

Arista Networks EOS 4.23.0.2F Release Notes

Version: 1.0

27 November 2019

Table of Contents

EOS 4.23.0.2F Release Highlights

New Platforms and Hardware

SNMP MIBs

SNMP Traps

Supported Hardware

Reference to MLAG ISSU Upgrade/Downgrade Table

Known Software Caveats in 4.23.0.2F

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-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

Limitations and Restrictions in 4.23.0.2F

Layer 3

Multi-agent

Rib

Layer 2

Layer 1

MLAG

Accelerated System Update

CVX

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

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

Conditions and Impacts

EOS 4.23.0.2F Release Highlights

New Platforms and Hardware

- PWR-721-DC-RED

SNMP MIBs

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

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

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

SNMP Traps

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

Supported Hardware

Fixed System

- | | | |
|-------------------------|-----------------------|-----------------------|
| • DCS-7150S-24-CL-F | • DCS-7280SR-48C6-F | • DCS-7050SX-64-SSD-R |
| • DCS-7150S-24-CL-R | • DCS-7280SR-48C6-M-F | • DCS-7050TX-128-F |
| • DCS-7150S-24-CL-SSD-F | • DCS-7280SR-48C6-M-R | • DCS-7050TX-128-R |

• DCS-7150S-24-F	• DCS-7280SR-48C6-R	• DCS-7050TX-48-F
• DCS-7150S-24-R	• DCS-7280TR-48C6-R	• DCS-7050TX-48-R
• DCS-7150S-24-CL-SSD-R	• DCS-7280TR-48C6-F	• DCS-7050TX-48-SSD-F
• DCS-7150S-52-CL-F	• DCS-7280TR-48C6-M-F	• DCS-7050TX-48-SSD-R
• DCS-7150S-52-CL-R	• DCS-7280TR-48C6-M-R	• DCS-7050TX-64-F
• DCS-7150S-52-CL-SSD-F	• DCS-7280QR-C72-F	• DCS-7050TX-64-R
• DCS-7150S-52-CL-SSD-R	• DCS-7280QR-C72-M-F	• DCS-7050TX-64-SSD-F
• DCS-7150S-64-CL-F	• DCS-7280QR-C72-M-R	• DCS-7050TX-64-SSD-R
• DCS-7150S-64-CL-R	• DCS-7280QR-C72-R	• DCS-7050TX-72-SSD-F
• DCS-7150S-64-CL-SSD-F	• DCS-7280SR2-48YC6-F	• DCS-7050TX-72-SSD-R
• DCS-7150S-64-CL-SSD-R	• DCS-7280SR2-48YC6-M-F	• DCS-7050TX-96-SSD-F
• DCS-7150SC-24-CLD-F	• DCS-7280SR2-48YC6-M-R	• DCS-7050TX-96-SSD-R
• DCS-7150SC-24-CLD-R	• DCS-7280SR2-48YC6-R	• DCS-7010T-48-F
• DCS-7150SC-64-CLD-F	• DCS-7020TR-48-F	• DCS-7010T-48-R
• DCS-7150SC-64-CLD-R	• DCS-7020TR-48-R	• DCS-7050SX-72-F
• DCS-7170-64C-F	• DCS-7280QRA-C36S-F	• DCS-7050SX-72-R
• DCS-7170-64C-R	• DCS-7280QRA-C36S-R	• DCS-7050SX-72-SSD-F
• DCS-7170-32C-F	• DCS-7280SR2A-48YC6-F	• DCS-7050SX-72-SSD-R
• DCS-7170-32C-M-F	• DCS-7280SR2A-48YC6-M-F	• DCS-7050SX-96-F
• DCS-7170-32C-M-R	• DCS-7280SR2A-48YC6-M-R	• DCS-7050SX-96-R
• DCS-7170-32C-R	• DCS-7280SR2A-48YC6-R	• DCS-7050SX-96-SSD-F
• DCS-7170-64C-M-F	• DCS-7280SRA-48C6-F	• DCS-7050SX-96-SSD-R
• DCS-7170-64C-M-R	• DCS-7280SRA-48C6-M-F	• DCS-7050TX-128-SSD-F
• DCS-7170-32CD-F	• DCS-7280SRA-48C6-M-R	• DCS-7050TX-128-SSD-R
• DCS-7170-32CD-R	• DCS-7280SRA-48C6-R	• DCS-7050TX-72-F
• PWR-1100AC-F	• DCS-7280TRA-48C6-F	• DCS-7050TX-72-R
• PWR-1100AC-R	• DCS-7280TRA-48C6-M-F	• DCS-7050TX-96-F
• PWR-750AC-F	• DCS-7280TRA-48C6-M-R	• DCS-7050TX-96-R
• PWR-750AC-R	• DCS-7280TRA-48C6-R	• DCS-7010T-48-DC-F
• PWR-500-DC-F	• DCS-7020TRA-48-F	• DCS-7010T-48-DC-R
• PWR-500-DC-R	• DCS-7020TRA-48-R	• DCS-7260QX-64-F
• PWR-500AC-F	• DCS-7280CR2A-60-F	• DCS-7260QX-64-R
• PWR-500AC-R	• DCS-7280CR2K-60-F	• DCS-7260QX-64-SSD-F
• PWR-745AC-F	• DCS-7280SRAM-48C6-F	• DCS-7260QX-64-SSD-R
• PWR-745AC-R	• DCS-7280SRAM-48C6-R	• DCS-7050SX-72Q-F
• FAN-7010	• DCS-7280CR2-60-F	• DCS-7050SX-72Q-R
• PWR-1900-DC-F	• DCS-7280CR2A-30-F	• DCS-7050TX-72Q-F
• PWR-1900-DC-R	• DCS-7280CR2K-30-F	• DCS-7050TX-72Q-R
• PWR-1900AC-F	• DCS-7280QRA-C36S-M-F	• DCS-7050QX2-32S-F
• PWR-1900AC-R	• DCS-7280QRA-C36S-M-R	• DCS-7050QX2-32S-R
• PWR-747AC-F	• DCS-7280SR2K-48C6-M-F	• DCS-7050SX2-128-F
• PWR-747AC-R	• DCS-7280SR2K-48C6-M-R	• DCS-7050SX2-128-R
• FAN-7002H-R	• DCS-7020SR-24C2-F	• DCS-7050TX2-128-F
• FAN-7000H-R	• DCS-7020SR-24C2-R	• DCS-7050TX2-128-R
• PWR-1600AC-F	• DCS-7280CR2M-30-F	• DCS-7050SX2-72Q-F
• FAN-7001D-F	• DCS-7280SRM-40CX2-F	• DCS-7050SX2-72Q-R
• FAN-7001C-F	• DCS-7280SRM-40CX2-R	• DCS-7060CX-32S-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
- 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-7280QR-C36-F
- DCS-7280QR-C36-M-F
- DCS-7280QR-C36-M-R
- DCS-7280QR-C36-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-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
- DCS-7050SX-64-F
- DCS-7050SX-64-R
- DCS-7050SX-64-SSD-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
- CCS-720XP-24Y6-F
- CCS-720XP-48Y6-F
- DCS-7050CX3M-32S-F
- DCS-7050CX3M-32S-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-F
- FAN-7002-R
- FAN-7002H-F
- FAN-7000H-F

Modular

- FAN-7002-F
- FAN-7002-R
- FAN-7002H-F
- FAN-7000H-F
- PWR-3K-AC-F
- PWR-3K-AC-R
- PWR-2700-DC-F
- 7500E-6CFPX-LC
- DCS-7500-SUP2
- DCS-7500-SUP2-D
- 7500R-36CQ-LC
- 7500R-36Q-LC
- 7500R-48S2CQ-LC
- 7504R-FM
- 7500R3K-36CQ-LC
- DCS-7508
- DCS-7504
- 7300-SUP
- 7300-SUP-D
- 7300X-32Q-LC
- 7300X-64S-LC

- PWR-2700-DC-R
- PWR-3K-DC-BLUE
- PWR-3K-DC-RED
- PWR-3KT-AC-BLUE
- PWR-3KT-AC-RED
- 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
- 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
- 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

Transceiver

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

- 40GBASE-SRBD
- 100GBASE-AOC
- CFP2-100GBASE-XSR10
- AOC-D-D-400G
- AOC-O-O-400G
- 10GBASE-ER
- 40GBASE-CR4

Reference to MLAG ISSU Upgrade/Downgrade Table

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

Known Software Caveats in 4.23.0.2F

Layer 3

- If a loopback interface is only configured with a link-local IPv6 address, OSPFv3 will not advertise the address. (131369)
- When using OSPFv3 and encryption, the Rib agent may stop responding and restart unexpectedly. (206494)
- 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)
- 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)

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

- In a EVPN VXLAN IRB setup, when service leafs (NELs) are used to interconnect data-centers (DCs), the ARP entry for a remote host (e.g. hostA) on a local DC may get erroneously removed on the service LEAF. Once this occurs, the service LEAF will not proxy reply to any further ARP requests for hostA from remote DCs and the ARP for hostA will not get resolved in remote DCs. (387954)
- 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)
- Setting `ipv6 icmp rate-limit-unreachable 0` and shutting down all the next hop interfaces for an IPv6 default route resilient ECMP causes the L3 Agent to restart. (395322)
- When an ARP entry is added while there is a matching, existing virtual IP with a non-/32 subnet, traffic loss may occur because traffic may be flooded instead of forwarded to the CPU. (402802)
- When multiple Connected Routes are deleted from the FIB that are Leaked across VRFs, the Attached Route in the Leaked-To VRFs remains as a DROP route without any ARP resolution.

The presence of the DROP route may prevent learning of the ARP if the Leaked-To VRF is required to trap the packet to CPU for the ARP resolution in the source VRF.

A workaround is to restart the Arp agent which will remove the DROP attached route from the Leaked-To VRF. (406752)

- If the number of VRFs is scaled above approximately 500-700, the VrfLeak agent may exhibit high CPU and stop processing VrfLeak policies. To address this, the VRF scale must be reduced. (407444)
- If V6 Attached Routes are not deleted by the L3 Forwarding Agent before the VRF deletion, then the Attached FEC doesn't get cleaned up. If the same Attached Route gets recreated in same or different VRF this Attached Route remains as a DROP route.

Work around is to restart L3 Forwarding Agent to delete the DROP route. (409067)

- OSPF graceful restart fails when multiple instances are configured and a hitless reload of the switch is performed. (409907)
- Large numbers of RSVP node protection enabled RSVP tunnels sharing the same outgoing interface can cause the layer 3 forwarding agent to transiently crash after a link down occurs on that interface if hardware fast failover is enabled as well. (409973)
- On an IS-IS router that has node protection with TI-LFA enabled, routing topology changes can cause the Mpls agent to restart unexpectedly (410606)
- Isis may restart if RSVP-TE is configured for Fast Reroute after it has formed adjacencies and received network LSPs. (412169)
- When TiLFA fast reroute is configured with local convergence delay, then Rib agent at Point of Local repair may restart when there is SPF scheduled to run for reason other than failure of the protected link , followed by link down of the protected link. (412231)
- When OSPF multiple instances are configured with unsupported area interface configuration, and then the multiple instances are deleted such that there is a single instance left, the OSPF agent may restart when OSPF interface configuration changes are applied. (413364)
- 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)
- A large configuration change or a forwarding plane reset may result in BFD sessions not coming up.
A workaround is to restart the Bfd agent. (414250)
- Removing SVI configuration at scale via a config-replace operation may cause the Bfd agent to crash. (414702)
- Receiving a trigger Resv message for an RSVP-TE bypass tunnel that is in active use may cause a brief drop in traffic forwarding over the bypass (415287)
- 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
- 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)
- Traffic will be dropped if it hits MPLS route which has ECMP nexthops with mixed swap and pop action when payload type is set to autoDecide. (417883)

- When "flow tracking hardware" is enabled, certain values of inactive timeout configured may result in continuous restart of StftAgent. The workaround is to configure inactive timeout between 5502 and 900000 milliseconds. (418692)
Series Affected: CCS-720XP Series
- L3 agent crashes and restarts continuously when more than 64 ECMP members are programmed in a group. (420390)
Series Affected: DCS-7260CX3 Series
- 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)
- 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)
- 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)

Multi-agent

- There may be some traffic loss on ASU or SSO restart at high scale (i.e with a large number of next hops). (362995)
- Performing a stateful swtichover in a switch with a large number of VRFs configured, with BGP config in all VRFs and IGPs configured in a large number of VRFs, may cause the system to restart unexpectedly, leading to traffic loss. (376773)
- The ConfigAgent may restart when there is churn in a large number of route next hops (394148)

- When deleting and re-adding a VRF in quick succession, the IpRib agent may program some IGP routes in the re-added VRF as drop routes (397760)
- During ASU if L2 EVPN BGP updates and withdrawals are being received continuously by the restarting speaker, then L2RibReconcile stage might timeout causing the device to reboot. The workaround is to avoid such continuous churn in the network during ASU. (397926)
- The "Not best" output in "show ip bgp <prefix> detail" and "show ipv6 bgp <prefix> detail" does not consider the IGP cost step in BGP best path selection. These commands will claim that affected paths are "Not best" due to some comparison performed after the IGP cost step. This is specific to the show command output; best path selection itself handles IGP cost correctly. (401354)
- When a BGP peer configuration modification results in a RibOut change, if while the peer is being integrated in its new RibOut, a route changes bestpath and the new bestpath is received from the same peer, then the expected withdrawal for the route may not be issued towards the peer. (409832)
- With BGP configured in VRF, traffic loss may be seen in some VRFs on "reload fast-boot" due to Forwarding state preserved(F) bit not being set in the open message (414457)
Series Affected: DCS-7280SR and DCS-7280TR Series
- Device reboot or Bgp agent restart can cause BgpCliHelper or Bmp agents to restart unexpectedly. A BgpCliHelper agent or Bmp agent restart may unexpectedly result in another restart. (419002)
- Extended traffic loss may occur during switchover on devices which have high IP route scale. (420564)
- 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)
- Bgp agent can restart under certain circumstances when "set metric igp-nexthop-cost" is configured. (421236)
- When OSPF graceful restart is configured in multiple VRFs in multi-agent mode and routes are redistributed into OSPF and SSU or SSO is executed, OSPF may withdraw 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)
- 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)

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 (standered, extended, large communities etc.) are included in the
outbound
route advertisements. (384629)

- The command 'show bgp ipv4|ipv6 unicast <prefix> detail' causes the Bgp agent to restart. Use 'show ip|ipv6 bgp <prefix> detail' instead. (394413)
- This issue can intermittently occur when we have an ECMP route to the
lending loopback of an adjacent neighbor over unnumbered interfaces. If

links bounce there is a race condition that may cause the arp entry for the neighbor on the affected link to be installed in the kernel by not be used by the rest of the system. If it occurs, a workaround is to remove the arp entry from the kernel... the next attempt to resolve the arp should succeed. (412349)

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)
- 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)
- Authentication failure VLAN for Dot1x does not work post ASU. Authentication failure VLAN needs to be removed and configured again post ASU for supplicants to be assigned under it. (397282)
- Dot1x agent may crash during ASU if some tagged supplicant has been authenticated on the dot1x enabled port without dynamic VLAN assignment. (397693)
- Dot1x agent may crash post ASU because of a change in the format of Calling-Station-Id to a standards compliant format. (410569)

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

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)
- During config-replace, MLAG port-channels may end up transmitting using members of the peer-link, misdirecting or blackholing traffic (405166)
Series Affected: CCS-720XP, DCS-7010T, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X2, DCS-7250X, DCS-7260X, DCS-7260X3 and DCS-7300 Series

- In a deployment with VXLAN routing on a MLAG pair, the VxlanSwFwd agent may experience high CPU utilization when the Vxlan interface is administrative shutdown.

The workaround is to re-enable the Vxlan interface or run the following on both MLAG peers:

```
$ bash python -m Acons VxlanSwFwd
$ cd /ar/VxlanSwFwd/VxlanSwFwd/virtualArpSm/tacVtiStatusSm/Vxlan1/
arpSyncUdpPam
$ __.mode = 'disabled'
```

If the Vxlan interface is later re-enabled, the VxlanSwFwd agent must then be restarted. (419220)

Accelerated System Update

- Performing SSU when /mnt/flash/persist/persistentRestartLog doesn't exist may result in extended traffic outage due to a watchdog reset that can occur. (158117)
- When 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
- Performing SSU with an MTU value less than 1280 on an IPv6 interface may result in a fatal error with extended traffic outage. (400408)
- When performing SSU on an MLAG setup, SSU can cause peer links on port-channel to flap, resulting in large traffic loss. (405788)
- Broadcast and unknown unicast Vxlan traffic may face an extended traffic loss of 20 seconds after ASU (409876)
 SKUs Affected: DCS-7280SR-48C6-F, DCS-7280SR-48C6-M-F, DCS-7280SR-48C6-M-R, DCS-7280SR-48C6-R, DCS-7280TR-48C6-F, DCS-7280TR-48C6-M-F, DCS-7280TR-48C6-M-R and DCS-7280TR-48C6-R
- Traffic loss of 1-2 seconds may be experienced on IS-IS Graceful Restart in scale scenarios following SSO or ASU (414174)
- 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
- With BFD scale, ASU may time out waiting for the Bfd agent (430963)

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)
- Restarting the OpenStack or Controllerdb agent may cause OpenStack dynamic configuration to flap on switches (181409)
- 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 a 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)
- The TopologyAggregator agent may fail on restart when an object of different value is given an already existing key. Workaround is to clear CVX state completely by going to CVX mode and shut/no shut. (393741)
- 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)

- 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)

MPLS

- The RSVP agent crashing or restarting could cause the Layer 3 platform agent to crash. (404206)
- If Rsvp agent is restarted by using 'kill' bash command (instead of Arista CLI commands, which is a recommended way of restarting agents), Rsvp agent may crash with assertion violation about 20 seconds after restart. (404207)
- LDP Pseudowires which are resolved via ISIS-SR tunnels may drop traffic on switches with only Jericho or Jericho+ linecards (406211)
- The Rsvp agent may restart unexpectedly when a large number of LSPs are using the same node protection bypass tunnel and that tunnel gets deleted. (415301)
- When EOS routers are deployed as both LSR and tail end, an RSVP primary LSP going through the EOS LSR may flap if its associated bypass LSP originating at that LSR flaps even when the bypass LSP is not in use. (417161)

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

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

- The CLI command "route-target import auto <ASNumber>" for VLAN VRF does not work and hence should not be used. Alternatively, the suggested configuration is "route-target import auto". In this case the AS Number is picked up from the BGP configuration.

Example:

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

Suggested Usage

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

- 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
ffffffffffffffff
[ 286.436777] IP: [<ffffffff8130fa93>] pci_free_cap_save_buffers+0x16/0x29
```

or

```
[ 336.206463] BUG: unable to handle kernel NULL pointer dereference at
0000000000000011
[ 336.300062] IP: [<ffffffff8130c63c>] _pci_find_saved_cap+0x18/0x31
(262388)
```

- 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)
- 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)
- OpenConfig affected by CVE-2019-6486 (356552)
- 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)
- 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)
- 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)

- PTP management packet handling in PTP boundary clock mode may result in a broadcast storm if the PTP topology has a timing loop in it.

A workaround is to use a PTP service ACL to block PTP management messages from looping between boundary clocks. (397865)

- The DHCP IPv4 Server can enter a deadlock state when show commands and configuration commands are run simultaneously. To clear the deadlock, execute `sudo kill -9 <pid>` where <pid> is the process ID of the Kea DHCP Server process. (399832)
- When the forwarding agent hitfully starts up, PTP packets sent downstream may have invalid timestamp and the switch may incorrectly advertise itself as the GM. Valid PTP packets and correct GM advertisement will converge post forwarding agent startup. (402786)
- In NETCONF output, YANG leafrefs to identityrefs are encoded incorrectly, with <Module> and <Name> elements. (402952)
- 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)
- In high scale situations, during an agent restart, mac table clear, or when learning a large a number of MACs, IpLocking may send too many lease queries to the DHCP server that the server may start dropping packets if the server load is too high. In this case, hosts with legitimate DHCP leases may be temporarily blocked until load on the DHCP server decreases. Workaround is for the affected host(s) to send an IP or ARP packet to the switch after a couple minutes. (409386)
- Pim SSM IPv4 Non-DR OIF installation feature may stop working resulting in duplicate multicast traffic when Ipv6 multicast routing is enabled. (410241)

- When more than 255 subnets are configured for DHCP Server, the agent will continuously restart. To correct the DHCP Server error state, the configured number of subnets must be changed to 255 or less and the DHCP Server ipv4SubnetLastUsedId state in Sysdb must be manually cleared. (411213)
- When "channel-group" command is applied through eAPI on an Ethernet interface, its switch port configuration may override the switch port configuration of the port-channel that the Ethernet interface is configured in.

Workaround is to add the required switchport configuration on all member interfaces of the port-channel. (412920)

- 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)
- A malicious gNMI client may cause significantly increased memory usage in the OpenConfig or Octa processes, resulting in a denial of service of gNMI, RESTCONF, and NETCONF access. (416066)
- DirectFlow 'action output nexthop <ip addr>' ignores TTL of incoming packet and forwards the packet even if the TTL has expired. (417994)
- OpenConfig repeatedly restarts on startup in cEOS-lab with the error "panic: Connection to @/Arista/sysdb closed: are we not configured as a daemon in the CLI? Or did Sysdb just die? Either way, check the contents of /var/log/agents/Sysdb-*". The Sysdb logs will contain a line similar to "Connection from <> to OpenConfig at <> closed locally (Processed an invalid mount profile. Component p4runtime not found in /ar/Sysdb)". (418027)
- System may unexpectedly reboot with an Unknown reload cause due to Uncorrectable ECC errors. (419257)

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-7020SRG-24C2-F, DCS-7020SRG-24C2-R, DCS-7020TR-48-F, DCS-7020TR-48-R, DCS-7020TRA-48-F, DCS-7020TRA-48-R, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050QX-32S-F, DCS-7050QX-32S-R, DCS-7050QX-32S-SSD-F, DCS-7050QX-32S-SSD-R, DCS-7050QX2-32S-F, DCS-7050QX2-32S-R, DCS-7050SX-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-7150SC-24-CLD-F, DCS-7150SC-24-CLD-R, DCS-7150SC-64-CLD-F, DCS-7150SC-64-CLD-R, DCS-7160-32CQ-F, DCS-7160-32CQ-R, DCS-7160-48TC6-F, DCS-7160-48TC6-R, DCS-7160-48YC6-F, DCS-7160-48YC6-R, DCS-7260QX-64-F, DCS-7260QX-64-R, 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-7280SE-72-F, DCS-7280SE-72-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 and DCS-7280TRA-48C6-R

- 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)
- OpenConfig may consume excessive amounts of memory when handling a gNMI Set RPC. (422344)
- 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)
- In OpenConfig, some VRF configuration can be missing from the YANG tree if hundreds of VRFs are configured. (426960)
- 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)

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)
- In an environment with a very large number of MSDP peers, during cleanup the Msdp Agent may crash. (354364)
- If large number of IGMP reports are received in a very short interval, some of them might get dropped resulting in the multicast route not being created. The routes will be re-learned on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIgmp", consider increasing the bandwidth for this Copp class. (356675)
- 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)
- When unicast routing is enabled and ip multicast static is configured on any interface, L2 multicast forwarding stops working as expected. Workaround is to enable multicast routing and configure PIM on a dummy interface in the affected VRF. (372312)
Series Affected: DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series
- 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)
- 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)
- 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)
- If a VRF is removed while it has active multicast routes, the Multicast Agent can restart unexpectedly. (418028)
- 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 queires are received in quick succession. This will result in traffic loss. (422416)

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)
- On the DCS-7500, DCS-7500E, DCS-7280E, DCS-7500R, and DCS-7280R series, in a switch configured as an MLAG peer, restart of the SandAcl agent can lead to unexpected packet behavior.
Workaround is to restart the SandAcl agent again. (175591)
- 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)
- When IP multicast routing is enabled, the switch might learn invalid MAC addresses from IPv6 multicast traffic based on the source IPv6 address in the packet. (226115)

Series Affected: DCS-7020, DCS-7280R2 and DCS-7500R2 Series

- 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 deployment where the VTEP does not receive any bridged traffic from the host. In such a scenario, the traffic toward the host will be flooded in the VLAN.

The work around for this is to force Neighbor Solicitation (NS) packets from the host for its gateway or send bridging data traffic from the host. (246676)

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

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

- The platform Slice agent may restart unexpectedly when sflow hardware acceleration is enabled or disabled. (247895)
- 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)

- 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)
- When both port-channel load balance for multicast and selective ingress replication is configured, groups that get ingress replicated incorrectly replicate packet to each member of port-channels that are member of the group.

As a workaround one of the features has to be disabled. (344650)

- With flowspec being enabled on a port-channel interface, a switch reload may result in some port-channel members not enabled for flowspec. The workaround is to reenable flowspec on the port-channel interface. Additionally, restarting SandAcl and other forwarding agents at the same time will cause unexpected SandAcl agent restarts. SandAcl agent will recover after the restarts. (349984)

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

- On DCS-7500RM-36CQ-LC linecards, rapidly enabling and disabling FEC on interfaces with 100GBASE-PSM4 transceivers can cause the interfaces to not come up, flap, or errdisable.

Workaround is to wait 20 seconds before applying another FEC configuration. (354921)

- 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)
- After SandFap slice agent restarts with PTP operational and using peer-to-peer delay mechanism, all PTP Peer delay messages ingressing PTP enabled ports will no longer reach the PTP agent. The workaround is to toggle PTP mode from disabled back to boundary. (370878)
- 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)
- Tap aggregation traffic steering policies which have key sizes of 80 bits are not installed into hardware. (382340)
- 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)

- The Type Of Service (TOS) field in MPLS-terminated IP packets may become corrupted when using a user-defined TCAM profile.

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

- With hardware accelerated sflow, the switch sometimes stops sampling traffic (394965)
- Traffic subject to CPU traffic policy `_policer_` action will not be `_forwarded_`. This includes multicast control traffic which is replicated to both the CPU but also forwarded (e.g. PIM, OSPF HELLOs, etc.). Unroutable unicast traffic is also flooded, and so will not be forwarded if matched by a CPU traffic policy policer action. An example policy of concern is

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

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

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

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

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

- On PTP enabled switches, CRC errors may occur in internal time synchronization that can cause PTP time to be unstable. PTP time will recover stability on its own after a few seconds. (399179)

- If the DirectFlow `priority` keyword is used with more than 1024 unique values, the SandAcl agent may restart causing all flows to be reprogrammed. (400125)
- In the output from "show plat fap directflow tcam" the "priority" field displayed is 1 higher than it should be. (400733)
- Insertion of a fabric module may cause brief traffic loss on the device (402508)
SKUs Affected: 7504R3-FM and 7508R3-FM
- Adding new VLAN members to a CPU traffic policy using the police action may result in policers not being allocated properly for those members. The workaround is to deconfigure and reconfigure the policy. (402804)
Series Affected: DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- 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)
- 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)
- On a linecard insertion, BGP Flowspec rules may fail to get installed for physical interfaces on the new linecard. (405109)
- 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)
- BGP Flowspec rules may not get installed correctly on an interface after recovering from error conditions. Error conditions include insufficient TCAM to program all the matches sent from BGP neighbors or matches specified with unsupported qualifiers or actions. (405932)
- When a traffic-policy with non-default actions is modified in a config-session the policy action stops working. (406486)

- PolicyMap configuration for QoS or IPv4 Traffic Steering in Tap Aggregation mode may not take effect after TCAM profile switching or Stateful Switch Over.

Remove the affected QoS or Traffic Steering policy configuration from interfaces and then reconfigure them. (407016)

- Host routes may not be installed in hardware when uRPF is enabled with "ip hardware fib optimize prefix-length 32 urpf" configuration. The workaround is to disable uRPF. (408073)
- When CPU Traffic policy is enabled and a Layer3 port-channel interface is configured in the system, a switch reload may result in some port-channel members not enabled for CPU traffic policy. The workaround is to reconfigure the policy or the interfaces. Additionally, restarting SandAcl and other forwarding agents at the same time will cause unexpected SandAcl agent restarts. SandAcl agent will recover after the restarts. (408220)
- With scaled Acl configurations, SandAcl agent restart may result in additional agent restarts. (408260)
- When configuring a LAG member it may get stuck programming RX as indicated by 'waiting for hardware to program port to RX (or receive)' in the output of 'show port-channel'

Unconfigure and reconfigure the interface as a LAG member (408448)

- When the Tc-Counters pmf profile is applied and packet-type mpls-over-gre outer-ip is enabled, bridged mpls-over-gre packets will be dropped (408835)
- 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)
- BGP Flowspec matches on LAG routed interfaces may fail to be installed on interfaces associated with a linecard OIR occurred or on member interfaces that were previously in a down state. (409795)
- TCAM reservations are not preserved after a system reboot. (410069)
- 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)
- 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)

- When new BGP flowspec rules are advertised, some of the existing installed flowspec rules may become uninstalled. This happens only when a new TCAM bank is allocated for flowspec usage. (412425)
- When performing speed change on interfaces which are LACP LAG members, some members may get stuck "Waiting for LACP response" as seen in the output of "show port-channel".

Restart SandLag agent (412728)

- On a linecard insertion, BGP Flowspec rules may fail to get installed for physical interfaces on the new linecard. (412802)
- Frequent BGP Flowspec rule churn may cause out of memory errors for the SandAcl agent, causing it to restart. (412931)
- When applying a CPU Traffic Policy to a VRF with a VLAN with members across multiple FAPs, the police action may not be programmed correctly if the VRF contains other interfaces on the same FAPs. (414185)
- 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)
- With PIM enabled for an SVI of a VXLAN VLAN, on hitful restart of SandAcl agent, PIM packets might be incorrectly flooded to remote VTEPs. Shut/no shut of SVI would resolve the issue. (415525)
- A switch reload with qos policy configured may result in SandAcl agent restart. (417935)
- 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)
- Rollback from an unsuccessfully-applied CPU traffic-policy will result in no CPU traffic-policy being applied, rather than reverting to any previous traffic-policy configuration. This only affects the 'cpu traffic-policy <> vrf all' command. (421391)

- 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)
- Static MACs that are removed from config when SandMact is down, are not cleaned up in hardware.
Workaround is to re-add the MACs and to delete them once SandMact is up. (422379)
- 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)
- 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

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
- With FlexRoute, and update-wait-install configuration, IPv4 BGP routes might not be advertised to a peer even after they are installed in hardware. (393706)

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)
- 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)

- The TOS field of VRF-mapped IP over MPLS packets may become corrupted when using the pbr-match-nexthop-group TCAM profile.

Remove the packet type 'mpls ipv4 forwarding routed decap' from all applicable features in the TCAM profile. More specifically, create a new TCAM profile based on pbr-match-nexthop-group and remove that packet type from the features 'acl vlan ip', 'acl port mac', 'pbr ip' and 'qos ip'. Alternatively, using the TCAM hardware rather than AlgoMatch, if applicable, may also avoid this issue. This can be done by running `hardware access-list mechanism tcam`. (412895)

DCS-7280R3 and DCS-7500R3 Series

- During restart of SandAcl agent, packets may not be subjected to ACL rules for a short period of time. (372971)
- Neighboring interfaces may flap if speeds or speed groups are configured on an interface. (394764)

- Configuring IPv4 and IPv6 rules or configuring mpls match rules on a PBR policy may cause the SandAcl agent to restart. (407070)
- Shared Egress Ipv4 RAcls can leak packets if the mac of the nexthop is not learnt. (412762)
- Peer link flap can cause interface to remain down. Workaround is to shut/no shut the interface. (413908)
- When traffic is running, speed change and speed group configuration can cause the forwarding agent to restart.

Workaround is to shut Ethernet interfaces before configuring speeds and shut all Ethernet interfaces of a speed group before configuring speed groups, no shut-ing them after speed/speed group configuration. (414432)

- Decap groups incorrectly uses uniform mode for QoS. (417136)
- When link down occurs on a 10G/40G interface, the other 10G/40G interfaces within the same speed group may flap. (418622)
- The traffic class in IPv6 packets egressing the switch after MPLS label termination might be incorrect. (420345)
- When changing the speed group of an interface with auto-negotiation configured to 10g, the forwarding agent can restart.

Workaround is to disable auto-negotiation, then configure the speed group to 10g. (420599)

- IPv6 over MPLS packets received with IPv4 Explicit NULL (i.e. Label=0) are misforwarded. (425904)
- PIM-SM v4 does not work on an mlag with IgmppSnooping enabled. (427859)

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

- 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)
- 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

- 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

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

- The forwarding agent may restart when mac address aging is running. The workaround is to disable the sram scanning in the forwarding agent. (165685)
Series Affected: DCS-7050X2 Series

- Some of the MACs learnt over remote VXLAN VTEPs go missing in hardware and are never relearned when a member in peer-link port channel is shut then unshut. (184732)
- 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

- The OpenFlow agent may fail to send a packet to a controller if the table profile in use is 'auto' and there is a lower priority flow that matches the traffic and drops it. To workaround this, use a different table profile. (217965)
- 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)

- 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

- Control-plane traffic on routed ports may experience traffic loss during switchover. (247668)

Series Affected: DCS-7300X 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)

- SW VXLAN forwarding does not work with selective Dot1q tunneling configuration. This may result in tunneled ARP packets getting mis forwarded to/from selective dot1q enabled ports with inner tags stripped. (273728)
- 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
- Shutting/Unshutting a MLAG downlink might cause the unknown unicast flooding on MLAG peer-link.
A possible workaround is to flush the MAC entries in the system (288719)
- 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: DCS-7050X3, DCS-7260X and DCS-7260X3 Series
- Some ports (esp. Et73-Et96 and Et103-Et104) might blackhole traffic or fail to linkup. (358361)
SKUs Affected: DCS-7050TX-128-F, DCS-7050TX-128-R, DCS-7050TX-128-SSD-F and DCS-7050TX-128-SSD-R
- Enabling EVPN MPLS may result in continuous StrataL2 forwarding agent restarts as it is not supported. (363491)
- Irrespective of IGMP snooping being enabled or disabled, IGMP packets will not be flooded to remote VTEPs. (368106)
Series Affected: 7300X3, DCS-7050X2 and DCS-7050X3 Series
- Executing "platform trident forwarding-table partition" in quick succession may result in unrecoverable traffic loss. The work around is to do remove

and reissue the command "platform trident forwarding-table partition".
(368260)

- Performing SSU may result in links remaining down, flapping or cold boot if any transceivers that do not support the default interface speed and no speed configuration are present on the interface. The default interface speed for QSFP28 interfaces is 100G, for QSFP+ interfaces it is 10G, for SFP28 interfaces it is 25G, and for SFP+ interfaces it is 10G.

A workaround is to explicitly configure the required speed on all the ports. (392658)

- When L2-EVPN is configured with VxLAN transport, if the StrataL2 agent restarts, it might result in some VTEPs missing from the flood list for some of the VxLAN VLANs. As a result, BUM traffic in those VxLAN VLANs will not reach those missing VTEPs. Workaround is to flap the VxLAN interface (vxlan1) using the shut and no-shut CLI commands. (400418)
- 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)
- 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)
- The switch may not reboot after an internal power failure, resulting in traffic loss until the switch is manually power cycled (405573)
SKUs Affected: 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 and DCS-7050TX-72-SSD-R
- When Vxlan is configured, doing configure replace again with vxlan config may result in dropping all Vxlan traffic. Workaround is to flap the Vxlan interface. (406674)
- 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
- When "hardware counter feature vni encap/decap" feature is enabled, SSO results in continuous StrataL2 crashes eventually resulting in full switch reload.
The workaround is to not enable the "hardware counter feature vni encap/decap" feature when sso is enabled. (408854)
- Enabling "hardware counter feature vtep encap/decap" may result in continuous StrataCounter agent crashes on stateful switchover.
If the switch gets into this condition, the workaround is to disable the "hardware counter feature vtep encap/decap". (410030)
- 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
- DirectFlow flows can be uninstalled from a linecard on reaching the idle timeout for the flow even though matching traffic continues to match the flow on another linecard. The flow can be restored by re-programming it or disabling the idle timeout.

Series Affected: 7300X3 (411851)
Series Affected: 7300X3, 7320X, 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)
- With IPv6 multicast configured, changing the speed of a port will lead to disruption of IPv6 multicast traffic (418250)
Series Affected: DCS-7260X3 Series
SKUs Affected: DCS-7050SX3-48YC8-F and DCS-7050SX3-48YC8-R
- Configuring more than 16K MPLS pop/swap labels will cause continuous restart of forwarding agent. (419243)

- Performing SSU with static NAT configuration may lead to fatal error reboot which will cause increased traffic outage. (419459)

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

- 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

- 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: 7320X, DCS-7250X, DCS-7260CX and DCS-7300 Series

- 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

DCS-7050X2 and DCS-7300 Series

- 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)
- 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)
- 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)

- Interfaces configured with Clause 73 autoneg might flap or fail to link up following an SSU (reload fast-boot) reload.

Workaround is to restart the forwarding agent hitfully using 'platform trident reset' command in enable mode. This will cause flaps and brief traffic disruption on all front panel ports. (409041)

Series Affected: DCS-7260CX3 Series

- After some speed-change operations, control-plane packets larger than 4466 bytes sent towards CPU could be discarded. (430651)

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)
- 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)
- Gratuitous ARP packets sent from CPU with source MAC as VARP Mac address may get encapsulated twice to remote VTEPs with SIP as local VTEP Ip and VARP Vtep Ip resulting in mac address move for the VARP mac address in remote VTEPs. This may result in momentary traffic loss. The workaround is to increase the GARP advertisement interval with command "ip virtual router mac advertisement-interval <interval>" to reduce the frequency of MAC address moves. This issue only affects configurations that contain ip virtual-router address. (289986)
- DirectFlow flow with L2 rewrites configured may not perform VLAN/SMAC/DMAC modification after port speed change. Toggling affected flow or modifying action's interface list should restore expected behavior. (390563)

Series Affected: DCS-7050SX3 Series

- MPLS IPv4/IPv6 ping is not supported (394778)
- When speed-change is done on an interface that is receiving PFC frames, then the effect of PFC frames is seen even if the PFC frames are no longer being received. In this condition, packets cannot be transmitted on some of the egress unicast and multicast queues. A workaround is to shutdown the interface before doing the speed-change. (397641)
- When performing SSU, there may be failure in table programming resulting in fatal error cold reboot with extended traffic outage. (399774)
- When performing SSU using "reload fast-boot", a delayed confirmation to proceed with the reload may cause a fatal error reboot resulting in extended traffic outage. The workaround is to use the "reload fast-boot now" command. (400649)
- When 'phone trunk untagged' and 'phone vlan' are configured, DHCP snooping is not supported on the VLAN specified as the phone VLAN. (408169)
- When a switchport is configured in "phone trunk untagged" mode, IGMP snooping does not work if the VLAN assigned based on source MAC is different from the port's native VLAN. (408180)
- Adding interfaces to a port-channel that is an output interface on a flow with VLAN/SMAC/DMAC modification may cause traffic to egress that port-channel without performing the modification. This can be resolved by toggling or otherwise modifying affected flows. (409975)
Series Affected: 7300X3 Series
- After performing ASU reload, OSPF sessions may not come up, resulting in complete traffic drop. (410242)
- QSFP28 interfaces with optical transceivers can experience FCS errors, CRC errors or link flaps. (410660)
SKUs Affected: DCS-7050SX3-48YC12-F
- 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
- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (416694)

- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (417223)
- When Evpn multihoming is configured 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)
- When Evpn multihoming is configured, BUM traffic to some Vxlan VLANs may get dropped and copied to cpu if the multicast index allocated is greater than 8191. The workaround is to restart L3 forwarding agent. (419207)
- When "platform trident forwarding-table flexible" is configured, L3 unicast forwarding to hosts / L3 multicast forwarding may not work resulting in unexpected behavior. The workaround is to remove the CLI and use "platform trident forwarding-table partition" instead. (421084)
- Performing a reload can result in some host routes being unprogrammed. (427790)
 Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX3 and DCS-7050X3 Series

7368 Series

- An interface may fail to link with a MAC Local Fault indication.
 The workaround is to run interface configuration command "shutdown" followed by "no shutdown". (341316)
 Series Affected: 7368 Series
- 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)
- Changing speed-group serdes configuration for multiple groups in a short time to 25G and 10G may cause Strata forwarding plane agent to restart resulting in link flaps and traffic loss. (409647)
 Series Affected: DCS-7060X4 Series
- The switch card may experience internal power failure. (427251)
 Series Affected: 7368 Series
 SKUs Affected: 7368X4-SC

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 the event of an SEU on the system FPGA, a power cycle may not correct the SEU if the system is configured to maintain PoE power during reboot.

The workaround is to configure the system to not maintain PoE power during a reboot with 'reboot action power-off' and then reboot the system. (402836)

Series Affected: CCS-720XP Series

- Speed changes resulting in a re-allocation of hardware resources may result in unexpected traffic loss. A forwarding agent restart is required to recover, which will result in a temporary traffic outage on all interfaces on the switch. (409773)
SKUs Affected: CCS-720XP-48ZC2-F
- When "flow tracking hardware" is enabled, some flow records may not get exported to collector. (418105)
Series Affected: CCS-720XP Series

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)
- 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)
- 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.

As a workaround, static ARP entry for the affected intercepted hosts can be added at the Service VTEP. (394247)

- (A1 silicon only) A single event upset (SEU) may cause Strata forwarding agent to restart unexpectedly. Traffic may be temporarily impacted. (412637)

DCS-7300 and DCS-7500 Series

- When many IPv6 routes are configured and the BFD protocol is 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)
- In L3 EVPN network, the IpRib agent may keep restarting after SSO switchover on the new Active node. (415096)

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)
- 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)

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)

Limitations and Restrictions in 4.23.0.2F

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, DCS-7060X2 and DCS-7060X3 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)

Multi-agent

- If a BGP peer is configured for BFD fallover, the BFD session is not cleaned up even after the BGP peer is unconfigured. (217199)
- Traffic destined to an L3EVPN route whose underlay has ECMP BGP-LU routes, may experience temporary traffic loss when the underlay goes from n way to n-1 way or lower. The duration of loss can vary depending on the scale of BGP-LU routes or L3EVPN routes. (251692)
- When the multi-agent routing protocol model is configured, the maximum number of next hops for an OSPF ECMP route is 128. (274888)
- With very a very low multi-hop BFD timer configured, a BFD session may flap when a path (belonging to an ECMP set), over which the BFD session is established, goes down. (287571)
- BGP session over indirectly connected interface that was down due to lack of IGP path to peering address may take around one minute to re-establish once the IGP path is available. (287628)
- Route-map specified as match-map for dynamic prefix-list is evaluated only against BGP routes. (345444)
- The IpRib agent runs out of addressable memory when a large number of routes are leaked to a large number of VRFs. (346575)

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

- 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)
- 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):

```
flow bypassPeerLink
    priority 65535
    table trident ifp
```

```
match input interface Po1
match ethertype ip (375185)
```

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

```
flow bypassFwIntf3
  priority 65535
  table trident ifp
  match input interface Po1103    >>>>>>> lag interface connected to FW
  match ethertype ip
!
flow bypassFwIntf4
  priority 65535
  table trident ifp
  match input interface Et31      >>>>>>> phy interface connected to FW
  match ethertype ip
! (375189)
```

- Add support for Port-Channel interfaces for ContainerTracer. (429898)

General

- CLI does not allow access to files whose names contain spaces. (539)
- 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)
- Copying boot-config containing a SHA-512 boot secret to switches with Aboot 2.0.7 or older will reject all boot passwords upon reload and the password recovery procedure will be required. (50695)
- Do not create an alias named "alias". (65392)
- When VmTracer is used in a network that includes Cisco devices, ESX hosts connected to the Arista switch will see CDP frames from the Arista switch as well as from the other Cisco devices. VmTracer will display the VM info when the ESX host reports Arista's CDP info, but then will delete it when the ESX host reports other CDP info.

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

- EOS SDK-based applications must be run via the 'daemon' command configuration. They cannot be run directly from bash. (158431)

- EOS extension scripts and agents which do not use EOS-SDK will need to be modified to work in 4.16.6M and beyond. (158467)
- 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)
- The boot config command takes effect immediately inside a config session and cannot be aborted (346171)

- 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. (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)
- Adding or Removing privilege level to an existing group in a LDAP group policy results in duplicate entries appearing in cli. (422323)
- 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)

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)

DCS-7150 Series

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

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

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

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

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

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

- Systems with A0 silicon, packets are not forwarded to the egress port when egress truncation and timestamping are enabled together. Removing the

configurations resumes packet forwarding. Truncation and timestamping are not supported together. (72756)

- The VXLAN flood VTEP list cannot be more than 128k VLAN-VTEPs, where VLAN-VTEPs is the number of VLANs multiplied by the number of VTEPs. (73228)
- Ingress port access control list (PACL) on a VXLAN-enabled interface is not supported. If ingress PAACL is configured on a VXLAN-enabled interface, in-hardware VXLAN encapsulation of incoming packets is not expected to work correctly. VXLAN encapsulation of packets in software is not affected. Removing ingress PAACL configuration on the affected interface restores correct VXLAN hardware encapsulation behavior on that interface. (76085)
- Multicast-routed frames get tx-mirrored even when they do not egress the tx-mirrored multicast destination due to VLAN suppression. This is a limitation of the switch ASIC. (87933)
- 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)
- 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)
- 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)
- 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 is not supported if collector is reachable via tunnel. (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)
- 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)
- 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)
- 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 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)

- 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)

DCS-7280R and DCS-7500R Series

- An IPv6oMPLS packet terminated by L3 EVPN MPLS configuration will not be forwarded. The workaround is to disable IPv6 exact match host routes using 'no platform sand ipv6 host-route exact-match' and disable IPv6 route optimizations for prefix length 128. (201414)
- 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)

- 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)
- 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)
- sh policy-map type control-plane copp-system-policy shows wrong hardware programming status. (367000)
- IP egress ACLs are not supported for MPLS terminated packets. (383263)
- Strict Priority scheduling for 400G interfaces don't work properly for 64 byte packets when traffic rate is above PPS limit of the chip. (388517)
- Cli to change tail-drop threshold for 200/400G ports is not supported by software. (396029)

- 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)
- 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)
- ISIS-SR egress LER (MPLS pop) operation is not supported. (418161)
- Set dcsp 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 milliseconds when a VxLAN nexthop (as part of ECMP) undergoes modifications. (268600)
- The DSCP value of BGP control plane packets set by the command 'bgp transport qos dscp <value>' is overwritten when 'qos rewrite dscp' is enabled and the outgoing routed interface is a SVI. The rewritten value is based on traffic-class to DSCP map configuration. (288617)
- If a different XP profile than the one that is currently running is configured, and ACL configuration/ACL viewing commands without running a chip reset/reboot (via platform xp xp0 reset), the command can cause XpAcl to restart. (418083)

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)

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-7010T 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-7010T, 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-7010T 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-7010T 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-7010T 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)
- 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

- 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

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

- 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
- NAT is not supported on ECMP interfaces (353704)
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)
- 50G speed is not supported on DCS-7050CX3M in this release (416146)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- 1G speed is not supported on DCS-7050CX3M in this release (416148)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- 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)

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)
- 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)

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

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)

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.