionUp, mplsLdpSessionDown
- MSDP-MIB
 - msdpEstablished, msdpBackwardTransition
- PIM-MIB
 - pimNeighborLoss
- ARISTA-BRIDGE-EXT-MIB
 - aristaMacMove, aristaMacLearn, aristaMacAge

Supported Hardware

Fixed System

- | | | |
|-------------------------|--------------------------|-----------------------|
| • DCS-7150S-24-CL-F | • DCS-7280QR-C36-F | • DCS-7050SX-64-F |
| • DCS-7150S-24-CL-R | • DCS-7280QR-C36-M-F | • DCS-7050SX-64-R |
| • DCS-7150S-24-CL-SSD-F | • DCS-7280QR-C36-M-R | • DCS-7050SX-64-SSD-F |
| • DCS-7150S-24-F | • DCS-7280QR-C36-R | • DCS-7050SX-64-SSD-R |
| • DCS-7150S-24-R | • DCS-7280SR-48C6-F | • DCS-7050TX-128-F |
| • DCS-7150S-24-CL-SSD-R | • DCS-7280SR-48C6-M-F | • DCS-7050TX-128-R |
| • DCS-7150S-52-CL-F | • DCS-7280SR-48C6-M-R | • DCS-7050TX-48-F |
| • DCS-7150S-52-CL-R | • DCS-7280SR-48C6-R | • DCS-7050TX-48-R |
| • DCS-7150S-52-CL-SSD-F | • DCS-7280TR-48C6-R | • DCS-7050TX-48-SSD-F |
| • DCS-7150S-52-CL-SSD-R | • DCS-7280TR-48C6-F | • DCS-7050TX-48-SSD-R |
| • DCS-7150S-64-CL-F | • DCS-7280TR-48C6-M-F | • DCS-7050TX-64-F |
| • DCS-7150S-64-CL-R | • DCS-7280TR-48C6-M-R | • DCS-7050TX-64-R |
| • DCS-7150S-64-CL-SSD-F | • DCS-7280QR-C72-F | • DCS-7050TX-64-SSD-F |
| • DCS-7150S-64-CL-SSD-R | • DCS-7280QR-C72-M-F | • DCS-7050TX-64-SSD-R |
| • DCS-7150SC-24-CLD-F | • DCS-7280QR-C72-M-R | • DCS-7050TX-72-SSD-F |
| • DCS-7150SC-24-CLD-R | • DCS-7280QR-C72-R | • DCS-7050TX-72-SSD-R |
| • DCS-7150SC-64-CLD-F | • DCS-7280SR2-48YC6-F | • DCS-7050TX-96-SSD-F |
| • DCS-7150SC-64-CLD-R | • DCS-7280SR2-48YC6-M-F | • DCS-7050TX-96-SSD-R |
| • DCS-7170-64C-F | • DCS-7280SR2-48YC6-M-R | • DCS-7010T-48-F |
| • DCS-7170-64C-R | • DCS-7280SR2-48YC6-R | • DCS-7010T-48-R |
| • DCS-7170-32C-F | • DCS-7020TR-48-F | • DCS-7050SX-72-F |
| • DCS-7170-32C-M-F | • DCS-7020TR-48-R | • DCS-7050SX-72-R |
| • DCS-7170-32C-M-R | • DCS-7280QRA-C36S-F | • DCS-7050SX-72-SSD-F |
| • DCS-7170-32C-R | • DCS-7280QRA-C36S-R | • DCS-7050SX-72-SSD-R |
| • DCS-7170-64C-M-F | • DCS-7280SR2A-48YC6-F | • DCS-7050SX-96-F |
| • DCS-7170-64C-M-R | • DCS-7280SR2A-48YC6-M-F | • DCS-7050SX-96-R |
| • DCS-7170-32CD-F | • DCS-7280SR2A-48YC6-M-R | • DCS-7050SX-96-SSD-F |
| • DCS-7170-32CD-R | • DCS-7280SR2A-48YC6-R | • DCS-7050SX-96-SSD-R |

- PWR-1100AC-F
- PWR-1100AC-R
- PWR-750AC-F
- PWR-750AC-R
- PWR-500-DC-F
- PWR-500-DC-R
- PWR-500AC-F
- PWR-500AC-R
- PWR-745AC-F
- PWR-745AC-R
- FAN-7010
- PWR-1900-DC-F
- PWR-1900-DC-R
- PWR-1900AC-F
- PWR-1900AC-R
- PWR-747AC-F
- PWR-747AC-R
- FAN-7002H-R
- FAN-7000H-R
- PWR-1600AC-F
- FAN-7001D-F
- FAN-7001C-F
- PWR-400AC-BLUE
- FAN-7011M-F
- FAN-7011M-R
- PWR-1021-AC-RED
- PWR-511-AC-BLUE
- PWR-511-AC-RED
- PWR-621-AC-RED
- PWR-400AC-RED
- FAN-7001DH-F
- PWR-1611-AC-RED
- PWR-1611-DC-RED
- 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-1011-AC-RED
- PWR-721-DC-RED
- PWR-621-AC-BLUE
- FAN-7012H-RED
- PWR-1011-AC-BLUE
- PWR-1011-DC-BLUE
- DCS-7280SRA-48C6-F
- DCS-7280SRA-48C6-M-F
- DCS-7280SRA-48C6-M-R
- DCS-7280SRA-48C6-R
- DCS-7280TRA-48C6-F
- DCS-7280TRA-48C6-M-F
- DCS-7280TRA-48C6-M-R
- DCS-7280TRA-48C6-R
- DCS-7020TRA-48-F
- DCS-7020TRA-48-R
- DCS-7280CR2A-60-F
- DCS-7280CR2K-60-F
- DCS-7280SRAM-48C6-F
- DCS-7280SRAM-48C6-R
- DCS-7280CR2-60-F
- DCS-7280CR2A-30-F
- DCS-7280CR2K-30-F
- DCS-7280QRA-C36S-M-F
- DCS-7280QRA-C36S-M-R
- DCS-7280SR2K-48C6-M-F
- DCS-7280SR2K-48C6-M-R
- DCS-7020SR-24C2-F
- DCS-7020SR-24C2-R
- DCS-7280CR2M-30-F
- DCS-7280SRM-40CX2-F
- DCS-7280SRM-40CX2-R
- DCS-7020SR-32C2-F
- DCS-7020SR-32C2-R
- DCS-7020SRG-24C2-F
- DCS-7020SRG-24C2-R
- DCS-7280CR3-32P4-F
- DCS-7280CR3-32P4-M-F
- DCS-7280CR3-32P4-M-R
- DCS-7280CR3-32P4-R
- DCS-7280CR3K-32P4-F
- DCS-7280CR3-32D4-F
- DCS-7280CR3-32D4-M-F
- DCS-7280CR3-32D4-M-R
- DCS-7280CR3-32D4-R
- DCS-7280CR3K-32D4-F
- DCS-7280PR3-24-F
- DCS-7280PR3-24-M-F
- DCS-7280PR3K-24-F
- DCS-7280DR3-24-F
- DCS-7280DR3-24-M-F
- DCS-7280CR3-96-F
- DCS-7280CR3K-96-F
- DCS-7050TX-128-SSD-F
- DCS-7050TX-128-SSD-R
- DCS-7050TX-72-F
- DCS-7050TX-72-R
- DCS-7050TX-96-F
- DCS-7050TX-96-R
- DCS-7010T-48-DC-F
- DCS-7010T-48-DC-R
- DCS-7260QX-64-F
- DCS-7260QX-64-R
- DCS-7260QX-64-SSD-F
- DCS-7260QX-64-SSD-R
- DCS-7050SX-72Q-F
- DCS-7050SX-72Q-R
- DCS-7050TX-72Q-F
- DCS-7050TX-72Q-R
- DCS-7050QX2-32S-F
- DCS-7050QX2-32S-R
- DCS-7050SX2-128-F
- DCS-7050SX2-128-R
- DCS-7050TX2-128-F
- DCS-7050TX2-128-R
- DCS-7050SX2-72Q-F
- DCS-7050SX2-72Q-R
- DCS-7060CX-32S-F
- DCS-7060CX-32S-R
- DCS-7060CX2-32S-F
- DCS-7060CX2-32S-R
- DCS-7260CX-64-F
- DCS-7260CX-64-R
- DCS-7260CX-64-SSD-F
- DCS-7260CX-64-SSD-R
- DCS-7060SX2-48YC6-F
- DCS-7060SX2-48YC6-R
- DCS-7260CX3-64-F
- DCS-7260CX3-64-R
- DCS-7050CX3-32S-F
- DCS-7050CX3-32S-R
- DCS-7050SX3-48YC12-F
- CCS-720XP-24ZY4-F
- CCS-720XP-48ZC2-F
- DCS-7050SX3-48YC8-F
- DCS-7050SX3-48YC8-R
- DCS-7260CX3-64E-F
- DCS-7260CX3-64E-R
- DCS-7060DX4-32-F
- DCS-7060PX4-32-F

- PWR-1011-DC-RED
- PWR-2411-AC-RED
- PWR-2411-DC-RED
- FAN-7000-F
- FAN-7000-R
- PWR-460AC-F
- PWR-460AC-R
- PWR-460DC-F
- PWR-460DC-R
- DCS-7280SE-64-F
- DCS-7280SE-64-R
- DCS-7280SE-72-F
- DCS-7280SE-72-R
- DCS-7280SE-68-F
- DCS-7280SE-68-R
- DCS-7280CR-48-F
- DCS-7050SX-128-F
- DCS-7050SX-128-R
- DCS-7050SX-128-SSD-F
- DCS-7050SX-128-SSD-R
- DCS-7250QX-64-F
- DCS-7250QX-64-R
- DCS-7250QX-64-SSD-F
- DCS-7250QX-64-SSD-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
- CCS-720XP-24Y6-F
- CCS-720XP-48Y6-F
- DCS-7050CX3M-32S-F
- DCS-7050CX3M-32S-R
- CCS-720XP-24Y6-R
- CCS-720XP-48Y6-R
- DCS-7160-32CQ-F
- DCS-7160-32CQ-R
- DCS-7160-48YC6-F
- DCS-7160-48YC6-R
- DCS-7160-48TC6-F
- DCS-7160-48TC6-R
- FAN-7002-R
- FAN-7002H-F
- FAN-7000H-F
- FAN-7002-F

Modular

- FAN-7002-R
- FAN-7002H-F
- FAN-7000H-F
- PWR-3K-AC-F
- PWR-3K-AC-R
- PWR-2700-DC-F
- PWR-2700-DC-R
- PWR-3K-DC-BLUE
- PWR-3K-DC-RED
- PWR-3KT-AC-BLUE
- PWR-3KT-AC-RED
- PWR-D1-3041-AC-BLUE
- 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-7508N
- 7500E-6CFPX-LC
- DCS-7500-SUP2
- 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
- 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
- 7500R3-24P-LC
- 7800R3-36P-LC
- 7800R3-48CQ-LC
- 7800R3-48CQM-LC
- 7800R3K-48CQ-LC
- DCS-7800-SUP
- DCS-7808-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
- 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
- 7300-SUP2-D

- DCS-7500-SUP2-D
- 7500R-36CQ-LC

- 7808R3-FM

- FAN-7002-F

Transceiver

- | | | |
|-----------------------|----------------------|------------------|
| • 1000BASE-BX10-D | • 25GBASE-AOC | • CAB-D-2Q-400G |
| • 1000BASE-BX10-U | • 25GBASE-CR | • CAB-D-4Q-200G |
| • 40GBASE-AOC | • 25GBASE-LR | • CAB-D-4Q-400G |
| • 10GBASE-LRL | • 25GBASE-SR | • CAB-D-8S-200G |
| • 10GBASE-ZR | • 100GBASE-SRBD | • CAB-D-D-400G |
| • 40GBASE-PLRL4 | • 100GBASE-SWDM4 | • CAB-O-2Q-200G |
| • 40GBASE-SR4 | • 100GBASE-ERL4 | • CAB-O-2Q-400G |
| • 40GBASE-XSR4 | • 100GE-DWDM2 | • CAB-O-4Q-200G |
| • 40GBASE-LR4 | • CFP2-100G-DWDM-DCO | • CAB-O-4Q-400G |
| • 40GBASE-LRL4 | • CFPX-100G-DWDM | • CAB-O-8S-200G |
| • 40GBASE-PLR4 | • CFPX-200G-DWDM | • CAB-O-O-400G |
| • CFP2-100GBASE-LR4 | • 1000BASE-LX | • OSFP-400G-2FR4 |
| • 10GBASE-AOC | • 1000BASE-SX | • OSFP-400G-FR4 |
| • 40GBASE-UNIV | • 1000BASE-T | • OSFP-400G-LR8 |
| • 100GBASE-CR4 | • 10GBASE-CR | • OSFP-400G-SR8 |
| • 100GBASE-LR4 | • 10GBASE-LR | • QDD-400G-FR4 |
| • 100GBASE-LRL4 | • 10GBASE-SR | • QDD-400G-LR8 |
| • 40GBASE-ER4 | • 10GBASE-SRL | • QDD-400G-SR8 |
| • 10GBASE-DWDM | • 100GBASE-XSR4 | • 100GBASE-FR |
| • 10GBASE-DWDM-T | • 100GBASE-DR | • CAB-Q-2Q-100G |
| • 10GBASE-DWDM-ZR | • 10GBASE-ERBD-D | • OSFP-400G-DR4 |
| • CFP2-100GBASE-ER4 | • 10GBASE-ERBD-U | • OSFP-400G-XDR4 |
| • 100GBASE-PSM4 | • 10GBASE-ERLBD-D | • QDD-400G-DR4 |
| • 100GBASE-SR4 | • 10GBASE-ERLBD-U | • QDD-400G-XDR4 |
| • 40GBASE-SRBD | • ADPT-O-Q-100G | • 100GBASE-LR |
| • 100GBASE-AOC | • AOC-D-D-400G | • SFP-10G-T |
| • CFP2-100GBASE-XSR10 | • AOC-O-O-400G | • 10GBASE-ER |
| • 100GBASE-CWDM4 | • CAB-D-2Q-200G | • 40GBASE-CR4 |

Reference to MLAG ISSU Upgrade/Downgrade Table

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

Enhancements in 4.23.2F

Layer 3

- Support for MTU signaling in RSVP on a Label Switching Router (365306)
- Added support for soft preemption on an RSVP configured transit router. When bandwidth on an interface is oversubscribed, LSPs signaled with soft preemption requested will be torn down after a delay to enable the headend to reroute the LSP in a non disruptive manner. The delay is configured using a soft preemption timer. (365319)
- Support for configurable maximum IS-IS LSP size. (368270)
- Support for interface-level user-defined circuit-id value for DHCP snooping v4 (388105)
- Support for sending interface name in remote-id option of DHCP relay v6 packet in clear text format. (398472)

Multi-agent

- Support for 'match interface' clause in route-map for BGP redistribute and network commands (141360)
- Support for setting route tag in neighbor route-map on the inbound side in BGP. (141362)
- Support for redistributing OSPF aggregates in multi-agent mode. (232224)
- Add RIP support in multi-agent model. (290423)
- Support for BGP "neighbor route-to-peer" (339077)
- Support ECMP paths in multi-agent mode in the sFlow BGP extension (352125)
- Support for redistributing connected, static, ISIS, OSPF, and BGP routes into BGP in the target VRF. (353021)
- Support for advertising IPv4 NLRI with an IPv6 Next Hop (RFC5549) over IPv4 IBGP sessions. Advertising over IPv4 EBGP sessions is not supported. (356674)
- Support for dynamically discovering IPv6 link-local BGP peers (357022)
- Support for per address-family bgp missing policy using "bgp missing-policy address-family all" configuration. (359885)
- Support for Evpn Vxlan All-Active Multi-homing Integrated Routing and Bridging (360145)

- MPLS EVPN interworking with L3VPN(RFC4364) (367813)
- Add a CLI to debug route map evaluation against a BGP path redistributed into BGP. (421313)

Rib

- ISIS-SR Graceful Restart support for Rib Agent Restart. (364523)

Layer 2

- Support for static lease command, "lease <V4ADDR> mac <MACADDR>" (420545)

Layer 1

- Support for "show interfaces phy diag error-correction histogram". This command displays the number of symbols corrected by reed-solomon coding on QSFP100 ports with the CRT50216 PHY at 100G-4 and 50G-2 speeds. (388815)
 SKUs Affected: 7500R3-36CQ-LC, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R and DCS-7280CR3-32P4-R
- Support for 50G-2 on QSFP28 ports. (420335)
 SKUs Affected: 7500R3-36CQ-LC, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R and DCS-7280CR3-32P4-R

MLAG

- Enhance MLAG Dual Primary detection feature to force dual primary state when all peer link transceivers/cables are pulled. (245410)

CVX

- Support for the Policy Control Service (PCS). (314230)
- Add support for EVPN VXLAN control plane provisioning from the Arista OpenStack ML2 driver. (425529)

MPLS

- MPLS forwarding over non-hierarchical MPLS routes is hitless when restarting the SandL3Unicast agent (412016)

General

- Support the redistribute of a device's system MAC over EVPN. (274143)
- Add support to configure PTP clock accuracy. (347080)
- This is an enhancement which gives capability for PBR to match on TTL (353504)

- Support for user set thresholds which produce a syslog message when crossed, indicating no PTP message has been received. The threshold can be set in two ways; a timer in seconds since the last received message, or by the number of missed messages.

Sync, follow up and announce messages supported for timer in seconds.

Sync, follow up, delay resp and announce messages supported for number of missed messages. (377097)

- Support for DHCP TFTP option 150. (414987)
- Support Python3 bindings for EosSdk APIs (432730)

Multicast

- Fast drop routes are installed dynamically when using static multicast. (407663)

Chassis Series

- Added the "reset system storage secure" command to securely erase storage devices (255978)

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

- Supports the ability to control the value of the CoS bits for CPU generated packets with VLAN tags on some interfaces on a device (223371)
- Support for configuring flow-spec in a non-default VRF. (290829)
- MACsec proxy support for L2 VXLAN (358157)
- Support header truncation in monitor sessions (359472)
- Support per port TC to COS mapping on switch ports, L2 sub-interfaces, and L2 sub-interfaces under port-channels (391489)
- Support prediction of forwarding destinations for certain packets specified as byte streams (398944)
- Support for match on inner VLAN (along with VLAN and CoS) in policy-maps on Sand platform. (402627)
- Add CLI to disable forced remote fault generation. (414345)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Support for DSCP mask in IPv4 ACL (422605)
- Support for SFP-10G-T on select SKUs. (425886)
SKUs Affected: 7500R2AK-48YCQ-LC, 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 and DCS-7280SR2A-48YC6-R

DCS-7280R and DCS-7500R Series

- Support for timestamping keyframes in Tap Aggregation mode (259705)
- Support for TCP MSS Clamping in hardware (356764)

DCS-7280R3 and DCS-7500R3 Series

- Support for auto-negotiation on front panel ports except QSFP100 ports using the CRT50216 PHY. (248088)
- Adding support for Policing via Qos policies. (292700)
- Support for User-Defined Fields in Access Control Lists (362707)
- Support for MPLS label lookup based on two MPLS labels. (373084)
- Support for the CLI command "<no|default> transceiver qsfp default-mode 4x10G". (389203)
- MPLS fast failover support (400197)
- Vxlan VARP support (403700)
- Drop action not available for PBR policies (404416)
- Vxlan single-homed EVPN support (416003)
- Support for configurable GRE protocol ID and metadata insertion for GRE tunnels used as monitor session destinations (417266)
- Policy Based Routing counters are configured by the CLI command "hardware counter feature pbr in". The legacy command "platform fap tcam counter feature" is no longer operative. (418573)
- Vxlan MLAG support (425980)
- Support LANS polling mode. (429040)

DCS-7050X[2], DCS-7060X[2], DCS-7250X and DCS-7300 Series

- Added a CLI controlling the inclusion of the System MAC of the switch in the LAG hash seed effective value. Effective value of LAG hash seed is composed of configured LAG hash seed and now optionally System MAC. (185236)
- IPv6 flow-label based load-balancing is now available as a configuration option for port channel load balancing and is enabled by default. (380507)

- Support 'Nexthop ECMP set' and 'Nexthop ECMP entries' streaming via Terminattr (387258)
- SNMP support for monitoring "Nexthop ECMP set" and "Nexthop ECMP entries". (387259)
- Support for SFP-10G-T. (412663)
 SKUs Affected: 7300X3-48YC4-LC, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7060SX2-48YC6-F and DCS-7060SX2-48YC6-R

DCS-7050X2 and DCS-7300 Series

- Support for Multi VTEP MLAG. (440978)

CCS-720XP Series

- Added the ability to configure "poe reboot action" per port (391302)
- Added support for PWR-721-DC-RED (412084)
- Added the "poe link down action" command to configure the behavior of PoE when a link goes down (413972)

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

- Mirror On Drop debug feature, which allows monitoring of unexpected TCP or UDP flow drops in the ingress pipeline. Upon detecting such dropped flow it generates reports, which are sent to pre-configured external collectors as UDP packets at configurable intervals. Reports contain a timestamp of when the drop was started/stopped, a drop reason vector, and up to 80 first bytes of the dropped packet. (343737)

7368 Series

- Support for subinterfaces. (390374)

CCS-720XP Series

- Support for variable polling interval. (271751)
- When the switch reboots, power to the PoE ports will be maintained by default. The behavior can be controlled with the 'reboot action' command in the 'poe' configuration mode. (359049)
 Series Affected: CCS-720XP Series

Transceiver Series

- EOS support for adding data path state, data path firmware fault, module state and module firmware fault for OSFP and QSFP-DD transceivers to the output of 'show transceiver status interface' CLI command. (366704)

vEOS Router Series

- Add support for Cloud HA in vEOS on GCP (362334)

Resolved Caveats in 4.23.2F

Layer 3

- If a sub-interface creation/deletion overlaps with it's VRF creation/deletion then the sub-interface may not be deleted properly in kernel. This prevents creation of same sub-interface in kernel again.

Workaround is to restart Ebra agent. (337707)

- When deleting and re-adding a VRF in quick succession either manually or via config-replace, BGP routes may remain unresolved.

The workaround is to delete and recreate the VRF as a first attempt, and restart the Bgp Agent as a last resort. (377648)

- At scale, when migrating BFD sessions from SwBfd to HwBfd, sessions may flap and take a few minutes to converge. (400256)
- When "flow tracking hardware" is enabled, some IPFIX data packets exported from HW may have an incorrect UDP checksum and hence could get dropped. There is no workaround to this issue. (416139)

Series Affected: CCS-720XP Series

- DHCP Relay Agent may exit unexpectedly when 'ipv6 dhcp relay install routes' is unconfigured with very high dhcp prefix delegation routes scale. (417272)
- The ISIS agent may restart if RSVP-TE is for Fast Reroute is configured at the same time as ISIS overload bit or te router-id. (417606)
- When OSPFv3 graceful restart is configured on an ABR and NSSA Type 7 external LSAs are translated to ASE Type 5 LSAs and SSU or SSO is executed, OSPFv3 may withdraw the translated ASE type 5 LSAs temporarily, causing traffic loss in OSPF peers. (421638)
- IPv6 nexthop address is not accepted by the Cli under nexthop-group configuration. (426019)
- In multi-agent routing protocol mode, if the BGP Prefix Independent Convergence (PIC) feature is enabled or disabled after BGP sessions are established, BGP next hops can mistakenly become unresolvable.

Restarting the Bgp agent can provide a workaround for the issue. (427029)

- 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)
- Deleting a VRF may lead to BFD sessions being stuck in their current state. Only BFD timeouts would result in a change of state of BFD sessions. A workaround is to restart the Bfd agent. (427612)
- Changing the VRF of an interface where BFD static neighbors are configured will result in BFD session failure. Subsequent configuration changes via 'bfd static neighbor' will result in duplicate BFD configuration displayed in the output of 'show active'. The workaround is to unconfigure BFD static neighbors before moving an interface to a different VRF. (430640)
- Changing the VRF of an interface where BFD static neighbors are configured will result in BFD session failure. Subsequent configuration changes via 'bfd static neighbor' will result in duplicate BFD configuration displayed in the output of 'show active'. The workaround is to unconfigure BFD static neighbors before moving an interface to a different VRF. (430672)
- When 'bfd per-link rfc-7130' is configured on static port-channel, BFD member session down may not block LAG member from joining port-channel. Configure 'bfd per-link rfc-7130' on a static port-channel may cause port-channel interface flap. (431738)
- BFD sessions may never come up after doing changing global BFD configuration (for example, doing a "configure replace" or "bfd shutdown" and "no bfd shutdown").

To workaround, restart Bfd agent. (432109)

- In BFD per-link rfc-7130 interop mode, when 'bfd neighbor' address does not match peer port-channel IP, unconfigure 'bfd per-link' on one side of BFD peers will cause LAG member to be removed from port-channel on the other peer. Make BFD configuration match on both sides to recover the port-channel. (434055)
- If an interface flaps, host routes may not be installed for static ARP/neighbor entries. This may also affect dynamic entries that were present before the flap and immediately relearned after the flap.

A workaround is to restart the Arp agent. Another workaround is to flap the interface again until the problem is resolved. (435806)

- Turning on the event monitor while using RSVP can cause ConfigAgent to crash. Shutting it down with `no event-monitor` will prevent the issue. (436254)
- If an interface flaps, host routes may not be installed in the hardware for static ARP/neighbor entries. This may also affect dynamic entries that were

present before the flap and immediately relearned after the flap.

A workaround is to flap the interface again until the problem is resolved. Another workaround is to restart the Layer3 Forwarding agent. (436597)

- V6 Attached route for a host over the Vxlan interface may not be programmed in the hardware for a remotely learned host. (438882)
- BGP agent may restart when auto-local-adr is configured on a non directly connected BGP session (442171)
- When IS-IS dynamic flooding is enabled, the IS-IS process may restart when there's a network topology change. (444542)

Multi-agent

- 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)
- Certain IPv4(6) BGP routes that were intermittently marked as looping due to a loop in the nexthop resolution may not get installed in RIB and FIB, eventhough the BGP path is active in the BGP table. This issue can cause loss of traffic for any incoming packets expecting to be routed based on the missing route. (415865)
- If Ethernet Segment Identifiers are moved between interfaces using configure replace, EVPN type 4 Ethernet Segment Routes may not be generated causing the switch to not participate in Designated Forwarder election for that Ethernet Segment. (420873)
- When OSPF graceful restart is configured in mulitple VRFs in multi-agent mode and routes are redistributed into OSPF and SSU or SSO is executed, OSPF may withrdaw the external LSAs for the redistributed routes temporarily, causing traffic loss in OSPF peers. (421252)
- Bgp agent may unexpectedly restart if import of both remote VPN peer paths and locally redistributed/network paths are configured (423096)
- In an EVPN deployment, a device reload may result in some MAC/IP Advertisement routes with non-zero IP addresses not being advertised. Refreshing the associated MAC address table or ARP table entry will enable advertisement to proceed. (424191)
- "show [ip | ipv6] bgp ... | json" commands which filter routes may, under certain timing conditons whereby a route is printed but a subsequent route filtered, produce malformed JSON and return JSON error. (424967)
- When configuring multiple VLANs at the same time to use BGP EVPN as the control plane, the BGP agent may crash due to race condition. (425573)

- If next-hop-self is explicitly configured for an eBGP peer (using 'neighbor <> next-hop-self' command) and the peer is also activated in EVPN address family, the BGP agent can unexpectedly restart. The explicit next-hop-self configuration for the peer should be removed to avoid this. (431680)
- The sFlow or MSDP agent, when configured with BGP extension, may unexpectedly restart when BGP routes are being removed. (432001)
- When using EVPN MPLS IRB, an ARP request from a host at a multi-homed site that is targeted at another host in the same site, may result in the the non-DF PE of the site sending an unexpected proxy reply towards the requesting host. The proxy reply sent by the non-DF PE may incorrectly modify state in any switches within the site, making them learn the target host MAC on the non-DF PE port, resulting in transient connectivity issues. (435029)
- When using EVPN MPLS IRB, an ARP request from a host at a multi-homed site that is targeted at another host in the same site, may result in the PE of the site sending an unexpected proxy reply towards the requesting host. The proxy reply sent by the PE may incorrectly modify state in any switches within the site, making them learn the target host MAC on the PE port, resulting in transient connectivity issues. (438142)
- If reachability of a path changes while awaiting convergence, the state may not be updated during post-convergence processing (439871)

Rib

- Under extreme circumstances, the area leader may exclude itself from the dynamic flooding topology. (425590)

Layer 2

- If VXLAN routing is configured on a MLAG, the VxlanSwFwd agent can crash when there UDP traffic received on port 51023 current used by VxlanSwFwd to synchronize ARP entries. (364633)
- If a non-existent interface, or an L2 subinterface configured in a VLAN with MLD Snooping enabled, and an IPv6 multicast group is statically configured on that interface, the MldSnooping agent may experience unexpected restarts. (420903)
- RADIUS Accounting Stop messages won't be sent for MBA supplicants in case of interface flapping. (422099)
- Supports IP Locking local interface in non-default VRF. (434731)

Layer 1

- Peer link flap can cause interface to remain down. Workaround is to shut/no shut the interface. (423313)
Series Affected: DCS-7060X4 Series
- Enabling or disabling MacSec on a port while traffic is running can result in the port getting into an unusable state where all traffic is dropped. Evora agent needs to be hitfully reset to recover the port. (436111)
SKUs Affected: 7500R-8CFPX-LC, 7500RM-36CQ-LC, DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R

MLAG

- Vxlan Controller Service flood lists will be lost during ISSU causing an interruption to BUM traffic (430867)
- When MLAG dual primary recovery delay is active, and the config for recovery delay is removed, the interfaces will errdisabled indefinitely. (433059)

Accelerated System Update

- Traffic loss of 1-2 seconds may be experienced on IS-IS Graceful Restart in scale scenarios following SSO or ASU (414174)
- Repeated downgrades and upgrades between < 4.22.1 and >= 4.22.1 may result in exceeding archiver quota and increased boot time. If performing an SSU upgrade it may lead to extended traffic outage. The workaround is to remove all the /mnt/flash/var_archive and /mnt/flash/archive directories, after a few cycles. (433495)

CVX

- 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)
- When a switch or CVX is upgraded or rebooted, CVX services can become partially disruptive (419781)
- In a Multi-VTEP MLAG deployment, rebooting a peer may result in the MACs learned on an orphan port with a VLAN to VNI binding to a local MLAG VTEP IP on the other peer to incorrectly update their move count. (427956)
- Add support for Port-Channel interfaces for ContainerTracer. (429898)

- In a Multi-VTEP MLAG deployment, rebooting a peer may result in MACs being learned on an orphan port with a VLAN to VNI binding to a local MLAG VTEP IP to not be advertised to CVX. (430536)
- The Macro-Segmentation Service does not support connecting a port-channel to the firewall when used with a CheckPoint firewall device with a sub-interface configured. (431709)
- In a MLAG deployment with Multi-VTEP MLAG enabled, a restart of the forwarding may result in the lose of VLAN to VNI bindings learned from the MLAG peer. (432294)
- Upgrading a Palo Alto Networks firewall from version 9.0.4 to 9.0.5 causes the Macro Segmentation System (MSS) on CVX to receive exceptions due to invalid firewall API replies, causing all MSS policies to be removed. (436030)

MPLS

- After rebooting or performing a configure replace, some MPLS routes may not be programmed.
A workaround is to restart the Mpls agent or Layer 3 forwarding agent. (428105)
- Hardware fast failover may converge slower with a larger number of routes failing over on port-channel interfaces. (439078)
Series Affected: DCS-7500E, DCS-7500R and DCS-7500R2 Series

General

- Inserting or removing a linecard can result in the system panicking with a signature similar to

[286.353571] BUG: unable to handle kernel paging request at ffffffff
[286.436777] IP: [<ffffff8130fa93>] pci_free_cap_save_buffers+0x16/0x29

or

[336.206463] BUG: unable to handle kernel NULL pointer dereference at 0000000000000011
[336.300062] IP: [<ffffff8130c63c>] _pci_find_saved_cap+0x18/0x31 (262388)
- The boot config command takes effect immediately inside a config session and cannot be aborted (346171)
- If the L3 forwarding agent is restarted when VXLAN MAC addresses are present and the route to remote VTEPs are down, then the L2 forwarding agent might unexpectedly restart in rare cases. (355110)

- 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)
- 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)
- If SSH or TELNET is enabled in a non-default VRF, restarting Sysdb can cause the service to be unavailable. This can typically happen on a standby supervisor configured with the RPR redundancy protocol, and may only be noticeable after a switchover. To recover, run 'sudo service xinetd restart' from bash. (420882)
- Adding or Removing privilege level to an existing group in a LDAP group policy results in duplicate entries appearing in cli. (422323)
- EventMgr Agent may crash after manually hot swapping a Line Card. This has been observed only on Fully loaded 8-slot modular systems so far. (422697)
- The Macro-Segmentation Service (MSS), when configured to communicate with a Checkpoint firewall device running in layer 3 mode, fails to ensure rule priority ordering is maintained. (424556)
- If "logging trap system severity emergencies" is configured, the SuperServer agent may restart repeatedly. Use "logging trap system severity alerts" to log both alert messages and emergency messages. (425696)
- In a VXLAN routing setup configured with Virtual Ip (Anycast IP) , when the local L3 VTEP (where the host is local) receives the ARP request for the gateway Anycast IP, then it will update the sender's ARP entry and reply to the sender host. However, the remote L3 VTEPs will receive the request due to VXLAN flooding, but will not update the sender's host arp entry, possibly resulting in an inconsistent ARP entry between the local and remote VTEPs. (429227)
- Configuring EOS via OpenConfig may fail when 'redistribute connected' is present in the BGP or ISIS configuration and static routes are not present.

This can be resolved by adding a single IPv4 or IPv6 static route. (434767)

Multicast

- With a scale of few thousand multicast route, and multicast traffic stopped and restarted after a while (> 5mins) some routes might get into a fastdrop state causing loss of multicast traffic. Clearing the affected multicast route can be used as a work around. (380214)

- If a VRF is removed while it has active multicast routes, the Multicast Agent can restart unexpectedly. (418028)
- The first SA msg send from one Msdp Peer to another could be delayed up to 60 seconds. (444449)

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

- The platform Slice agent may restart unexpectedly when sflow hardware acceleration is enabled or disabled. (247895)
- With MPLS pop routes and IP egress ACL configuration, restarting the L3 forwarding agent may lead to another unexpected restart. (349534)
- 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)
- With PFC + Sflow configured on interface, TC0 data traffic will see drops if sflow interface has congestion. (398779)
- 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)
- PTP time will temporarily be incorrect if the SandFap agent hangs for more than 32 seconds. The PTP time will recover on its own after a few minutes. (403095)
- In the event of a SSO, reload, or SandTunnel agent restart, sub-interface, VLAN Translation, LDP MPLS Pseudowires, and other features that require egress VLAN tagging may mis-forward and not tag egress traffic.

The workaround is to re-configure these features to correct hardware programming corruption cause by the SSO, reload, or SandTunnel agent restart event. (405157)

- Advertising BGP Flowspec rules with police action while having insufficient TCAM space may cause the SandAcl agent to restart. (405381)
- Advertising BGP Flowspec rules with police action while having insufficient TCAM space may cause the SandAcl agent to restart. (405663)
- 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)

- If a large number of BGP Flowspec rules (> 5000) with policer actions configured are programmed and a port channel with BGP Flowspec applied to is shut and unshut, the filtering actions may fail to program and the SandAcl agent might restart unexpectedly.
To recover from this condition, reduce the number of policer actions configured to 5000 or fewer. (411635)
- Control plane connectivity on ports in 802.1q tunnel mode will be affected if the port mode is set to 'dot1q-tunnel' after the access VLAN is configured. The workaround is to set switchport mode to dot1q-tunnel first, then set switchport access vlan. Similarly, the access VLAN needs to be cleared before clearing the port mode when unconfiguring. (419766)
- When a BGP Flowspec rule changes both its priority and the action(s) to be applied on matching traffic, the new action(s) may not be honored. (421495)
- EventMgr Agent may crash after a manual linecard hotswap. A temporary workaround is to remove EventMgrSandVoqDeletePlugin via "rm /usr/lib/EventMgrPlugin/EventMgrSandVoqDeletePlugin" and then restart EventMgr via "agent EventMgr terminate". This will remove the capability of VOQ Delete event management. (424250)
- Applying a CPU Traffic Policy with policers while policer resources are exhausted may cause the SandAcl agent to restart unexpectedly. (426395)
- Configuring the maximum TCAM bank limit while having a number of banks allocated for BGP Flowspec larger than the configured limit may cause both the SandAcl and SandTcam agents to restart unexpectedly and continuously. To recover from this, the device must be rebooted. (432441)
- 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)
- TCAM banks will not be released if the max bank CLI command was configured while more banks were allocated than the configured limits. For the banks to be released, unconfigure the max bank CLI and try reapplying and removing the TCAM entries. (434549)
- When some number of TCAM banks are already allocated for a feature, configuring the maximum TCAM bank limit lower than the allocated number may cause the SandAcl agent to restart repeatedly. To recover from this, the device must be rebooted. (434560)
- The Igmp agent will not receive any IGMP packets where source MAC address matches xx:xx:04:xx:xx:xx on an SVI. (436955)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

DCS-7280R2 and DCS-7500R2 Series

- While a system with both 7500R and 7500R2 linecards is booting up, under the rare condition that if routes are programmed while the linecards were booting up, the switch may be misforwarding traffic. Restarting the SandL3Unicast agent recovers from the condition. 'platform sand l3 lpm encoding jericho-plus' to prevent the problem. (433628)

DCS-7280R and DCS-7500R Series

- With MPLS 6 label push configured with a port-channel/port-channel sub-interface as nexthop, one of the label pushed becomes null instead of configured value. Issue is not seen on every time port-channel is flapped. Workaround is to restart the "SandL3Unicast" agent. (433872)

DCS-7020 Series

- Configuring a "tunnel tos" value greater than 15 in the Ipsec tunnel interface configuration will cause tunnels not to establish. (423116)
SKUs Affected: DCS-7020SRG-24C2-F

DCS-7280R3 and DCS-7500R3 Series

- Decap groups incorrectly uses uniform mode for QoS. (417136)
- The traffic class in IPv6 packets egressing the switch after MPLS label termination might be incorrect. (420345)
- IPv6 routing into VXLAN is not supported. (422440)
- A restart of the SandFapNi agent may result in incorrect forwarding of packets that are routed over a VXLAN tunnel. (424625)
- IPv6 over MPLS packets received with IPv4 Explicit NULL (i.e. Label=0) are misforwarded. (425904)
- Enabling TX mirroring on a trunk port will cause a traffic storm on that port.
Break the traffic loop by removing the monitor session. (427586)
- PIM-SM v4 does not work on an mlag with IgmpSnooping enabled. (427859)
- Nexthop groups counters doesn't work for direct IP nexthop entries (without any tunnel encapsulation). (428144)
- BFD sessions may flap unexpectedly with default timeout setting. When that happens, the SandFapNi-FixedSystem agent or the SandFapNi-Linecard<X> agent may show the following warning message:
TX packet DMA has recovered after being stuck for 1 cycles. (TX could be

RX too).

To work around the issue, set the BFD total timeout to be longer than one second, e.g. interval: 500ms, multiplier: 3. (429928)

- MPLS Pseudowire tunnel resolving through nexthop-group may not be programmed correctly in hardware. (430182)
- Even with a MAC ACL to deny IS-IS packets, MAC address of the IS-IS neighbor is learnt. (431856)
- IPv6 egress PACL may not take effect when both IPv4 and IPv6 egress PACLs are applied on the same interface. The workaround is to remove the Egress Ipv6 Acls from the interface and reapply. (433886)
- 40G ports may not come up due to missing serdes tuning (434732)
- On a modular Jericho2 system with more than 5 Jericho2 linecards, configuring MTU on any interface causes SandL3Unicast to crash. (443043)

DCS-7160 Series

- 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)

DCS-7170 Series

- 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)
- 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-7050X[2], DCS-7060X[2], DCS-7250X and DCS-7300 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)
- Changing MMU Profile configuration while the PFC watchdog has detected stuck queues on an interface can result in packets getting sent out from the interface instead of being dropped. (400780)
- When running DirectFlow at scale, removing an IP/Mac ACL may result in (repeated) slice restarts (408779)
- Control plane traffic may not be transmitted on a port-channel during stateful switchover if only one linecard is inserted. (411680)
Series Affected: DCS-7300X Series
- Multicast routed packets with an IIF as a VXLAN VLAN can cause the packet to be looped back through the MLAG port-channel of the MLAG peer. (423278)
Series Affected: DCS-7250X, DCS-7260CX, DCS-7300 and DCS-7320 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)
- On MLAG VRRP setup, if there is a VRRP flap that results in MLAG pair transitioning from backup to master and back, VRRP MAC entries can be programmed incorrectly in hardware. (434866)

DCS-7010T Series

- A single-event upset (SEU) on FP_METER_TABLE may cause the Strata agent to restart. This may disrupt traffic flow temporarily. (389323)

DCS-7060X and DCS-7300 Series

- A single-event upset (SEU) on L3_ECMP_RRLB_CNT can lead to a forwarding agent restart, after which the switch forwards traffic normally. (348722)

CCS-720XP Series

- All fans will spin to 100% when one out of two power supplies is not inserted. (435602)
Series Affected: CCS-720XP Series

DCS-7260X3 Series

- Changing the speeds on an entire core from 100G to 50G using interface ranges can lead to a forwarding agent restart. (405636)

- If a subordinate interface (/3) is errdisabled due to a speed misconfiguration during a hitless upgrade or hitless slice agent restart, then a speed change on the interface will cause a hitful Strata forwarding agent restart and impact traffic on all ports.

The workaround is to change the speed on the master port to a different speed and back before configuring the speed on the subordinate port. (408967)

- After a hitless software upgrade from 4.20.x, interfaces Ethernet6/1-4, Ethernet9/1-4, Ethernet23/1-4, Ethernet28/1-4, Ethernet37/1-4, Ethernet42/1-4, Ethernet56/1-4, Ethernet59/1-4 can be unstable after speed change and could be wedged.

The workaround is to perform a hitful restart of the Strata forwarding agents following hitless software upgrade which will flap all interfaces in the system. (417205)

- Inserting or removing a 40G transceiver while an interface is configured with "speed forced 50gfull" will cause a forwarding agent restart.

Workaround is to change the configured speed before inserting or removing the transceiver. (425865)

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

- 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)
- MPLS IPv4/IPv6 ping is not supported (394778)
- When L3 forwarding agent restarts, some of the host routes may not be programmed into hardware resulting in software forwarding for those IP addresses.
The workaround is to reload the system. (398288)
- When the Strata agent is not running, the CPU will not be able to send or receive packets until Strata starts running again. (413020)
- 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
- Configuring 50g on the master lane and 10g on all subordinate lanes of a QSFP port will cause continuous slice agent restarts.

Workaround is to remove invalid speed configuration. (415691)

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

- 1G speed is not supported on DCS-7050CX3M in this release (416148)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- When Evpn multihoming is configured and when vlan-vni map is removed, multicast resources allocated for that VLAN may not be freed resulting in resource allocation failure. The workaround is to restart the L3 forwarding agent. (417616)
- LAG and ECMP hashing on RRoCe, IPV4/TTL and/or IPV6 hop limit packet fields may fail to function. (424477)
- Performing a reload can result in some host routes being unprogrammed. (427790)
Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX3 and DCS-7050X3 Series
- LAG and ECMP hashing on RRoCe, IPV4/TTL and/or IPV6 hop limit packet fields may fail to function. (435585)
- In case of online module insertion, slice agent could crash due to incorrect state in software and recover. (438368)
Series Affected: 7300X3 Series
- 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)
- When the system is running with scaled L3 ARP or multicast route entries, changing routing table or forwarding table partitions may result in additional restart of StrataL3 forwarding agent. (443267)
- Creating a DirectFlow flow with an MLAG interface included in the output interface action can result in a continuous restart of the StrataL3 agent. (446134)

7368 Series

- The switch card may experience internal power failure. (427251)
Series Affected: 7368 Series
SKUs Affected: 7368X4-SC

CCS-720XP Series

- MacBasedAuth Configuration is not honored after switch is rebooted on Sand based platform. User can enable MacBasedAuth by flapping MacBasedAuth Config on the port after switch restarts. (423084)

- When the configuration has "dot1x authentication failure action traffic allow vlan #", "switchport phone trunk untagged" and "switchport phone vlan VLAN#", the Dot1x agent crashes while handling the failure of a tagged EAPOL authentication attempt, if the VLAN in the tag is the same as the one configured in "switchport phone vlan" (435372)
- Dynamic Authorisation(DM/CoA Messages) requests for MBA users will be discarded if userName format is configured for MBA supplicants. Use default values of userName format to use Dynamic Authorisation for MBA supplicants. (441137)

DCS-7050X and DCS-7300 Series

- In Macro-Segmentation Service (MSS) deployment with a firewall running in transparent virtual wire mode, an intercepted host might not be able to communicate with destinations beyond its own L2 segment. (420119)

DCS-7300 and DCS-7500 Series

- BFD sessions may flap when a standby supervisor is inserted after the active supervisor is already online. (431652)
SKUs Affected: 7300-SUP, DCS-7500-SUP2, DCS-7500E-SUP and DCS-7516-SUP2

Transceiver Series

- When operating a 400GBASE-DR4 in 4 x 100G-2 breakout mode, administratively disabling an interface may cause another interface on the same 400GBASE-DR4 to go link down. (428140)
- When operating a 400GBASE-DR4 in 4 x 100G-2 breakout mode, if one interface goes down, it may cause another interface on the same 400GBASE-DR4 to go link down.
Restore link on all 4 interfaces of the 400GBASE-DR4 to recover the port. This will not prevent any future link down from affecting multiple interfaces on the same 400GBASE-DR4. (428829)

Known Software Caveats in 4.23.2F

Layer 3

- If a loopback interface is only configured with a link-local IPv6 address, OSPFv3 will not advertise the address. (131369)
- 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)
- High rate of ISIS-SR MPLS route churn may cause the Mpls agent to restart. (283149)
- 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)
- The Macro-Segmentation Service (MSS) for Layer 3 firewalls may not correctly program traffic redirection/offload flows if the service is enabled prior to DirectFlow being enabled on all switches.

A workaround is to shutdown the MSS service on CVX, enable DirectFlow on all supported switches and re-enable the MSS service. (366022)

Series Affected: DCS-7050, DCS-7050T, DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260X, DCS-7260X3, DCS-7300 and DCS-7300X Series

- BFD configurations with large numbers of sessions may see BFD session flaps during initial session bringup shortly after boot. (375773)
- With a large number of sessions active, a configuration change may cause some sessions to become stuck in incorrect states. Restart the Bfd agent to work around the problem. (377125)
- The Arp agent restarts when a virtual IP is added and there is a matching stpername ARP entry in the kernel but not in the ARP agents internal data structures. Usually, the Arp agent will recover normally. In rare situations, the Arp agent may repeatedly restart. In this case, adding a static ARP/neighbor entry via the CLI will fix the issue with no impact on the system. Note that adding a static ARP/neighbor entry via the CLI for every virtual IP address will avoid this problem completely. (389761)
- 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. (402802)

- 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)
- 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)
- Despite there being no support for TI-LFA protection of proxy node segments, backup paths computed towards the Segment Routing Mapping Server could be installed for IS-IS SR LFIB route and tunnels that correspond to proxy node segments. (429305)
- 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 carries length of the value field 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, SandAcl agent may crash continuously if more than one flow-spec rule were received that differ only in match operator length field.
For example: Two rules noted below are expected to match 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 adding a new unnumbered interface to a route with an ECMP set of nexthops, some traffic may be briefly lost if it attempts to traverse the newly added interface before ARP has a chance to resolve the nexthop. (434726)
- If routing is disabled while the platform L3 agent is coming up, the platform agent may restart. (439604)
- Upon configuring a large scale of VARP subnets, the Acl agent may start crashing with a SIGQUIT. There currently is no workaround. (442400)
- Hardware accelerated BFD sessions on an physical interface with dot1q encapsulation may fail to come up.
A workaround is to disable BFD hardware acceleration. (444671)
- 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)

- OSPF may incorrectly originate an Indication LSA in an area which has LSAs with the DC-bit cleared. (458064)
- When OSPF is translating Type 7 LSAs to Type 5 LSAs, and it stops being a translator, it may not flush the translated Type 5 LSAs when the Type 7 LSAs are flushed. (458780)
- Removal and re-addition of existing SRLG configuration on an interface in quick succession may result in Rib agent exiting unexpectedly (458856)

Multi-agent

- Performing a stateful swtichover 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 ConfigAgent may restart when there is churn in a large number of route next hops (394148)
- Changing IPv6 next hop of 6PE routes via the outbound route map causes incorrect label advertisement of explicit-v4-null (0) (442283)
- Routes may not be updated in the kernel when its next hop gets resolved via a drop route. (444682)
- Applying an inbound route map to change the next hop of a 6PE route makes the route invalid (445187)
- Mpls agent may consume high memory and can restart unexpectedly if there is a large amount of churn in vpn-ipv4 or vpn-ipv6 nexthops present on the router. (450048)

Rib

- When both IPv4 LU and IPv6 LU are enabled, the BGP tunnel table may enter an inconsistent state that requires a clearing of neighbors. (384365)
- 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 (standed, extended, large communities etc.) are included in the outbound route advertisements. (384629)

- The dynamic flooding algorithm to compute the flooding topology picks a start node to begin the computation. When that node is partitioned from the Area Leader, the nodes in the Area Leader's partition will not be in the flooding topology. This causes the problem of that partition being exposed to a congestive collapse. This problem lasts until the LSPs of the incorrect start node are timed out, which is typically 20 minutes but could be different if lsp refresh interval is configured).

Suggested workaround: To reduce the probability of this happening, prioritize the Area Leader to be the node with the highest degree of connectivity. Configuration statement:

```
arista(config-router-isis)#area leader priority ?  
  <0-255> Area leader priority
```

This workaround will not work in all situations since connectivity of a node will change and there may be multiple nodes with the same degree of connectivity. But it does reduce the probability of hitting this bug for a given topology. (429480)

- In the presence of tie breaker configuration for ECMP paths, the Rib agent may exit unexpectedly while handling BGP updates for a prefix having ECMP, when some paths of the given prefix are undergoing metric change. Workaround is to remove tie breaker configuration. (444000)

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)

Layer 1

- Speed change on interfaces with "speed-misconfigured" errdisable state after hitless reload causes forwarding agent restart. This leads to the flap of all connected interfaces. The workaround is to recover all "speed-misconfigured" interfaces by configuring speeds of affected /2, /3 and /4 ports to match the speed of the corresponding /1 ports before hitless reload. (352115)

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

- The forwarding agent may continuously restart if the speed is changed from 100G to 50G when a 50G QSFP transceiver is inserted. This only affects 2-lane 50G transceivers, and does not affect 100G transceivers operating in 50G mode. A workaround is to configure the speed before inserting the transceiver. The forwarding agent restarts can be stopped by temporarily configuring the affected interfaces in 25G mode. (424990)

- Speed changes on SFP ports cause all four ports in the port group to flap. (460034)

SKUs Affected: CCS-720XP-24Y6-F, CCS-720XP-24Y6-R, CCS-720XP-24ZY4-F, CCS-720XP-48Y6-F, CCS-720XP-48Y6-R, CCS-720XP-48ZC2-F, DCS-7050SX3-48YC8-F and DCS-7050SX3-48YC8-R

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)

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)
- SSU with Splunk enabled may result in extended boot times, which could lead to protocol timeouts or switch cold rebooting. (343629)
- When running 'reload fast-boot' or 'reload hitless', ports with optical transceivers may flap, causing traffic disruption (348172)

SKUs Affected: DCS-7050SX-128-F, DCS-7050SX-128-R, DCS-7050SX-128-SSD-F, DCS-7050SX-128-SSD-R, DCS-7050SX2-128-F, DCS-7050SX2-128-R, DCS-7050TX-48-

F, DCS-7050TX-48-R, DCS-7050TX-64-F, DCS-7050TX-64-R, DCS-7050TX-72Q-F and DCS-7050TX-72Q-R

- When performing SSU on an MLAG setup, SSU can cause peer links on port-channel to flap, resulting in large traffic loss. (405788)
- During SSU, the ARP timeout may exceed 60s since kernel bootup is slower. Workaround is to first upgrade to 4.21.3F and then to the target release. (423992)

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

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)
- MSS revokes redirect rules from a MLAG peer when the last member interface on any of the port channels connected to the service device goes down. This results in redirected traffic traversing that MLAG peer being dropped.

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. (230653)

- 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)
- If there are large number of intercept hosts in L2 policies in the order of 10k, then shutting down management cvx feature may not cleanup all DirectFlow rules (285821)
- ConfigAgent prints an error in its log that it cannot find Aegis toggle. No functionality should be effected. (372579)
- The Macro-Segmentation Service (MSS) interfacing with a firewall running in L2 transparent mode does not enforce a firewall policy that requires intercepting traffic on a VNI for the first time while one of the peers of the MLAG Service VTEP is being reloaded. MSS starts intercepting that set of hosts when the peer reload completes. (373587)
- In a CVX deployment, the Controllerdb agent can experience a memory leak which may trigger a Controllerdb restart which may trigger a leader re-election. (379620)

- When a security policy on a CheckPoint firewall includes unsupported services, the Macro-Segmentation Service (MSS) intercepts traffic for all services for the specified end-points in the policy. (433067)
- The MssPolicyMonitor agent throws an exception when it received an unexpected reply from a Palo Alto Networks firewall (435713)
- Vxlan Controller Service may ignore vARP MACs published by switch for VNIs not shared by all VTEPs publishing vARP VTEP IP. (448453)
- The OpenStack agent may restart when topology changes (451613)
- In a CVX + EVPN DCI topology, if a VTEP is moved from one CVX cluster to the other, EVPN on the new CVX cluster can fail to advertise IMET routes for that VTEP. (453153)

MPLS

- When the refresh overhead reduction extension is used, some RSVP messages may never be sent out (shown as "State: path-only" under "show mpls rsvp session").
As a workaround, the extension can be turned off and on again. (441063)

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)

- When routing is disabled, ARP control packets share the CPU queue with other CPU bound packets. This could lead to MLAG disassociation and network outages.

The workaround is to either enable IP routing or have a static ARP entry on the MLAG peers. (90138)

- 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. (136944)

Series Affected: 7368 Series

SKUs Affected: DCS-7010T-48-DC-F, DCS-7010T-48-DC-R, DCS-7010T-48-F,

DCS-7010T-48-R, DCS-7020SR-24C2-F, DCS-7020SR-24C2-R, DCS-7020SR-32C2-F, DCS-7020TR-48-F, DCS-7020TR-48-R, DCS-7020TRA-48-F, DCS-7020TRA-48-R, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050QX-32-F, DCS-7050QX-32-R, DCS-7050QX-32-SSD-F, DCS-7050QX-32-SSD-R, DCS-7050QX-32S-F, DCS-7050QX-32S-R, DCS-7050QX-32S-SSD-F, DCS-7050QX-32S-SSD-R, DCS-7050QX2-32S-F, DCS-7050QX2-32S-R, DCS-7050SX-64-F, DCS-7050SX-64-R, DCS-7050SX-64-SSD-F, DCS-7050SX-64-SSD-R, DCS-7050SX-72-F, DCS-7050SX-72-R, DCS-7050SX-72-SSD-F, DCS-7050SX-72-SSD-R, DCS-7050SX-72Q-F, DCS-7050SX-72Q-R, DCS-7050SX-96-F, DCS-7050SX-96-R, DCS-7050SX-96-SSD-F, DCS-7050SX-96-SSD-R, DCS-7050SX2-128-F, DCS-7050SX2-128-R, DCS-7050SX2-72Q-F, DCS-7050SX2-72Q-R, DCS-7050SX3-48YC12-F, DCS-7050TX-48-F, DCS-7050TX-48-R, DCS-7050TX-48-SSD-F, DCS-7050TX-48-SSD-R, DCS-7050TX-64-F, DCS-7050TX-64-R, DCS-7050TX-64-SSD-F, DCS-7050TX-64-SSD-R, DCS-7050TX-72-F, DCS-7050TX-72-R, DCS-7050TX-72-SSD-F, DCS-7050TX-72-SSD-R, DCS-7050TX-72Q-F, DCS-7050TX-72Q-R, DCS-7050TX-96-F, DCS-7050TX-96-R, DCS-7050TX-96-SSD-F, DCS-7050TX-96-SSD-R, DCS-7050TX2-128-F, DCS-7050TX2-128-R, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-F, DCS-7060SX2-48YC6-R, DCS-7150S-52-CL-F, DCS-7150S-52-CL-R, DCS-7150S-52-CL-SSD-F, DCS-7150S-52-CL-SSD-R, DCS-7160-32CQ-F, DCS-7160-32CQ-R, DCS-7170-32C-F, DCS-7170-32C-M-F, DCS-7170-32C-M-R, DCS-7170-32C-R, DCS-7170-32CD-F, DCS-7170-32CD-R, DCS-7170-64C-F, DCS-7170-64C-M-F, DCS-7170-64C-M-R, DCS-7170-64C-R, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7280CR2A-30-F, DCS-7280CR2K-30-F, DCS-7280CR2M-30-F, DCS-7280QR-C36-F, DCS-7280QR-C36-M-F, DCS-7280QR-C36-M-R, DCS-7280QR-C36-R, DCS-7280QRA-C36S-F, DCS-7280QRA-C36S-M-F, DCS-7280QRA-C36S-M-R, DCS-7280QRA-C36S-R, DCS-7280SE-64-F, DCS-7280SE-64-R, DCS-7280SE-68-F, DCS-7280SE-68-R, DCS-7280SR-48C6-F, DCS-7280SR-48C6-M-F, DCS-7280SR-48C6-M-R, DCS-7280SR-48C6-R, DCS-7280SR2-48YC6-F, DCS-7280SR2-48YC6-M-F, DCS-7280SR2-48YC6-M-R, DCS-7280SR2-48YC6-R, DCS-7280SR2A-48YC6-F, DCS-7280SR2A-48YC6-M-F, DCS-7280SR2A-48YC6-M-R, DCS-7280SR2A-48YC6-R, DCS-7280SR2K-48C6-M-F, DCS-7280SR2K-48C6-M-R, DCS-7280SRA-48C6-F, DCS-7280SRA-48C6-M-F, DCS-7280SRA-48C6-M-R, DCS-7280SRA-48C6-R, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F, DCS-7280SRM-40CX2-R, DCS-7280TR-48C6-F, DCS-7280TR-48C6-M-F, DCS-7280TR-48C6-M-R, DCS-7280TR-48C6-R, DCS-7280TRA-48C6-F, DCS-7280TRA-48C6-M-F, DCS-7280TRA-48C6-M-R, DCS-7280TRA-48C6-R, DCS-7304, DCS-7308, DCS-7504N, DCS-7508N, DCS-7512N and DCS-7516N

- Enabling the Macro-segmentation service can result in duplicate packet delivery of broadcast, multicast or unknown unicast traffic sourced from the host whose traffic is intercepted. Additionally, one copy can be delivered to the intercepted host that is the source of the traffic. (212480)
- 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)

- Any update in OpenConfig removes the "shutdown" configuration from Management interfaces (250235)
- The CLI command "route-target import auto <ASNumber>" for VLAN VRF does not work and hence should not be used. Alternatively, the suggested configuration is "route-target import auto". In this case the AS Number is picked up from the BGP configuration.

Example:

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

Suggested Usage

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

- The MacroSegmentation Service (MSS) feature doesn't work in conjunction with the VARP VTEP IP configuration required for VXLAN routing where a subset of VTEPs are L2 VTEPs. (263532)
- Support for VRRPs in OpenConfig not fully implemented (268736)
- OpenConfig fails because of an invalid type in ISIS with the log message 'expected attribute type nominal.U8, got uint8' (293256)
- 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)
- OpenConfig affected by CVE-2019-6486 (356552)
- CVE-2019-6486: A malicious peer using P-521 or P-384 elliptic curve encryption may cause a TerminAttr DoS. (356638)

- 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
- Nitrogen ODL gives an error reading the path /interfaces/interface/arista-vxlan (370169)
- When configuring VXLANs from the EOS CLI, "ip address virtual" is supported and mapped to YANG, but "ip virtual-router address" is not supported. (372193)
- Mlag agent may crash if the system is being put in Maintenance Mode if interface profile has shutdown max-delay enabled. Workaround is as follows:
1. Create a maintenance interface profile which does not have shutdown max-delay configured with profile name being lexicographically largest.

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


2. Create an interface group constituting all the physical ethernet interfaces of peer-link and attach the above interface profile that does not have shutdown max-delay configured.
3. Create a unit with this group and quiesce it. (380387)
- 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)
- While receiving large number of Igmp joins, some of the IGMP joins might end up getting dropped causing loss of traffic for the affected groups. (413880)
- DirectFlow 'action output nexthop <ip addr>' ignores TTL of incoming packet and forwards the packet even if the TTL has expired. (417994)
- The Aaa agent might restart when LDAP servers are unreachable (436291)

- ZTP can stop cycling speeds on front-panel ports whose capabilities have no supported speeds for a possible speed-group configuration; this will not affect management ports. For example: a 40G-only transceiver inserted into a port with a default speed-group setting of 25G. (444861)
- A YANG operation to configure VLANs on an EVPN instance of type 'VLAN aware bundle' will only configure one VLAN with an ETID of zero, even if multiple VLANs with ETID of zero are specified. (448468)
- 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)
- Removing a bash action from an event-handler via "action bash" can result in inability to copy and paste that particular event-handler config. A workaround is to use "action bash :". (452832)
- If any interface is configured with, "ip address unnumbered," configuration changes made with gNMI, NETCONF, or RESTCONF do not apply. To workaround the issue "ip address unnumbered" needs to be removed from the configuration via the CLI or EAPI. (457161)
- ISIS configuration on a non default VRF via OpenConfig fails (460407)

Multicast

- An S,G route may not properly cleanup previously sent S,G joins. This would prevent S,Gs from being deleted in the upstream router. A work around is to flap the interface in which the join is sent on. (291756)
- If large number of IGMP reports are received in a very short interval, some of them might get dropped resulting in the multicast route not being created. The routes will be re-learned on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIgmp", consider increasing the bandwidth for this Copp class. (356675)
- When using bess as the forwarding scheme, if the multicast routes are greater than 2500 in scale then they do not get displayed in "show platform sfe counters" and "show ip mfib software" (371843)
- 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)
- When using Igmp host proxy, longer than expected traffic gaps might be seen when proxying source specific joins. (408784)

- Changing PIM DR priorities while PIM SSM IPv4 Non-DR OIF installation feature is enabled on multiple interfaces with thousands of multicast routes on 7050 and 7300 platforms results in device getting into a bad state with ACLs. Once the device is in this state, the device needs to be rebooted. Disable the feature and change the priorities and then reenables the feature in order to not get into this state. (416636)
- In an environment with a very large number of MSDP peers, during cleanup the Msdp Agent may crash. (418190)
- When IGMP snooping proxy is not enabled and there is multiple IGMP queries on a VLAN, the membership of the VLAN may prematurely age out if multiple general queries are received in quick succession. This will result in traffic loss. (422416)
- Pimsm could restart when link is flapped while route to RP changes. (443240)
- When using Igmp Snooping proxy in the network, DCS-7170 devices might cause multicast traffic loss. (448375)
Series Affected: DCS-7170 Series
- When static multicast routes are configured, in some cases the IgmpHostProxy agent may not create reports for all of the mroutes. As a workaround, after this occurs, deleting and recreating the static mroute should fix the IgmpHostProxy agent state. (449133)

DCS-7150 Series

- If a multicast boundary is applied on an SVI which also has IGMP snooping enabled, IGMP snooping will not take effect. (138610)

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

- In scenarios with continuous MAC address flapping where VXLAN remote VTEPs are being added and removed, the forwarding state for hosts pointing to some remote VTEPs may be incorrect. This may lead to brief misforwarding of packets or a restart of the SandMact agent (68683)
- 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)
- When a sub interface under lag member configured with shaping is removed, IP ACLS stops working on port-channel sub-interfaces. (162785)
- 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)
- 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)
- If an ACL is attempted to be applied on a subinterface, whose parent port is not a routed port, the CLI will timeout. (229537)
- PBR policies applied on interfaces do not work if the interface has some subinterfaces that also have PBR policies applied. (243286)
- 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 depolymnt 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 errornously 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)

- Unexpected DSCP rewrite may occur on a port if the port has an internal VLAN that was previously assigned to a port which had DSCP rewrite map applied.

The workaround is to apply and remove a dummy map on that interface using "[no] qos rewrite traffic-class to dscp". (274484)

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

- 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)
- PIM multicast traffic may experience a brief interruption when one or more linecards on the same modular chassis enters or leaves TAP Aggregation mode. (279972)

- 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)
- SandTunnel Agent may restart unexpectedly with evpn mpls vlan aware configuration. (344053)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- 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)
- 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, if VXLAN header stripping is enabled on a Tap port and recirculation is required, policy maps applied on such a Tap port may not take effect. Moreover, unrelated traffic from other Tap ports may be misdirected to the tool group to which the policy map may direct the stripped VXLAN traffic. (371416)
- 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)
- When the DirectFlow action 'output nexthop' is applied, the TTL for the redirected bridged packets is set to 255. (375100)
- 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)
- In the presence of virtual MAC with mask configuration, packets that undergo recirculation may not be properly forwarded. The workaround is to eliminate the configuration that causes the recirculation to happen. (387535)

- If VXLAN is configured, PBR and QOS policies won't apply to traffic ingressing on a MLAG peer-link. (390804)
- Traffic subject to CPU traffic policy `_policer_` action will not be `_forwarded_`. This includes multicast control traffic which is replicated to both the CPU but also forwarded (e.g. PIM, OSPF HELLOs, etc.). Unroutable unicast traffic is also flooded, and so will not be forwarded if matched by a CPU traffic policy policer action. An example policy of concern is

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

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

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

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

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

- Insertion of a fabric module may cause brief traffic loss on the device (402508)
SKUs Affected: 7504R3-FM and 7508R3-FM
- When `configure replace` is used to configure large acl configurations, the configuration might time out and the new configuration may not apply. The workaround is to make smaller changes in the acl configuration. (404797)
- With scaled Acl configurations, SandAcl agent restart may result in additional agent restarts. (408260)
- TCAM reservations are not preserved after a system reboot. (410069)

- 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)
- When the DirectFlow 'priority' keyword is set greater than 65534 the configured priority will not be set to the correct value and the SandAcl agent may continually restart. (412206)
- 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)
- In an RFC 5549 environment, if an IPv4 packet egresses out of an interface with only IPv6 address configured, the deny log rules in the ipv4 egress acl do not work. The packet gets denied, but the syslog messages are not generated. There is no workaround. (415305)
- A switch reload with qos policy configured may result in SandAcl agent restart. (417935)
- Storm control maximum burst size is fixed to 10kB which can lead to storm control drops for bursty incoming traffic. (423231)
- With WRR scheduling on many ports and linerate across multiple traffic classes there can be a small amount of egress buffer drops. (425123)
- Sometimes on reload MACsec sessions on MACsec proxy ports stay down when FIPS restriction is enabled. (430728)
- In Tap Aggregation mode, there is a small timing window where MAC addresses could be learned on an interface which is a member of a port-channel if its link state is toggled repeatedly. The effect is that traffic will be dropped if the destination MAC matches these learned addresses.

Use 'clear mac address-table dynamic' to clear the MAC address table. (432544)

Series Affected: DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7500R and DCS-7500R2 Series

- The following commands, "show ip[v6] hardware fib routes", "show ip hardware ale l2-adj", and "show ip[v6] hardware fib fec" when used with very large output may crash SandL3Unicast agent when interrupted via control-c or by quitting pager. (433841)
- Series Affected: DCS-7020, DCS-7048, DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series

- If BGP Flowspec is enabled in multiple VRFs, Flowspec matches from one VRF may apply to interfaces in another VRF where Flowspec is enabled. (439105)
- The SandAcl agent may restart repeatedly when multiple UDF payload matches with the same header and offset but different masks are configured in access-lists. (439400)
- 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)
- If L3 Forwarding agent restarts in the presence of VXLAN MAC addresses, it might lead to unexpected restart of L2 Forwarding agent. This might result in temporary flooding of traffic toward some destinations. (440226)
- When the outgoing interface for a Vxlan underlay route changes from one forwarding chip to another, it can impact VXLAN overlay route convergence. (440233)
- Performing many updates to a PBR policy in a short amount of time may cause the SandAcl agent to restart unexpectedly. (441912)
- IP multicast packets may get mis-routed in an MLAG configuration when the peer chassis is rebooted in a multi FAP system.
A restart of the SandMcast agent will clear the left over entries from the hardware. (444856)
- 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
- When VXLAN routing is configured, restart of the SandL3Unicast agent may lead to a crash and hitful restart of the SandL3Unicast agent. (448317)
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- V6 unknown multicast packets are not headend replicated to remote VTEPs. (448950)
- 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

- VXLAN routing over MACsec proxy port stops working when a proxy port that is not bound to an external interface is updated to be bound to an external interface. (455759)
- Sflow samples of VXLAN packets terminated by the switch, will have their VXLAN headers stripped. (457096)
- "match nexthop-group" filter is not supported in PBR rules when the corresponding nexthop-group is referenced by ECMP routes. (458189)
- Pre existing DirectFlow flows may temporarily fail to be effective while new flows are being installed. (460230)

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

- 100GBASE-PSM4 modules may not link up after a reboot. The workaround is to remove and reinsert the module. (354176)

SKUs Affected: 7500RM-36CQ-LC

- 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-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

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)
- In some yet unknown cases, IPsec engine may lockup. Currently, only a restart of SandFap Agent will fix the problem. (425295)

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

DCS-7020RA, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards Series

- VXLAN packets that are decapped and bridged are incorrectly subject to the IP ACL applied on the ingress port. (221457)
- Packets may be unexpectedly dropped if ACLs are applied and an interface has a double tag VLAN mapping. (229765)
- The HaloFpga agent will continuously restart if SandFap agent is restarted.

The workaround is to powercycle the linecard or powercycle the switch if it's a fixed system. (252069)

- With AlgoMatch enabled, the tc-counters TCAM profile applied, an IPv6 security ACL applied on an interface, and an IPv6 PBR policy applied on the same interface as the ACL, actions for that ACL (deny, count, etc) will not apply to incoming traffic. (399725)

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)
- TTL exception handling during MPLS termination may not occur after a parity error in the IPPD_FWD_ACT_PROFILE hardware table.

To address these errors, schedule downtime and restart the SandTunnel agent. (424100)

- Disabling MPLS routing (globally or under an interface) does not take effect. (426023)
- A statically configured MAC entry may be replaced by a dynamically learned MAC entry leading to transient mis forwarding of packets destined to that MAC address. This transient condition gets corrected in the next poller cycle. (428580)
- SandAegis agent may crash when the last IPv6 rule is deleted from a traffic-policy which has both IPv4 and IPv6 rules (435733)
- SandTunnel agent may restart when SandFapNi agent restarts, such as when a linecard is removed and re-inserted in a chassis. (441928)
- Egress MPLS tunnel counter may not work after the reachability of tunnel nexthop goes through churn. (443473)
- On Switch with Jericho2 FAP(s), if Pimsm spt threshold is configured as infinity, Pimsm multicast traffic forwarding may not work properly. (447366)
- PFC configuration doesn't take effect if configured after LLFC is configured first on a port. (457400)
- After a restart of the forwarding agent packets originating on the CPU are dropped. (458206)
- Link flaps may occur when using 400g optics. (459729)
SKUs Affected: 7800R3-36P-LC

DCS-7280E and DCS-7500E Series

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

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

- IPv4 egress RACLs will not filter multicast traffic in shared mode. In unshared mode, multicast traffic is filtered only if ingress replication is enabled. (131660)
- An IPv6 Egress ACL applied to a Port Channel is not applied to a newly added member interface, if the member interface already has the same ACL applied to it. The workaround is to remove the ACL from the interface before adding it to the PortChannel. (141736)
- SAND_POSSIBLE_STUCK_PORT may be logged on an Ethernet interface after switch initialization, which can cause an interface to not transmit packets despite being up.

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

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

DCS-7160 Series

- Under certain race conditions between LAG configuration changes and XpCentral agent restart, the system could get into a state where XpCentral agent continuously undergoes a restart.

The workaround is to remove all LAG interfaces and re-configure them after XpCentral agent has restarted successfully. (161558)

- When mirroring destination is configured to CPU, under heavy traffic scenarios, it may be possible that switch engine ASIC may leak some packet buffers, and may eventually run out of them causing switching to stop forwarding packets on one or more ports.

Workaround is to not configure mirroring destination as CPU for a prolonged duration under sustained linerate traffic. (225142)

- 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 crash when the switch is reloaded if PBR policies with nexthop VRF actions are applied on an interface in the non-default VRF. (262713)
- Value of outer/underlay UDP source port of VXLAN encapsulated packet may vary based on overlay packet size for a given overlay packet flow (identical overlay/inner 5 tuple). This could lead to unnecessary ECMP load balancing of egress VXLAN encapsulated packet which can further lead to performance issues for a overlay TCP session/flow. (422669)
- 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)

DCS-7170 Series

- If a LAG was created with member-ports in shutdown state, outgoing traffic can incorrectly get hashed to those member links leading to traffic loss. (250543)
- On the 7170, when dynamic NAT is configured on the layer 3 interfaces used for ECMP, return traffic does not get NAT applied unless it comes in on the same ECMP member of the outbound NATted traffic. (398104)
Series Affected: DCS-7170 Series
- The maximum MTU supported is 9210 bytes instead of 9214 bytes. (457427)
- When performing egress multicast nat, only the 19 least significant bit of the mac address is updated. (459603)

DCS-7170 Series

- If a mac learnt on a vxlan interface on an MLAG peer is promoted due to an event such as a MLAG peer-link going down, it may cause the forwarding plane agents to restart. (261152)

DCS-7050X[2], DCS-7060X[2], DCS-7250X and DCS-7300 Series

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

- Traffic streams matching DirectFlow flows with VLAN or MAC rewrite actions could face brief disruption across a StrataL3 agent restart. (111833)
Series Affected: DCS-7050 and DCS-7050X Series
- The WRED drop decision obtained from the ECN configuration on tx-queue 0 is also used to determine whether packets should be dropped on UC8. This can result in control packets enqueued on UC8 getting dropped if UC0 is experiencing WRED drops for non ECT packets enqueued on it. A workaround is to use a queue other than tx-queue 0 for applying ECN. (125803)
Series Affected: DCS-7050 Series
- MAC ACLs and PACLs cannot co-exist on the same interface if QoS policy ACLs exist in the switch (186949)
- On an MLAG pair of switches, with fast-MAC redirect enabled, the Macro-Segmentation Service (MSS) can fail to recover from a MLAG port channel member link flap.

The workaround is to shut/ no-shut the port channel connected to the intercepted host experiencing traffic loss. (213199)

Series Affected: DCS-7050X and DCS-7060X 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: DCS-7010, DCS-7050, DCS-7050T, DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7260X, DCS-7260X3, DCS-7300, DCS-7300X and DCS-7320 Series
- If nexthop resources are exhausted as indicated by syslog VXLAN_HWHER_NEXTHOP_RESOURCE_FULL, forwarding of BUM traffic to that VTEP will be affected. (246391)
Series Affected: DCS-7050X and DCS-7060X Series
- During an SSO failover, Strata Fabric agent may not detect hardware, resulting in a fatal error reboot with increased traffic outage. (250129)
Series Affected: DCS-7300 Series
- In an L3 EVPN setup, traffic that is sourced by a VTEP that is a pure VXLAN router (i.e., it performs no VXLAN bridging and has no hosts attached to

it) might incur some loss at the receiving VTEP, where the packets are decapsulated and forwarded.

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

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

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

- If SSO redundancy mode is configured and a linecard is removed before it is fully initialized, RPR switchover may be performed. (283088)
Series Affected: DCS-7300X Series
- 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
- 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
- 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)

- 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-7300 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)
- Restart of StrataL2 may lead to continuous StrataL2 restarts. Workaround is to reload the system. (408138)
- After lowering the priority of a flow, traffic may continue to match on that flow even if existing flows now have higher priority. Updating flows priority again or toggling DirectFlow should restore expected behavior. (408312)

- 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

- Flow control transmit PAUSE packets won't be sent in case of congestion in current device or upstream device (410180)

SKUs Affected: CCS-720XP-24Y6-F, CCS-720XP-24ZY4-F, CCS-720XP-48Y6-F, CCS-720XP-48ZC2-F, DCS-7050SX3-48YC8-F and DCS-7260CX3-64E-F

- 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, DCS-7300, DCS-7300X and DCS-7320 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)

- 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)
- If the number of next-hop groups configured exceeds 512, then STRATA-3-MPLS_MAC_DA_PROFILE_TABLE_FULL syslog message is logged and L3 forwarding agent may continuously restart.

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

- Ethernet49-54 may not link up with peers over long copper cables. Workaround is to enable autoneg on both ends of the link with 'speed auto <speed>' command. (421121)
SKUs Affected: DCS-7060SX2-48YC6-F and DCS-7060SX2-48YC6-R
- After a speed change, flows with mirroring actions on affected interfaces may not be mirrored as expected. Strata Slice agent can be restarted to restore expected behavior. (422221)
Series Affected: CCS-720XP, DCS-7050SX3 and DCS-7260CX3 Series
- 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
- 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)
- Restarting the StrataL3 forwarding agent, or performing an SSU operation may cause the StrataL3 forwarding agent to restart repeatedly. The switch has to be rebooted to recover from this condition. (450331)

DCS-7050X2 and DCS-7300 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-FixedSystem agent to restart following which the interface continues to transmit properly. (245501)

- If the switch receives a VLAN-tagged VXLAN encapsulated packet destined to the switch's VTEP IP address, but the packet's destination MAC is not the switch's MAC, such a packet is flooded to all ports in the ingress VLAN including the ingress physical port. (261494)

DCS-7010T Series

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

DCS-7060X and DCS-7300 Series

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

DCS-7050 and DCS-7300 Series

- 25G optics are not supported. Inserting one would lead the forwarding agent to crash continuously. (199036)

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)

- When a QSFP port is set to 4x25G rate, inserting a QSFP which is only capable of 40G can trigger a forwarding agent restart.
A workaround is to configure the port for 40G or 4x10G rate before inserting the transceiver. (405185)
Series Affected: DCS-7260CX3 Series
- Forwarding agent may restart when a 40GBASE-SRBD transceiver is inserted after the hitless reload. To work around, configure "speed forced 40gfull" on the desired interface prior to inserting the transceiver. (417311)

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

- Egress ACL on SVI for BUM traffic will not match ACL rules for packets being sent across fabric interfaces. (284543)
- When performing SSU, there may be failure in table programming resulting in fatal error cold reboot with extended traffic outage. (399774)
- After performing ASU reload, OSPF sessions may not come up, resulting in complete traffic drop. (410242)

- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (416694)
- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (417223)
- If a hitless forwarding agent restart is attempted while the hardware configuration does not match the running configuration, then it may cause an additional forwarding agent restart and loss of traffic. This may also happen if the forwarding agent is restarted while some interfaces have been disabled due to insufficient logical ports. (417835)

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

- When "platform trident forwarding-table flexible" is configured, L3 unicast forwarding to hosts / L3 multicast forwarding may not work resulting in unexpected behavior. The workaround is to remove the CLI and use "platform trident forwarding-table partition" instead. (421084)
- If an interface has a 1g transceiver and other interfaces in the same speed group are configured for 25g, the slice agent will restart continuously.

Workaround is to configure all ports in the speed group for 10g. (433880)

- When L3 forwarding agent restarts, some host routes may be programmed incorrectly resulting in drops. The work around is to reload the system. (447854)
- When L3 forwarding agent restarts, some host routes may be programmed incorrectly resulting in drops. The work around is to reload the system. (451933)

7368 Series

- 802.3X pause is not supported (348758)
- L3 flooded packets do not preserve VLAN priority values. (365998)
- MPLS traffic will be dropped if the nexthop MAC address ages out or not known.

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

- If test patterns are configured on certain interfaces in startup-config, or in a configuration used with "configure replace" which also changes interface speeds, Strata forwarding agent may restart causing all links to flap upon booting or upon replacing configuration, respectively. (435141)

Series Affected: DCS-7060X4 Series

CCS-720XP Series

- If link speed is changed on a 2.5G interface, Strata forwarding plane agent may unexpectedly restart with "Process died with signal 3 \ (SIGQUIT\)" signature. Subsequent restart of the Strata agent will resolve the issue. (368424)
- In multi-host authenticated mode, switch may send out untagged unicast identity requests to the tagged EAPOL supplicants which may cause them to not respond to the identity requests breaking their authentication. (447422)
- Dynamic Authorisation response packets(i.e. ACK/NAK sent by NAS to AAA) may be ignored by AAA server if NAS is configured to use loopback interface. (457009)

DCS-7050X and DCS-7300 Series

- 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)
- A single event upset (SEU) on an MMU configuration table may incorrectly clear the entry configuration. (187144)
- 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)
- Sometimes, an interface can stop transmitting despite there being queued packets on the interface's egress queues. Doing a "shutdown" followed by "no shutdown" on the interface can resolve the problem. (400610)

DCS-7300 and DCS-7500 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 crash repeatedly in rare corner case and Management0 interface is not reachable in Modular systems. Doing switchover will address the issue. (369565)

DCS-7020RA, DCS-7160, DCS-7170, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards 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)

Containerized EOS Series

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

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)

vEOS Router Series

- The OpenFlow agent fails to handle packet-in messages (145001)
- 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)
- 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 a vEOS-DPDK system, if the MTU of the Ipsec or GRE tunnels is configured to be larger than the MTU of the underlying interface, packets won't be forwarded correctly. The work-around is to configure the tunnel MTU to be just below the MTU of the underlying interface. For example, the MTU for a GRE tunnel should be 28 bytes less and for an Ipsec Tunnel it should be 108 bytes less to account for the overhead of the encapsulation headers. (343876)

- 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)

BGP EVPN L2/L3 VXLAN/MPLS

- When a secondary IP address is configured on a VXLAN interface, the switch may fail to send EVPN MAC address advertisement routes for locally learned MAC addresses. (446070)
- When using EVPN MPLS, if the ArpSuppression agent restarts, it may fail to populate some of the ARP entries corresponding to the remote bindings from the received MAC-IP routes. (457156)

Limitations and Restrictions in 4.23.2F

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 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)
- BMP is not supported for router bgp instances configured in non-default VRFs. (210363)
- SSO is not supported when BFD is configured for ISIS IPv6 sessions. (239704)
- An LACP port channel configured with BFD per-link and echo may experience BFD session flap on existing member links if new member link(s) are added to the port channel while the BFD peer device undergoes Stateful Switchover (SSO). (255042)
- Eos does not support programming IPv6 ECMP paths to kernel (258387)
- The DSCP value of BGP control plane packets set by the command 'bgp transport qos dscp <value>' will be overwritten when the platform 'qos rewrite dscp' is enabled and the outgoing interface is a SVI based on traffic-class to dscp map present on the platform. (269764)
- VRF support is 256 VRFs on 32G memory platforms, 64 VRFs on >= 8G memory platforms, 32 VRFs on < 8G memory platforms. (271478)
- 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)

- The command "[no] service routing configuration route-map no-set-community explicit-members" affects the behaviour of the route-map command "no set community [community-list] [member]+ [action]". If "service routing configuration route-map no-set-community explicit-members" is enabled, only "set community" statements in the routemap that are referred to by the list of members in the "no set community" command will be affected. The default behavior is to remove all "set community community-list" commands from the routemap. (330839)
- 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-7050, DCS-7050T, 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)

Multi-agent

- If a BGP peer is configured for BFD fallover, the BFD session is not cleaned up even after the BGP peer is unconfigured. (217199)
- Traffic destined to an L3EVPN route whose underlay has ECMP BGP-LU routes, may experience temporary traffic loss when the underlay goes from n way to

n-1 way or lower. The duration of loss can vary depending on the scale of BGP-LU routes or L3EVPN routes. (251692)

- When the multi-agent routing protocol model is configured, the maximum number of next hops for an OSPF ECMP route is 128. (274888)
- With very a very low multi-hop BFD timer configured, a BFD session may flap when a path (belonging to an ECMP set), over which the BFD session is established, goes down. (287571)
- Route-map specified as match-map for dynamic prefix-list is evaluated only against BGP routes. (345444)
- The IpRib agent runs out of addressable memory when a large number of routes are leaked to a large number of VRFs. (346575)
- Routes which are imported into a target VRF using a "leak routes source-vrf" policy cannot be matched on the basis of the source-protocol from the source-VRF as part of the "rib fib policy" (403739)

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)

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)

- If 'traffic unprotected allow' is configured in a Macsec profile, agent restart or Stateful Switchover may result in traffic loss (362766)
- 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)

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

MLAG

- Spanning-tree mode backup does not work in an MLAG configuration. (15151)
- 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)

Accelerated System Update

- Aborting 'reload fast-boot' command with Ctrl-C is unsupported (344574)
- Rib agent can restart when the scale of IS-IS adjacencies are huge and multiple IS-IS adjacencies flap and/or multiple prefix-segments are flapped (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)

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)

- 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)
- Mss does not support traffic interception for the hosts connected to Service VTEP. (354754)
- 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)
- 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):
- 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 bypassPeerLink
    priority 65535
    table trident ifp
    match input interface Po1
    match ethertype ip (375185)
```

```
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)
```

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)

General

- CLI does not allow access to files whose names contain spaces. (539)
- 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)
- Private VLAN is not supported in the EOS image distributed to US customers. (159408)
- 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)
- Copy commands with sftp: or scp: URL do not work via eAPI by passing password through the "input" parameter. SSH keys have to be setup to run the commands without a password. (283716)
- EOS support alphanumeric VLAN names, however ODL gives an error when reading the leaf /interfaces/interface/routed-vlan/config/vlan if the name contains alphabetic characters.

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

- Enabling next-hop sharing (using "ip hardware fib next-hop sharing" command) could result in slower convergence in the case of link failures. (369305)
- For static tunnel interfaces, duplicate or conflicting GRE tunnels are tunnels that have the same source, destination, tunnel mode (GRE) and key configurations as an existing tunnel. A duplicate tunnel is always in the down state. When there are duplicate static tunnel interfaces and a config-replace is executed, it is non-deterministic which one of the conflicting tunnels will be in the up state. (373135)
- When an interface profile is used to configure an interface's access VLAN, the VLAN is not created automatically. The VLAN must be manually created with the vlan command. (376621)

- When using interface profiles, configuring interfaces may take a few minutes, depending on the number of interfaces being configured. It may be faster to apply an empty interface profile to the interfaces, and then update the profile. (382120)
- The CLI commands "locator-led [multifan | powersupply | chassis]" will not cause multifan, powersupply and chassis LEDs on the front panel of the system to flash, as they did on previous generations of fixed systems. The beacon LED on the back panel of the chassis is purple by default. The CLI commands "locator-led chassis" causes it glow red instead of expected blue. (395892)
 SKUs Affected: DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3K-32P4-F, DCS-7280PR3-24-F and DCS-7280PR3-24-M-F
- The configuration lock feature does not prevent changes to the configuration via SNMP. (402556)
- If an EOS CLI command is used to create a reference to a VRF that does not exist, a subsequent OpenConfig YANG operation may cause the VRF to be created in EOS. (411908)
- Large amounts of monitoring traffic processed at the control plane via aggressive sflow sampling or mirroring to CPU could affect other control plane functions due to CPU bandwidth contention. (415992)
- IP Locking may send out additional queries when STP changes state (even possible in portfast config which is common for the IPLocking use-case). (417806)
- When "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)
- In the 'show interfaces interactions' CLI, the list of speeds an interface is limited to does not account for the limitations that result from speed-group configuration. (434363)
- Configuring both TACACS+ and LDAP with SSL profiles may cause the Aaa agent to continuously restart (439333)
- 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)

- Installing 64-bit rpm fails with 32-bit EOS.swi on older supervisors (457100)

Multicast

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

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

- After ASU reload, if the first PIM hello packet received from the neighbor gets lost and If the neighbor has a high hello query-interval, it will result in traffic loss. Workaround is to use default PIM hello query-interval and increase the PIM hello query-count. (341447)
- With scale greater than 5000 multicast IPV6 routes, show platform sfe mroute ipv6 <vrf> commands might not display the right number of routes when performing multiple addition/deletion operations across VRFs. This is applicable only if user level multicast forwarding agent is configured to be used instead of kernel. (378841)
- Bessd might crash 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)

DCS-7150 Series

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

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

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

ARP resolution completes.

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

- Source port loopback suppression at high traffic rate can reduce the egress bandwidth of a port. This will happen when a port forwards incoming multicast traffic at a high rate to a multicast group which also has this port as a member. This port will receive back copies of the replicated traffic and drop them due to loopback suppression if there are no receivers on it at the egress stage of the port. (40294)
- NAT and VXLAN cannot be configured on the switch at the same time. (47258)
- When agile ports are configured for 40G with clause 73 autoneg (speed auto 40gfull) while using CR cables, the ports cannot be used in a port channel.

The workaround is to force the speed to '40gfull' on the ports. (55517)

- Systems with A0 silicon, packets are not forwarded to the egress port when egress truncation and timestamping are enabled together. Removing the configurations resumes packet forwarding. Truncation and timestamping are not supported together. (72756)
- The VXLAN flood VTEP list cannot be more than 128k VLAN-VTEPs, where VLAN-VTEPs is the number of VLANs multiplied by the number of VTEPs. (73228)
- Ingress port access control list (PACL) on a VXLAN-enabled interface is not supported. If ingress PACL is configured on a VXLAN-enabled interface, in-hardware VXLAN encapsulation of incoming packets is not expected to work correctly. VXLAN encapsulation of packets in software is not affected. Removing ingress PACL configuration on the affected interface restores correct VXLAN hardware encapsulation behavior on that interface. (76085)
- Multicast-routed frames get tx-mirrored even when they do not egress the tx-mirrored multicast destination due to VLAN suppression. This is a limitation of the switch ASIC. (87933)
- 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

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

- The L3 MTU on interfaces is not enforced on SVIs. (11659)

- 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)
Series Affected: DCS-7500E Series
- 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)
- Due to hardware limitation, if an egress IP ACL is applied on an interface which is part of a VLAN with an egress router ACL, the filtering behavior for egress traffic of that interface may be undefined. (82094)
- For packets that hit rules across multiple ACLs, the hit counts are accounted for only in a single ACL. PACL counters take precedence over RACLs and MAC ACLs, and RACL counters take precedence over MAC ACLs. (85440)
- When an egress IPv6 RACL is applied to a VLAN, IPv6 packets may not be subjected to MTU checks. (88778)

- An egress ACL on a destination port of a monitor session only takes effect for Rx mirrored packets which are IP forwarded. The egress ACL will not work for bridged packets or Tx mirrored routed packets. (89915)
- When a shaper is configured the output of the shaper has a tolerance of +30% with 100-byte frames. This reduces to +5% as the frame size is increased to 600 bytes. The shaper variance reduces when the port happens to be the endpoint of a tunnel. (93546)
- The egress per-queue byte counter is not exact for certain packet types. For routed packets the egress per-queue byte counter may be offset by +/-2 bytes. For any packets originating from the CPU, the per-queue byte counter may be offset by -12 bytes. (97660)
- Packets that attempt to do a GRE key lookup where the GRE key lookup misses will have egress IPv6 ACLs applied with the ingress VLAN (102455)
- GRE key forwarding packets that are recirculated will have no QoS policy applied to them. (102458)
- If a double tagged frame is received on an untrusted or DSCP-trusted interface, the CoS is retained if the outgoing frame is tagged, unless the frame is routed (131614)
- IPv6 Egress ACL deny logging is not supported on subinterfaces. (146487)
- PBR policy is supported only on interfaces in the default VRF. If a PBR policy is configured on an interface in a user defined VRF, the behavior is undefined. (148895)
- Disabling IPv4 routing on an L3 interface is not supported when IPv6 routing is enabled on that L3 interface. When "no ip routing ipv6 interfaces" is configured to enable the forwarding of IPv4 packets over IPv6 nexthops over interfaces that do not have IPv4 address, but only a IPv6 address, it will also allow routing IPv4 traffic over these interfaces. (151356)
- 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)

- 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)

- Linecards with HW revision < 02.00, IP-in-IP encapsulated packets might not get forwarded after termination if the packet size falls within the range 223–287 bytes. (193573)
SKUs Affected: 7500E-12CM-LC, 7500E-36Q-LC, 7500E-48S-LC and 7500E-72S-LC
- 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)
- On devices with Algomatch enabled, decapsulation of an outer IPv6 header will not work. (215833)
- 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

- Vxlan decapped packets will not match rules that match on 'tcp established' in egress acls (226582)
- 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)
- After applying or removing shaping configuration on an L2 subinterface, the L2 subinterface must be flapped for the configuration to take effect. (278730)
- 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)
- Enabling Route Exact Match may cause security ACLs (Port ACLs, Router ACLs & MAC ACLs) to stop working when AlgoMatch hardware is also enabled. (348149)
- Filtered mirroring is not supported on linecards or fixed-systems operating in TAP Aggregation mode when traffic steering is enabled. (349883)
- 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)
- Software forwarding is not supported for GRE or IP-in-IP tunnel terminated packets using decap-groups. (358369)
- The loop protect feature does not work on ports of Algomatch enabled linecards and fixed systems. (361130)
- 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)
- Different streams being rate-limited by the same policer might exhibit bias towards some streams and fair policing across all the streams is not guaranteed. (380599)
- Configuring more than 250 BFD sessions on a physical interface may result in many of those sessions flapping because of hardware drops on the interface. (389387)
- Drop Precedence based tail drop thresholds can be ignored when multiple queues are congested and buffering resources are low. (395693)
- If a security-ACL is configured on a sub-interface beyond the supported limit of acls, SandAcl might restart continuously until the number of configured acls is brought back within the supported limits (398815)
- Queue build-up doesn't show up for 64 byte packets in case of congestion via following commands:-
 - 1) platform fap diag diag cosq non_empty_queues
 - 2) show interfaces queue length

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

- PIM sparse mode feature which allows installation of outgoing interfaces on the non-DR for a faster failover does not work on L3 sub-interfaces. This feature relies on an egress ACL to stop multicast traffic from going out of the interface. Multicast egress ACLs do not work on L3 sub-interfaces due to a platform limitation. If the feature is configured, there will be duplication of traffic. The workaround is not use the feature on such interfaces. (405696)
- When a policy-map with match on inner VLAN is programmed on an interface, if single tagged packet traffic is sent to that interface, the behaviour of that packet is indeterministic as inner VLAN tag is calculated using an offset and vlan-tag-format qualifier doesn't differentiate between single and double tagged packets. (415587)

- When MPLS label is pushed on the traffic ingressing CoS trust port then EXP bits are not derived from TC based on packet's COS. (417108)
- Byte counts are not supported for IPSec tunnels. (419031)
- SflowAccel supports up to 200 VRFs. (423892)
- PAUSE frames with correct DMAC=01:80:C2:00:00:01, Ethertype=0x8808, opcode=0x0001 are never forwarded and are always consumed by the hardware. (426047)
- Header truncation for mirrored packets can not be used when the destination of a monitor session is a GRE tunnel. (428547)
- A port transmitting packets smaller than 100B that were tunnel terminated on the switch might be limited to ~90% linerate (429797)
- 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)

DCS-7280R and DCS-7500R Series

- An IPv6MPLS 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)
- 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)
- VXLAN and GRE tunnels cannot be configured at the same time. (248239)
- 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)

- MPLS ping and traceroute are not supported for VRF label routes. (343526)
- If mirroring to CPU is configured for an interface and an incoming MPLS packet on that interface matches a configured "mpls static vrf-label" entry, the packet will not be mirrored to CPU. (359362)
- 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-7048, DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series

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
- The current implementation of IPSec supports only encryption/authentication algorithm combinations of AES128/SHA1 or AES256/SHA256. No other combinations are currently supported. (359986)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Using an Ipsec tunnel interface for policy-based routing nexthop is not supported. (378861)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

- Ipsec tunnel interfaces are not supported as part of nexthop-group. (378890)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Dynamic NAT, AKA Port Address Translation, is not supported for IPSec tunnels. (384624)

DCS-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)
- 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)
- 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 linerate) to port which has been configured with Strict-Priority scheduling. (406550)
- Traffic drop may happen when we sent traffic to all 8 traffic-classes (aggregate rate nearing linerate) to port which has been configured with WRR scheduling. (406561)
- Only port-channel interfaces are supported as peer links. Single member port-channel interfaces are supported. (409958)
- Per TC based VOQ tail-drop thresholds isn't supported in software. (413344)
- Set dcspl will not apply to mpls routed packets that egress with mpls labels (418791)

DCS-7160 Series

- If there is a port ACL (PACL) or router ACL (RACL) that matches on L4 information for TCP/UDP/ICMP, and if the incoming packet has IP options, 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)
- ACLs with Vxlan is not supported, and may result in unexpected forwarding behavior depending on the configuration and pipeline profile. (217888)
- 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 mililiseconds of packet loss when VxLAN ECMP nexthop changes. (219888)
- There may be traffic loss of up to 100 mililiseconds 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)

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
- BFD sessions are not supported on Port-channel interfaces. The sessions will not come up. (250315)

- 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)
- After applying or removing shaping configuration on an L2 sub-interface, the L2 sub-interface must be flapped for the configuration to take effect. (281312)
Series Affected: DCS-7170 Series
- After setting CoPP rate = 0, a few packets (<10) might be let into the CPU, but everything will be dropped afterwards. (350676)
- Full mroute scale may not be achievable under certain scenarios due to hash collision for the multicast route table. (380545)
- When egress NAT and MLAG are both used, 'show ip mroute' may show NAT'ed addresses in addition to the receiving mroute addresses. (394571)

DCS-7050X[2], DCS-7060X[2], DCS-7250X and DCS-7300 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)
- Running the 10G SFP+ ports of the DCS-7050Q at 100Mbps is not supported (30130)
Series Affected: DCS-7050Q Series
- 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)
- L3 packets destined to a next hop whose MAC address is not learned and which fail the uRPF check configured on the incoming SVI are flooded in the ingress VLAN. (75660)
Series Affected: DCS-7050 Series

- 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)
- 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)
- 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 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)
- 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 switch may not be able to bridge or route VXLAN traffic if it hits one or all of following conditions:
 - Nexthop resources are exhausted as indicated by syslog
VXLAN_HWHOP_NEXTHOP_RESOURCE_FULL.
 - MAC table overflow as indicated by show platform trident mac-address-table missing.

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

- Flexible port-channel hashing may not hash accurately when the inner L4 header option is configured for IPV6 GRE packets. (267487)
- If a packet stream is policed by multiple policers, the policed rate may be lower than any of the configured policer rates. (271046)
- To perform filtered mirroring of a packet based on VLAN from an interface configured to perform VLAN translation, the ACL must be programmed with the translated VLAN. The mirrored packet will be the same as the packet at the source interface. (278599)

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

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

SKUs Affected: DCS-7050SX3-48YC8-F 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 hitless from EOS-4.22 or earlier results in /2 and /4 ports on Ethernet 1-12, 21-44 and 53-64 to remain inactive when /1 and /3 are configured in 10G or 25G. In order to activate /2 and /4 ports, configure 40G, 50G or 100G /1, then configure 10G or 25G back. These ports are also activated following a forwarding agent restart. (409057)
Series Affected: DCS-7260CX3 Series
- "reload fast-boot" is not supported with virtual IP address config. (411323)
- Speed change on a port with macsec enabled is not supported and could cause the port to get into a persistently un-authenticated state. Workaround is to disable macsec on the port, change the speed and then re-enable macsec on the port. (415737)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- Performing SSU from any release prior to 4.22.2 can lead to extended outage and mis-forwarding. (424812)
SKUs Affected: DCS-7050SX3-48YC8-F and DCS-7050SX3-48YC8-R
- IP ACL match over MPLS Encapsulated Packets is not supported. (426413)
- During Mpls PHP operation inner payload type gets ignored. (428546)
- When 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-7050CX3, DCS-7050SX3, DCS-7050X3, DCS-7060X3, DCS-7060X4 and DCS-7260CX3 Series
- With VXLAN dataplane learning disabled, locally learned MAC addresses will not age out if VXLAN encapsulated packets from remote VTEPs are received with inner packet's source mac as the locally learned MAC. The locally learned MAC entry needs to be cleared manually by issuing 'clear mac address-table dynamic vlan <VLAN> address <ADDRESS>'. (460607)

DCS-7050X2 and DCS-7300 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)
- 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)

DCS-7010T Series

- IEEE 1588 PTP transparent clock does not decrement TTL for PTP routed packets (55078)
- Enabling counters for Static or Twice NAT connection can cause FocalPointV2 agent restarts, when the number of connections exceed 275. (109048)

DCS-7060X and DCS-7300 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)

DCS-7260X3 Series

- LANZ is not supported. (217543)
- Configuring 100G, 50G, 40G, 25G and 10G simultaneously might result in unexpected packet forwarding behavior. (235075)

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

- 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 and 7300X3-48YC4-LC
- 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)
```

- Setting the default CoS value for a phone VLAN, using the CLI "switchport default phone vlan cos <>", takes effect only if the phone trunk port is in CoS trusted mode. If the phone trunk port is in CoS untrusted mode, then the port's default CoS value will be used for the phone VLAN. (390588)
- MAC Egress ACL configurations matching IP packets may not work when the IPv4/v6 per-interface counter feature is also configured. (411243)
- Leaf stateful switch upgrade is not supported on this platform. (415340)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- CLI output of "show platform trident l3 shadow mpls-vc-and-swap-label" will show push
as MplaLabelAction even for swap operation. (416134)
- 10/25G speeds do not support Macsec in this release (418855)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- 10/25G speeds not supported (420239)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- MPLS packet undergoing swap operation does not honor the MTU configured on ingress/egress interface. (422668)
- Only 2K MPLS routes are supported (425691)
Series Affected: CCS-720XP Series
- If the number of MPLS next-hop groups configured exceeds 512, then STRATA-3-MPLS_MAC_DA_PROFILE_TABLE_FULL syslog message is logged and L3 forwarding agent may continuously restart.

The workaround is to scale down the number of MPLS next-hop groups configured to stop the agent from continuously restarting. (428162)

- 50g port speed is not supported in this release (445999)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R

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)

CCS-720XP Series

- When hardware or software flow tracking is enabled, Vxlan encapsulated TCP flows may get exported with incorrect tcpFlags. The tcpFlags field track outer header's TCP flags hence for Vxlan packets it is undefined. There is no workaround to this issue. (416079)
- On receiving COA-REQ for changing VLAN for EAPOL supplicant, the NAS deliberately fails authorization for the first time in the new vlan. (449890)
- The per-interface configuration to ignore reauthorization timeouts from the AAA server ("dot1x timeout reauth-timeout-ignore") is not designed to work if manual reauthorization is requested from the command line. The reason for this is that manual reauthentication is designed to force the supplicant to contact the AAA server to retrieve the latest credentials. (453044)

DCS-7050X and DCS-7300 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)
- EOS will fail to detect modules, such as some 100GBASE-DWDM4 modules, that require deassertion Low Power Mode (LPMODE) pin to be de-asserted. (160016)
Series Affected: DCS-7050S and DCS-7050T Series
- 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)

DCS-7300 and DCS-7500 Series

- BFD on a statically-configured port channel may flap and traffic may be lost on a linecard insertion of which an interface is a member of the port channel.

Use LACP or disable the port-channel members on the inserting linecard before the linecard insertion (361336)
- Non per-link BFD sessions over port-channels may flap when a linecard where some of the port-channel's members reside is powered off with "no power enable". (363122)
- When a system shuts down due to a power overload, EOS does not log a 'power overload' reload cause (371160)

DCS-7020RA, DCS-7160, DCS-7170, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards 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)

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)

- 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
- 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)

vEOS Router Series

- 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 deattching a network interface. (306132)

- Site-to-site path monitored under "router path-selection" may experience flapping, when a lot of packets on the path (>10%) are dropped due to too much traffic injection into the virtual router. Workaround is to use QoS service policy to assign Inband Telemetry control packets (default UDP destination port is 4789) with a higher priority. (434087)

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)

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.