

# **Arista Networks EOS 4.23.7.1M Release Notes**

Version: 1.0

24 March 2021

# Table of Contents

## **EOS 4.23.7.1M Release Highlights**

New Platforms and Hardware

## **SNMP MIBs**

## **SNMP Traps**

## **Supported Hardware**

## **Reference to MLAG ISSU Upgrade/Downgrade Table**

## **Resolved Caveats in 4.23.7.1M**

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

DCS-7280R and DCS-7500R Series

## **Known Software Caveats in 4.23.7.1M**

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

Containerized EOS

CVP

Network Provisioning

CVX

General

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

MPLS

Multicast

vEOS Router / CloudEOS

DCS-7150 Series

DCS-7170 Series

DCS-7170 Series

CCS-720XP Series

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

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

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

DCS-7500E Series

DCS-7280R and DCS-7500R Series

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

DCS-7020 Series

DCS-7500R3 and DCS-7800R3 Series

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

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

DCS-7280R and DCS-7280R2 Series

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

DCS-7010 Series

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

DCS-7260X3 Series

7368 and DCS-7060X4 Series

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

DCS-7050X2 Series

CCS-720XP Series

DCS-7050X3 Series

Transceiver Series

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

DCS-7160 Series

#### **Limitations and Restrictions in 4.23.7.1M**

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

CVX

General

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

MPLS

Multicast

vEOS Router / CloudEOS

DCS-7150 Series

DCS-7170 Series

DCS-7170 Series

CCS-720XP Series

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

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

CCS-720XP Series

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

DCS-7500E Series

DCS-7280R and DCS-7500R Series

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

DCS-7020 Series

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

DCS-7280R and DCS-7280R2 Series

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

DCS-7010 Series

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

DCS-7260X3 Series

7368 and DCS-7060X4 Series

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

DCS-7050X2 Series

Transceiver Series

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

DCS-7160 Series

## **Conditions and Impacts**

# EOS 4.23.7.1M Release Highlights

## New Platforms and Hardware

- None

## SNMP MIBs

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

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

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

## SNMP Traps

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

## Supported Hardware

### Fixed System

- |                   |                     |                        |
|-------------------|---------------------|------------------------|
| • DCS-7150S-24-CL | • DCS-7280SR-48C6   | • DCS-7050QX-32S-SSD-F |
| • DCS-7150S-24    | • DCS-7280SR-48C6-F | • DCS-7050QX-32S-SSD-R |

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

- PWR-745AC
- PWR-745AC-F
- PWR-745AC-R
- FAN-7010
- PWR-1900-DC
- PWR-1900-DC-F
- PWR-1900-DC-R
- PWR-1900AC
- PWR-1900AC-F
- PWR-1900AC-R
- PWR-747AC
- PWR-747AC-F
- PWR-747AC-R
- FAN-7002H-R
- FAN-7000H-R
- PWR-1600AC
- PWR-1600AC-F
- FAN-7001D
- FAN-7001D-F
- FAN-7001C
- FAN-7001C-F
- PWR-400AC-BLUE
- FAN-7011M
- FAN-7011M-F
- FAN-7011M-R
- PWR-1021-AC-RED
- PWR-511-AC-BLUE
- PWR-511-AC-RED
- PWR-621-AC-RED
- PWR-400AC-RED
- FAN-7001DH
- FAN-7001DH-F
- PWR-1611-AC-RED
- PWR-1611-DC-RED
- FAN-7011H
- FAN-7011H-F
- FAN-7011H-R
- PWR-511-DC-BLUE
- PWR-511-DC-RED
- PWR-1511-AC-BLUE
- PWR-1511-AC-RED
- PWR-1511-DC-BLUE
- PWR-1511-DC-RED
- PWR-1011-AC-RED
- PWR-721-DC-RED
- PWR-621-AC-BLUE
- FAN-7012H-RED
- DCS-7280CR2-60
- DCS-7280CR2-60-F
- DCS-7280CR2A-30
- DCS-7280CR2A-30-F
- DCS-7280CR2K-30
- DCS-7280CR2K-30-F
- DCS-7280QRA-C36S-M-F
- DCS-7280QRA-C36S-M-R
- DCS-7280SR2K-48C6
- DCS-7280SR2K-48C6-M-F
- DCS-7280SR2K-48C6-M-R
- DCS-7020SR-24C2
- DCS-7020SR-24C2-F
- DCS-7020SR-24C2-R
- DCS-7280CR2M-30
- DCS-7280CR2M-30-F
- DCS-7280SRM-40CX2
- DCS-7280SRM-40CX2-F
- DCS-7280SRM-40CX2-R
- DCS-7020SR-32C2-F
- DCS-7020SR-32C2-R
- DCS-7020SRG-24C2
- DCS-7020SRG-24C2-F
- DCS-7020SRG-24C2-R
- DCS-7280CR3-32P4
- DCS-7280CR3-32P4-F
- DCS-7280CR3-32P4-M-F
- DCS-7280CR3-32P4-M-R
- DCS-7280CR3-32P4-R
- DCS-7280CR3K-32P4
- DCS-7280CR3K-32P4-F
- DCS-7280CR3K-32P4-R
- DCS-7280CR3-32D4
- DCS-7280CR3-32D4-F
- DCS-7280CR3-32D4-M-F
- DCS-7280CR3-32D4-M-R
- DCS-7280CR3-32D4-R
- DCS-7280CR3K-32D4
- DCS-7280CR3K-32D4-F
- DCS-7280CR3K-32D4-R
- DCS-7280PR3-24
- DCS-7280PR3-24-F
- DCS-7280PR3-24-M-F
- DCS-7280PR3K-24
- DCS-7280PR3K-24-F
- DCS-7280DR3-24
- DCS-7280DR3-24-F
- DCS-7260QX-64-SSD-R
- DCS-7050SX-72Q
- DCS-7050SX-72Q-F
- DCS-7050SX-72Q-R
- DCS-7050TX-72Q
- DCS-7050TX-72Q-F
- DCS-7050TX-72Q-R
- DCS-7050QX2-32S
- DCS-7050QX2-32S-F
- DCS-7050QX2-32S-R
- DCS-7050SX2-128
- DCS-7050SX2-128-F
- DCS-7050SX2-128-R
- DCS-7050TX2-128
- DCS-7050TX2-128-F
- DCS-7050TX2-128-R
- DCS-7050SX2-72Q
- DCS-7050SX2-72Q-F
- DCS-7050SX2-72Q-R
- DCS-7060CX-32S
- DCS-7060CX-32S-F
- DCS-7060CX-32S-R
- DCS-7060CX2-32S
- DCS-7060CX2-32S-F
- DCS-7060CX2-32S-R
- DCS-7260CX-64
- DCS-7260CX-64-F
- DCS-7260CX-64-R
- DCS-7260CX-64-SSD-F
- DCS-7260CX-64-SSD-R
- DCS-7060SX2-48YC6
- DCS-7060SX2-48YC6-F
- DCS-7060SX2-48YC6-R
- DCS-7260CX3-64
- DCS-7260CX3-64-F
- DCS-7260CX3-64-R
- DCS-7050CX3-32S-F
- DCS-7050CX3-32S-R
- DCS-7050SX3-48YC12
- DCS-7050SX3-48YC12-F
- CCS-720XP-24ZY4
- CCS-720XP-24ZY4-F
- CCS-720XP-48ZC2
- CCS-720XP-48ZC2-F
- DCS-7050SX3-48YC8
- DCS-7050SX3-48YC8-F
- DCS-7050SX3-48YC8-R

- PWR-1011-AC-BLUE
- PWR-1011-DC-BLUE
- PWR-1011-DC-RED
- PWR-2411-AC-RED
- PWR-2411-DC-RED
- FAN-7000
- FAN-7000-F
- FAN-7000-R
- PWR-460AC
- PWR-460AC-F
- PWR-460AC-R
- PWR-460DC
- PWR-460DC-F
- PWR-460DC-R
- DCS-7280TR-48C6
- DCS-7280QRA-C36S
- DCS-7020SR-32C2
- DCS-7280SE-64
- DCS-7280SE-64-F
- DCS-7280SE-64-R
- DCS-7280SE-72
- DCS-7280SE-72-F
- DCS-7280SE-72-R
- DCS-7280SE-68
- DCS-7280SE-68-F
- DCS-7280SE-68-R
- DCS-7280CR-48
- DCS-7280CR-48-F
- DCS-7280QR-C36
- DCS-7280QR-C36-F
- DCS-7280QR-C36-M-F
- DCS-7280QR-C36-M-R
- DCS-7280QR-C36-R
- DCS-7280DR3-24-M-F
- DCS-7280CR3-96
- DCS-7280CR3-96-F
- DCS-7280CR3K-96
- DCS-7280CR3K-96-F
- DCS-7280CR3MK-32P4
- DCS-7280CR3MK-32P4-F
- DCS-7280CR3MK-32P4-R
- DCS-7050TX-128
- DCS-7050TX-72
- DCS-7050TX-96
- DCS-7050CX3-32S
- CCS-720XP-24Y6
- CCS-720XP-48Y6
- DCS-7050SX-128
- DCS-7050SX-128-F
- DCS-7050SX-128-R
- DCS-7050SX-128-SSD-F
- DCS-7050SX-128-SSD-R
- DCS-7250QX-64
- DCS-7250QX-64-F
- DCS-7250QX-64-R
- DCS-7250QX-64-SSD-F
- DCS-7250QX-64-SSD-R
- DCS-7050QX-32
- DCS-7050QX-32-F
- DCS-7050QX-32-R
- DCS-7050QX-32-SSD-F
- DCS-7050QX-32-SSD-R
- DCS-7050QX-32S
- DCS-7050QX-32S-F
- DCS-7050QX-32S-R
- DCS-7260CX3-64E
- DCS-7260CX3-64E-F
- DCS-7260CX3-64E-R
- DCS-7060DX4-32
- DCS-7060DX4-32-F
- DCS-7060PX4-32
- DCS-7060PX4-32-F
- CCS-720XP-24Y6-F
- CCS-720XP-48Y6-F
- DCS-7050CX3-32S-SSD-F
- DCS-7050CX3-32S-SSD-R
- DCS-7050CX3M-32S
- DCS-7050CX3M-32S-F
- DCS-7050CX3M-32S-R
- CCS-720XP-24Y6-R
- CCS-720XP-48Y6-R
- DCS-7160-32CQ
- DCS-7160-32CQ-F
- DCS-7160-32CQ-R
- DCS-7160-48YC6
- DCS-7160-48YC6-F
- DCS-7160-48YC6-R
- DCS-7160-48TC6
- DCS-7160-48TC6-F
- DCS-7160-48TC6-R
- FAN-7002
- FAN-7002-F
- FAN-7002-R
- FAN-7002H
- FAN-7002H-F
- FAN-7000H
- FAN-7000H-F

## Modular

- FAN-7002
- FAN-7002-F
- FAN-7002-R
- FAN-7002H
- FAN-7002H-F
- FAN-7000H
- FAN-7000H-F
- PWR-3K-AC
- PWR-3K-AC-F
- PWR-3K-AC-R
- PWR-2700-DC
- DCS-7508N
- 7500E-6CFPX-LC
- DCS-7500-SUP2
- DCS-7500-SUP2-D
- 7500R-36CQ-LC
- 7500R-36Q-LC
- 7500R-48S2CQ-LC
- 7504R-FM
- 7508R-FM
- 7512R-FM
- DCS-7512N
- 7800R3-48CQ-LC
- 7800R3-48CQM-LC
- 7800R3K-48CQ-LC
- 7808R3-FM
- DCS-7800-SUP
- DCS-7508
- DCS-7504
- 7300-SUP
- 7300-SUP-D
- 7300X-32Q-LC
- 7300X-64S-LC

- PWR-2700-DC-F
- PWR-2700-DC-R
- PWR-3K-DC-BLUE
- PWR-3K-DC-RED
- PWR-3KT-AC-BLUE
- PWR-3KT-AC-RED
- PWR-D1-3041-AC-BLUE
- PWR-2900AC
- 7500E-36Q-LC
- 7500E-48S-LC
- 7500E-72S-LC
- 7504E-FM
- 7508E-FM
- DCS-7500E-SUP
- DCS-7500E-SUP-D
- 7500E-12CM-LC
- 7500E-6C2-LC
- DCS-7504N
- 7500E-48T-LC
- 7500E-12CQ-LC
- 7500RM-36CQ-LC
- 7500R-8CFPX-LC
- 7500R2-18CQ-LC
- 7500R2-36CQ-LC
- 7516R-FM
- DCS-7516-SUP2
- DCS-7516-SUP2-D
- DCS-7516N
- 7500R2A-36CQ-LC
- 7500R2AK-36CQ-LC
- 7500R2AK-48YCQ-LC
- 7500R2M-36CQ-LC
- 7500R2AM-36CQ-LC
- 7500R3-36CQ-LC
- 7504R3-FM
- 7508R3-FM
- 7500R3K-36CQ-LC
- 7500R3-24P-LC
- DCS-7808-CH
- 7800R3-36P-LC
- 7300X-64T-LC
- 7304X-FM
- 7308X-FM
- DCS-7304
- DCS-7308
- 7316X-FM
- DCS-7316
- 7320X-32C-LC
- 7324X-FM
- 7328X-FM
- 7300X3-32C-LC
- 7304X3-FM
- 7308X3-FM
- 7300X3-48YC4-LC
- 7368
- 7368-16C
- 7368-SUP
- 7368-SUP-D
- 7368X4-SC
- 7308X3-FM2

## Transceiver

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

- 100GBASE-SR4
- 40GBASE-SRBD
- 100GBASE-AOC

- 10GBASE-ERBD-U
- 10GBASE-ERLBD
- 10GBASE-ERLBD-D

- SFP-10G-T
- 10GBASE-ER
- 40GBASE-CR4

## Reference to MLAG ISSU Upgrade/Downgrade Table

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

## Resolved Caveats in 4.23.7.1M

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

#### DCS-7280R and DCS-7500R Series

- When a device with VLAN translation entries configured on LAGs where the VLAN tag matching rules overlap with bridging VLANs on the same LAG is rebooted, there may be a window of time in which packets are mapped to incorrect bridging VLANs (571307)

## Known Software Caveats in 4.23.7.1M

### Accelerated System Update

- Performing SSU when /mnt/flash/persist/persistentRestartLog doesn't exist may result in extended traffic outage due to a watchdog reset that can occur. (158117)
- When performing ASU2 on an MLAG setup with VXLAN configuration, ASU2 can result in a fatal error resulting in a cold boot being performed resulting in large traffic loss. (245555)  
Series Affected: DCS-7060X and DCS-7060X2 Series
- SSU with Splunk enabled may result in extended boot times, which could lead to protocol timeouts or switch cold rebooting. (343629)
- When performing SSU on an MLAG setup, SSU can cause peer links on port-channel to flap, resulting in large traffic loss. (405788)

### BGP EVPN L2/L3 VXLAN/MPLS

- A remote VTEP reached through unnumbered BGP over IPv6 will be incorrectly displayed as a configuration error with "FAIL" state due to "No route" (525842)
- Some MAC addresses may get blacklisted when the BGP agent is restarted on a multihomed peer. The workaround is to clear the blacklist using the "clear bgp evpn host-flap" command. (538177)
- During transition of VXLAN Encapsulation from dual stack to IPv4, the routes may not be installed correctly. The work-around is to reset the BGP session by issuing 'clear ip bgp' CLI command. (555495)

## Containerized EOS

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

## CVP

### Network Provisioning

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

## CVX

- If a VmTracer session configures all the available VLANs on the switch, then the switch may errdisable the routed ports by freeing up the internal VLANs. (139294)
- On all switch series configured as a MLAG pair where CVX is used as the VXLAN control plane, on rare occasion after a MAC address move between switches the MAC address can be advertised to the CVX server by two different VTEPS. This can cause packets destined to the MAC address to be blackholed. The workaround is to clear the mac address on both the advertising VTEPS so that the mac address will be relearned properly. (158130)
- The MacroSegmentation Service (MSS), working with L2 transparent firewalls, requires the firewall interface be statically identified on CVX via configuration commands (as opposed to using LLDP to detect them dynamically). This prevents issues such as traffic loss in the event that a member port of a aggregated link goes down. (275384)
- Macro-segmentation Service (MSS) might stop intercepting traffic when the last member of the Near service LAG goes down.

To work around this failure, use the interface mapping from the MSS dynamic device-set configuration to statically map service device interfaces to their corresponding switch port channels. (275656)

- When a security policy on a CheckPoint firewall includes unsupported services, the Macro-Segmentation Service (MSS) intercepts traffic for all services for the specified end-points in the policy. (433067)
- When a group on NSX-T containing more than 1000 hosts, policy rules for the group on CVX will only generate ACLs for the first 1000 hosts in that group. (459644)

- The Policy Control Service (PCS) fails to apply any security policy configured on NSX-T on a switch interface after the switch configuration is modified to remove all tags from an interface followed by re-adding tags, using the configuration replace feature. (463958)
- When 'client state preserve' feature is configured on CVX, state of the primary switch of an MLAG VTEP will not be preserved upon disconnection. If the primary switch did not reboot and disconnected from CVX due to network faults, it may cause disruption of VXLAN control plane leading to VXLAN traffic loss.

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

- In a Macro-Segmentation Service (MSS) deployment with Panorama, uncommitted changes to the list of managed devices may be used by MSS to program rules on the switches. (536412)
- When the Policy Control Service is enabled on CVX, under certain circumstances such as re-numbering IP on a host connected to a tagged port, the groups received from NSX controller may not be resolved into IP addresses. This may result in incorrect programming of ACLs on the switches.

Workaround is to disable and re-enable 'service pcs' on CVX. (538537)

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

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

## General

- When a policy map of type PBR is attached to an interface that is either in down state or a switchport (or both), and if the policy map is too large to fit into available hardware resources, the CLI will not report the error and roll back the configuration. Later, when the interface becomes an operational routed interface, the policy map will not be programmed into hardware. However, "show policy-map type pbr" shows the policy applied to the interface.

To fix the discrepancy, remove extra rules from the policy map to make it fit into the hardware. (89931)

- OpenFlow Flows may not be programmed in the right priority order after changing OpenFlow Table Profile config to "auto" mode. (118767)

SKUs Affected: DCS-7260CX3-64, DCS-7260CX3-64-F, DCS-7260CX3-64-R,  
DCS-7260CX3-64E-F and DCS-7260CX3-64E-R

- While running a combination of both UDP and TCP traffic (e.g. BFD and BGP), a race condition may lead to a kernel panic due to reusing previously freed memory. (131245)
- 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

- When configuration of a service-policy of type pdp on an interface fails due to lack of hardware resources, the interface may not be protected by any PDP policy. The workaround is to remove the configuration of the PDP service-policy and to configure the PDP service-policy default-pdp-policy on the interface. (216365)
- The CLI command "route-target import auto <ASNumber>" for VLAN VRF does not work and hence should not be used. Alternatively, the suggested configuration is "route-target import auto". In this case the AS Number is picked up from the BGP configuration.

Example:

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

Suggested Usage

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

- The VM Tracer agent may restart when connected to a VMware vCenter instance with a VMKernel not assigned with an IPv4 address. (258986)
- The MacroSegmentation Service (MSS) feature doesn't work in conjunction with the VARP VTEP IP configuration required for VXLAN routing where a subset of VTEPs are L2 VTEPs. (263532)
- Support for VRRPs in OpenConfig not fully implemented (268736)
- 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 multiple front-panel ethernet interfaces are configured as reflector interfaces for the same flow of traffic to be reflected, loops can be formed and hosts can flap from one interface to another when the reflector configuration is subsequently changed. The workaround is to only configure one reflector interface to handle traffic reflection. (360869)  
Series Affected: DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7500E, DCS-7500R and DCS-7500R2 Series
- When configuring VXLANs from the EOS CLI, "ip address virtual" is supported and mapped to YANG, but "ip virtual-router address" is not supported. (372193)
- Mlag agent may crash if the system is being put in Maintenance Mode if interface profile has shutdown max-delay enabled. Workaround is as follows:  
1. Create a maintenance interface profile which does not have shutdown max-delay configured with profile name being lexicographically largest.  
  
profile interface zzz\_no\_shutdown\_delay  
rate-monitoring threshold 4294967295  
2. Create an interface group constituting all the physical ethernet interfaces of peer-link and attach the above interface profile that does not have shutdown max-delay configured.  
3. Create a unit with this group and quiesce it. (380387)
- 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)
- DirectFlow 'action output nexthop <ip addr>' ignores TTL of incoming packet and forwards the packet even if the TTL has expired. (417994)
- When an OpenConfig NETCONF client is used to add new classes to policy map, the new classes are added before any previous classes. (451025)
- ISIS configuration on a non-default VRF via OpenConfig fails. (460407)

- In rare conditions, the kernel may panic with the message containing "BUG at ../mm/slub.c:3334" (476225)
- Any configured QoS mappings for traffic class to COS and traffic class to DSCP bits are not applied on packets that are reflected via an Interface Reflector that is configured in MAC swap and direction out mode. The source and destination MAC addresses are swapped as expected, but the COS and DSCP bits of the reflected packets are copied unmodified from the original ingressing packet. (476507)  
 Series Affected: DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series
- When IP Locking is configured with more than one DHCP server is configured, if a host initially leases an IP from one server and leases a second IP from one of the other configured server, the host may experience packet drops for up to 150 seconds after leasing the second IP. (476860)
- On system reload, the kernel may crash with the "PANIC: double fault" message. (493516)
- EosSdk agents using a certain combination of Sysdb mount paths may cause the Sysdb agent to restart on connecting to Sysdb. If the agent is configured and stored in startup-config this can cause Sysdb to get stuck in a restart loop. (496371)
- Adding and deleting Nexthop groups via EosSdk APIs, may increase process memory usage over a period of time. (498368)
- For an egress policy-map, attaching a policer to a class-map that requires more hardware TCAM rules than that can be fit in a single TCAM slice is not supported. This is a hardware limitation. (529081)
- When communication between CVX and switch is secured using SSL profile, Controllerdb agent on CVX or Sysdb agent on switch can restart unexpectedly. (537015)
- When a member interface of a PTP-enabled port channel is removed from the port channel, a subsequent speed change on the former member interface may cause the Strata agent to unexpectedly restart. A workaround would be to disable PTP on the port channel before removing the member interface. (540175)  
 Series Affected: DCS-7050CX3 and DCS-7050SX3 Series
- The "arp reply relay" feature does not support a VXLAN domain with L2 VTEPs. (540369)
- EOS is not vulnerable to CVE-2020-26116. However, custom agents using the python3 runtime can be vulnerable. http.client in Python 3.x before 3.5.10, 3.6.x before 3.6.12, 3.7.x before 3.7.9, and 3.8.x before 3.8.5 allows CRLF

injection if the attacker controls the HTTP request method, as demonstrated by inserting CR and LF control characters in the first argument of `HTTPConnection.request`. The fix for this CVE-2020-26116 prevents header injection in http methods. (543494)

- EOS is not vulnerable to CVE-2019-20907. However, custom agents using the python3 runtime can be vulnerable.

In `Lib/tarfile.py` in Python through 3.8.3, an attacker is able to craft a TAR archive leading to an infinite loop when opened by `tarfile.open`, because `_proc_pax` lacks header validation. The fix for CVE-2019-20907 updates the python3 tarfile module to prevent an infinite loop. (543497)

- Modular switch may restart in case of Fabric or Linecard failure. (554503)
- When using OSPFv3 and encryption, the Rib agent may stop responding and restart unexpectedly. (560294)
- BmpAgent, Sflow or BgpCliHelper may abort and restart with the agent log message "DirectRegionHugeMap::remove() i:42 is not region:XXX" (563870)
- In PTP transparent clock mode, a PTP packet updated by the switch may have its VLAN tag stripped when it is forwarded. This can happen if the VLAN-tagged PTP packet ingresses a PTP-enabled VLAN trunk port, egresses a PTP-enabled VLAN trunk port, and there are other PTP-enabled ports configured with a switchport access VLAN. (566037)

## Layer 1

- When a QSFP port is set to 4x25G rate, inserting a 40G QSFP transceiver can trigger a single or continuous forwarding agent restart. A workaround is to configure the port for 40G or 4x10G rate before inserting the transceiver. (405185)

Series Affected: 7300X3, DCS-7050X3 and DCS-7260CX3 Series

- Configuring speed, forward error-correction or speed-groups on a large number of interfaces at once may cause the forwarding agent to restart and all interfaces to flap. The workaround is to configure interfaces in smaller batches. (494284)

Series Affected: 7368 and DCS-7060X4 Series

## Layer 2

- If Ebra agent restarts on a VTEP that is running EVPN control plane with route type 5, then it can result in VxLAN traffic loss. (276063)
- MACsec delay protection feature is not working. (282186)

- When running rapid-PVST at high scales of interfaces and vlans, the device may occasionally not transmit a BPDU, which can lead to a BPDU timeout on the neighbor. (442663)
- Performing an SSO switchover on an MLAG primary peer running Multiple Spanning Tree Protocol may result in STP topology reconverging. (470962)
- Restarting STP agent with a high number of interfaces and vlans in PVST mode may result in STP topology reconverging. (472675)
- On a device that's an MLAG or had MLAG configured if egress sFlow is enabled, interfaces can go to blocking with syslog SPANTREE-4-INTERFACE\_SELF\_LOOPED.

To workaround, disable egress sFlow. (546134)

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

- LACP port-channel member interface(s) may unexpectedly leave the port-channel due to "partner not in sync" syslog reason without member interface link status change. The member interface will rejoin the port-channel immediately after. However, the event may cause port-channel interface to flap if there is only one member interface configured in it. (546649)
- On a device that's an MLAG or had MLAG configured in the past and configured with rapid-pvst STP mode, if more VLANs are configured than platform limit, then the StrataVlanTopo agent may continuously restart. Device won't log SPANTREE\_TOPOLOGY\_LIMIT\_EXCEEDED to indicate platform topology limit exceeded.  
To workaround, reduce number of VLANs to platform limit specified in Arista product documentation. (548254)

## Layer 3

- If a loopback interface is only configured with a link-local IPv6 address, OSPFv3 will not advertise the address. (131369)
- During reload or configuration changes, BFD session may be stuck in init state on one side. Restart Bfd agent to work around the issue. (195545)
- The Ldp agent may restart if configured with a large number of MPLS ECMP routes with many next hops. (273767)
- With scaled VRF config, in the order of 1K VRFs, if the Arp agent restarts, it may fail to come up. (275514)
- In certain scale situations upon a sysdb restart, arp entries can be learned in kernel but not show up under 'show arp'. When this happens, the work around

is to  
restart the Arp agent. (275749)

- Configuring 'max-metric router-lsa' on an OSPF router may cause summary LSAs to be generated with maximum metric. (276629)
- 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)
- BFD configurations with large numbers of sessions may see BFD session flaps during initial session bringup shortly after boot. (375773)
- When an ARP entry is added while there is a matching, existing virtual IP with a non-/32 subnet, traffic loss may occur because traffic may be flooded instead of forwarded to the CPU.

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

- When graceful restart is enabled under "router isis" and other protocol routes are redistributed into IS-IS then on ASU or SSO, some traffic that is forwarded using redistributed routes might be lost. (413706)
- Removing SVI configuration at scale via a config-replace operation may cause the BFD agent to unexpectedly restart. (414702)
- If a decap group with multiple IPv6 addresses are configured in either a config-session or config-replace, forwarding agent may restart unexpectedly, or the decap-group may not get programmed correctly. (421010)
- The MLAG agent may exit unexpectedly after upgrading from 4.22.0 or an older release if the MLAG peer switch is still running the older software version and DHCPv6 prefix delegation routes are being sync'd across the peers. (425810)
- When a IS-IS/BGP/OSPF peer is torn down on the receipt of the BFD session down notification on a link that is still physically up, there could be traffic loss on traffic destined to prefixes which have ECMP paths that also include the aforementioned link. (434576)
- In situations involving a route with a next hop IP getting deleted and added back without a next hop IP, static ARP entries falling in the same subnet as the

prefix of  
this route may not be correctly added.

Workaround is to restart the ARP agent. (436002)

- When a series of back-to-back configuration changes are applied to a BFD session, the session may not have the latest configured value. To work around the issue, un-configure and re-configure the BFD session. (458955)
- The output of 'show ip nat sync peer' may incorrectly indicate a connected state, even though an incorrect peer IP address has been configured. (464011)
- When there're a lot of NAT translations synchronized through NAT Sync in a very short time period, NAT Sync may hang and stop synchronizing. Restarting Nat agent with "agent Nat terminate" in CLI Enable mode may help. Retrying the command several times may be needed. (467300)
- DHCP relay agent might exit unexpectedly if several VRFs are continuously deleted and re-created. (471169)
- With 'ipv6 dhcp relay install routes' configured on more than 100 SVIs, DHCP relay agent may exit unexpectedly when a IPv6 DHCP helper address is configured simultaneously on all SVIs. (471338)
- In situations in which a large number of routes are forwarded over the same dynamic tunnel interface, the Arp agent may crash if the forwarding information of these routes are changed. The Arp agent will restart normally afterwards. (472837)
- When adding a VRF during config replace or quickly deleting and re-adding a VRF, the internal state of the ARP agent may go out of sync and ARP entries may be missing from "show arp." Workaround is to delete the VRF and re-add it. (474469)
- When the MAC address changes on a BFD peer, the session may flap. This can happen when the interface fails over from active to standby. (477559)
- Cspf agent may restart when TI-LFA protection is configured and the number of segments requiring protection is over 2K (481703)
- In EOS, host route computation is facilitated by the Arp agent. In situations with scaled neighbor entries, particularly when all or most of the entries map to the same route, the Arp agent may encounter a heartbeat timeout, leading to the Arp agent restarting. Typically, the Arp agent will run fine after the restart. (484901)
- In EOS, host route computation is facilitated by the Arp agent. In situations where the FIB routing table is at internet scale, particularly if all or most of the routes are tunneled, the Arp agent may encounter a

heartbeat timeout, leading to the Arp agent restarting. Typically, the agent will run fine after restart. (484902)

- TI-LFA backup paths will be computed slowly when there are large number of node-segments/anycast-segments/adjacency-segments that are protected. (489097)
- All ECMP static vxlan evpn routes for a prefix get deleted when one of the contributing static vxlan evpn routes is deleted. (493505)
- An LSR which is the penultimate hop in an RSVP LSP may incorrectly report a "bandwidth unavailable" error if the tunnel destination IP is an interface IP that does not match the last hop in the ERO.

A workaround is to use a loopback interface as tunnel destination or the interface IP that is present in the ERO. (493578)

- BFD sessions may continuously flap after a speed change on an interface. A workaround is to restart StrataBfd. (499185)
- When changing the OSPF Router Id, OSPF will not flush its existing self-originated LSAs from the routing domain before the new Router Id takes effect (499551)
- Isis agent will restart if an interface has ipv4 unnumbered configuration and ipv6 global address configuration and has enabled ipv4 and ipv6 traffic engineering. (501374)
- When the IP reachability of the TI-LFA protected link is leaked from one level into another level on the router in level-1 and level-2, then on link down local convergence delay timer is immediately aborted for the protected prefixes. This may cause the traffic loss due to micro loop. (504629)
- Protocol messages exceeding 1500 bytes are ignored, in particular BUNDLE and SREFRESH messages, which is the case for jumbo frames. The symptom can be dropped LSPs or long convergence times. (521363)
- An RSVP RESV message is immediately sent as a reply to a NACK object, which can lead to a packet storm if an RSVP neighbor incorrectly uses NACK objects as an error reporting mechanism. Configuring RSVP with `refresh method explicit` will prevent this issue but disable Refresh Overhead Reduction. (522041)
- When large number of distribute lists are configured in RIP then it may result in route flaps and the protocol agent CPU utilization can increase significantly. (536500)
- If neither Graceful Restart nor Hold Timer are negotiated for a peer, BMP will not send any Route Monitoring Messages for that particular peer. There is no workaround. (539641)

- KernelFib agent may restart on interface flap, if there are a large number of route resolving through that interface (539757)
- RSVP tunnels might not get established or may use non-best paths when IS-IS system ID or OSPF router ID is changed for any node in the IGP network and the best path goes via that node. One workaround is to change the TE router ID also for the same node. (546904)
- When TE router ID of one of the two routers which have the same TE router ID is changed to be same as that of a third router, Cspf agent may sometimes restart (565678)
- Rib agent may restart when BGP PIC ("BGP additional-paths install") is configured and a BGP path's next hop resolves over a BGP route with a backup path.  
Workaround is to unconfigure the additional-paths configuration. (566312)

## Multi-agent

- If a VRF is deleted while the Bgp agent is in the middle of processing paths in BRIB, it may restart unexpectedly. (209143)
- Performing a stateful swtichover in a switch with a large number of VRFs configured, with BGP config in all VRFs and IGP's configured in a large number of VRFs, may cause the system to restart unexpectedly, leading to traffic loss. (376773)
- When sampling packets with the BGP extension enabled, if the address family of the route does not match the address family of the next hop, the sFlow agent can restart unexpectedly.  
When sampling packets with sampled flow tracking, if the address family of the route does not match the address family of the next hop, the SftAgent can restart unexpectedly.

Use "no sflow extension bgp" to stop the sFlow agent agent from restarting. There is no workaround for sampled flow tracking. (427674)

- When a large number of updates are backed up in the transmit socket buffers, it may sometimes lead to the BGP neighbors session getting stuck in Idle (PendingNotification) state when the session reset is attempted and BGP is unable to send a notification. Not being able to send out updates may lead to traffic misforwarding or drops. (466575)
- Sometimes BGP (multi-agent) would advertise an incorrect originator-ID when sending updates to iBGP clients when configured as route-reflector. The problem only manifests when different paths (could be paths belonging to the same prefix or different prefixes) have the exact same path attributes, including NextHop. The problem corrects itself if the paths are resent to the route-reflector from the source whose paths had the incorrect

originatorId. This problem can manifest with add-path configuration as well when the additional paths have the same path attributes. (472025)

- During BGP graceful restart, 'update wait-for-convergence' behavior is unnecessarily forced for AFI-SAFIs which do not have graceful restart configured (either specifically per AFI-SAFI or globally). (478038)
- Bgp agent may restart unexpectedly when 'bgp always-compare-med' option is configured and large number of competing paths for the same prefix from several neighbors with different MED values are present (482275)
- During SSU or SSO, 'update wait-for-convergence' behavior is unnecessarily forced for AFI-SAFIs which do not support SSU or SSO. (483591)
- The configuration 'bgp missing-policy direction out action deny' may be ignored briefly when a route map applied on outbound to a BGP peer is deleted.

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

## Rib

- In the ribd routing protocol model, prior to EOS 4.21.3F release, the command  
"neighbor <peer|peer-group> send-community extended" (under "router bgp")  
will result in both standard and extended communities being sent  
(in outbound route advertisements) to the corresponding peer.

Starting with EOS release 4.21.3F this behavior has changed. The command  
"neighbor <peer|peer-group> send-community extended" will result in \*only\*  
extended communities sent to the peer.

If the previous behavior is desirable, then the command  
"neighbor <peer|peer-group> send-community" without additional keywords  
can be used. This command will ensure that all applicable community types  
are  
included in outbound route advertisements for the peer.

This change can be done in any release (running EOS version 4.18.1F or  
newer)  
prior to the upgrade to EOS 4.21.3F (or newer) release.

EOS 4.21.3F or newer releases allow much more granular control of what  
community  
types (standered, extended, large communities etc.) are included in the

outbound  
route advertisements. (384629)

- The Rib agent may restart when large number of interfaces go down at the same time and there is a large scale of BGP routes reachable over these interfaces. (549816)

## MLAG

- If non-MLAG reload-delay is configured to be lower than the MLAG reload-delay when LACP standby is used, traffic entering the non-MLAG interfaces destined towards hosts on the MLAG interfaces will be lost during the LACP standby period. (83330)
- MLAG ISSU breaks on a switch running a version of EOS < 4.20.5 if it receives ARP entries from the VXLAN Control Service (VCS) while the peer has been upgraded to EOS version >= 4.20.5 (497776)

## MPLS

- When the refresh overhead reduction extension is used, some RSVP messages may never be sent out (shown as "State: path-only" under "show mpls rsvp session").  
As a workaround, the extension can be turned off and on again. (441063)
- MPLS agent might restart due to route churn in FIB containing large number of routes. (526406)
- Excessive toggling of the 'mpls ip' or 'ip routing' commands may cause the LDP label range to be exhausted.

Wait for the LDP agent to shutdown after disabling 'mpls ip' before reenabling it to ensure the LDP labels are released. (539317)

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

## 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 unexpectedly restart. (354364)
- If large number of IGMP reports are received in a very short interval, some of them might get dropped resulting in the multicast route not being

created. The routes will be re-learned on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIgmp", consider increasing the bandwidth for this Copp class. (356675)

- Fast reload fails when multicast is configured on a non-default VRF that is not PIM enabled. The default VRF is always PIM enabled, but a non-default VRF is PIM enabled when it contains at least one PIM interface. Workaround is to enable a PIM interface on each non-default VRF. (398256)
- In an environment with a very large number of MSDP peers, during cleanup the MSDP Agent may unexpectedly restart. (418190)
- High CPU utilization might be observed for the forwarding agent when PIM Bidir routes are present. (479119)
- When using Sfe based forwarding configuration, the BessMgr agent may restart unexpectedly when multicast packet size is greater than 1500 bytes (481192)
- If MLD Snooping is configured, IPv6 multicast packets may not be bridged as expected in a PIM SM or PIM SSM non DR VLAN. (493226)  
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- BessMgr agent could restart unexpectedly when resolving a large number of unresolved packets at once. (495863)
- Removing router pim bsr also remove router pim sparse mode configuration. Workaround is to just remove specific configuration under router pim bsr mode. (516188)
- Using "software-forwarding sfe" for software forwarding of multicast packets does not work on x86\_64 version of EOS. (537616)

## vEOS Router / CloudEOS

- The OpenFlow agent fails to handle packet-in messages (145001)
- When several IPsec tunnels are up and traffic is passing through them, removal and addition of a IPsec license may cause the the Ipsec agent to restart.

The workaround is to shutdown all the interfaces before the IPsec license is removed and reapplied and then do a "no shutdown" on the tunnel interfaces. (248213)

- A reboot due to a kernel panic can happen during first 10 minutes of vEOS bootup on Microsoft Azure. The kernel panic occurs because of a task blocking without being scheduled for more than 300 seconds in "D" state. (249618)

- Interface may be allowed to be configured with un-supported speed of NIC using "speed <speed>", which will cause an unexpected PhyEthtool agent restart. Do not use "speed <speed>" command to configure interface speed that NIC does not support. (264395)
- When running vEOS devices in Sfe mode with interfaces having an MTU of 9214, users may encounter a scenario where not all routes are propagated through the network. The workaround is to have interfaces use an MTU of 9202 or less. (289886)
- On vEOS Router in Sfe/DPDK mode with many GRE/IPSec tunnels, changes to configuration may cause a long time to be applied to the datapath. (349552)
- When loading a vEOS instance in an Azure cloud, it is possible for the device to experience a kernel panic just after bootup, causing the device to reboot a second time. (354411)
- In an EVPN IRB deployment on vEOS, changing the VXLAN source-interface interface may result in EVPN ip-prefix routes not getting installed into one or more VRFs for one or more remote VTEPs. (369816)
- When multiple IPSec connections are created, some with AES encryption and some with NULL encryption and NULL authentication, bessd will crash when decrypting the IPsec packets encrypted with NULL encryption and NULL authentication. (378590)
- CloudEOS supports upto 600K BGP routes with 8G of system memory; If 800K BGP routes are programmed, out-of-memory condition will kill random processes.

Out-of-memory issue can be mitigated by increasing the system RAM to 12G to program 800K BGP routes. (478214)

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

The workaround is to restart the software forwarding engine (Sfe) agent. (523066)

- Changing the mapping of Vrf to VNI (VXLAN Network Identifier) twice will cause the software forwarding engine (Sfe) agent to restart. There will be a brief moment of traffic loss which will recover eventually. (539064)

## 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)
- On DCS-7150 series, FocalPointV2 agent might restart if a misbehaving host creates a large number of new connections in quick succession, without sending data.

This issue can be prevented by configuring NAT connection limits for such hosts. (149634)

## DCS-7170 Series

- On the DCS-7170 series MLAG protocol traffic shares the same CPU queue with traffic destined to CPU. A burst of such traffic, destined to CPU, can lead to MLAG flaps. It will recover automatically when the burst is over. (449993)
- On the DCS-7170 series L3 protocol traffic shares the same CPU queue with traffic destined to CPU. A burst of such traffic, destined to CPU, can lead to protocol flaps. It will recover automatically when the burst is over. (464981)
- On DCS-7170 series, a policer configuration with burst unit in mbytes will result in continuous crash of BfnQos agent. Remove the configuration to recover, and the issue can be worked around by specifying the burst in kbytes. (469790)
- IPv4 Multicast routes may be prematurely aged out (472845)
- Hardware errors in the switching ASIC may get silently ignored after an initial reporting. (484277)

## DCS-7170 Series

- When a monitor destination is continually congested due to fan-in/high rates of traffic, and LAG is used in the system, the slice agent may restart when there is LAG membership change or port status change. (516260)
- This is a debug issue. Some dropped packets are not accounted while they should be counted as "Invalid Egress Port" drops on the output of "show platform barefoot drop counters". The packets are dropped by the traffic manager but retrieving this counter is not available on some releases. (535991)
- When any form of NAT is disabled, the first packet for translation is sent to CPU and forwarded in software and after the translation is programmed in the forwarding ASIC, the subsequent packets are forwarded by the ASIC. If the NAT destination is reachable via ECMP paths, the path chosen by the first packet by software forwarding may not be the same as subsequent packets forwarded in hardware. This can potentially cause issues if the NAT destination is to load-balancer or an anycast address that maintains NAT flow state is maintained per receive port.

The workaround for this case is to have a single ECMP path for the NAT destination. (543692)

## 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)
- The forwarding slice agent may restart unexpectedly when a SEU is encountered in the BSC\_AG\_AGE\_TABLE during initialization. This may result in brief traffic outage. Switch is expected to start normal operation after the agent restart. (417588)
- When two or more hardware flow trackers are configured the forwarding agent may restart unexpectedly. In this case even if the extra flow tracker is removed the forwarding agent may continue to restart.

Workaround is as follows:

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

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

```
(config)# no flow tracking hardware
```

and reconfigure flow tracking hardware and add only one flow tracker. (426694)

- When a device behind phone/hub/switch connected to Arista switch is authenticated via Dot1x MAC Based Authentication, it cannot be moved to different port if dynamic VLAN is not received or dynamic VLAN is same as native VLAN. (477280)  
Series Affected: CCS-720XP Series
- Authentication of any data device behind a tagged phone will fail after the phone is authenticated and authorised via EAP in phone VLAN on that port. Also, if a data device authenticates before tagged phone doing EAP, the authorisation of phone in phone VLAN will fail. (490365)
- On phone trunk ports which are not controlled by Dot1x, phones do not honour port-security max-address limit and get installed in the L2 table even after the port-security limit is breached (494095)
- NAS sends two EAPOL identity requests in response to EAPOL start from supplicants. This can cause authentication timeouts for some Windows supplicants. (498610)

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

- Some front panel BASE-T interfaces might fail to forward traffic after a reload or linecard insertion (for modular systems), if they are configured with `speed 100full` command. The issue only affects the following 2.5GBASE-T interfaces:

Et1, Et6, Et9, Et14, Et17, Et22, Et25, Et30, Et33, Et38, Et41, Et46, Et49, Et54, Et57, Et62, Et65, Et70,

Workaround is to run `shutdown` followed by `no shutdown` on the affected interfaces. (519774)

SKUs Affected: CCS-720XP-24ZY4 and CCS-720XP-48ZC2

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

Workaround is to reload or power cycle the switch. (523649)

SKUs Affected: CCS-720XP-24Y6 and CCS-720XP-48Y6

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

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

SKUs Affected: CCS-720XP-24Y6 and CCS-720XP-48Y6

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

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

- The AlgoMatch forwarding agent will restart when config replace is used to configure access-lists with large number of rules. (240540)
- Following a physical error on link between Jericho and Algomatch FPGA the SandHalo agent may restart, causing ACLs to be removed and reinstalled. (283917)
- When an acl configuration is applied to a large number of SVIs in a configuration session, the acl configuration fails. (367513)
- A SEU can occur on the Jericho side of the ELK interface to the Halo FPGA. These error interrupts can happen for other reasons like oversubscribing the ELK interface, but if they persist, they are usually the result of a SEU. This results in dropped control packets from SandHalo and SandHalo will unexpectedly restart continuously. A reboot or linecard reset is required to clear this condition. (486369)

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

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

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

- When many IPv6 routes are configured and the BFD protocol in use, BFD sessions may flap when a linecard is removed or shut down. (279129)
- ManagementActive Agent repeatedly restarts and Management0 interface is not reachable in Modular systems. Doing switchover will address the issue. (369565)

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

- When the SandMcast and the SandFap agents are restarted, the VLAN floodset may not contain all the ports for several minutes. (67439)
- When in MLAG mode without ip routing enabled, under some control plane oversubscription scenarios the MLAG peer link might come down. The workaround is to either enable IP routing or have a static ARP entry on the MLAG peers for each other. (90247)
- An IPv6oGRE packet terminated using decap-group with a host route destination IPv6 address will be dropped.

The workaround is to disable exact match host routes using "no platform sand ipv6 host-route exact-match". (95873)

- On DCS-7280E, DCS-7500E series, DCS-7050X, DCS-7250X, and DCS-7300 series switches, the vlan tag on IGMP messages(report, query, leave) generated by the switch is not translated in the egress direction by the switchport vlan mapping command. (113104)
- Some packet loss may be observed when IP multicast traffic has high fanout on a single physical port (e.g. trunk port with many receivers for the same multicast group on different VLANs), even when the total egress packet rate is well below linerate. The default egress buffer settings may not perform well under this scenario. This issue can be resolved by adding "platform arad buffering egress port multicast max-queues-full 1" to the switch config and rebooting the switch. (115343)
- MAC mirroring ACLs do not match IPv6 packets. (137537)

- On a device with with a large number of active VLANs, doing "shutdown" and "no shutdown" on many interfaces may cause the SandMcast agent to restart. (175529)
- Mirrored traffic is not subject to egress security ACLs which are applied to mirroring destination ports. (190637)
- PBR is not supported along with IPv6 in IPv6 packets. (193914)
- On 100M interfaces acting as IEEE 1588 slave ports, restarting IEEE 1588 may cause ingress timestamps to spike by as much as 100s of ms.

A workaround would be to reset the interface after restarting IEEE 1588. (218386)

- Acl and related features that use TCAM for rule matching do not work correctly on packets with IPv6 extension headers if the rules have layer 4 match criteria such as tcp flags and source/destination ports. (221161)
- 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)  
 Series Affected: DCS-7020, DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series
- VTEP cannot learn a remote host's MAC address from VXLAN encapsulated IPv6 Neighbor Advertisement packets received from that host. This might be an issue in a pure IPv6 depolyment where the VTEP does not receive any bridged traffic from the host. In such a scenario, the traffic toward the host will be flooded in the VLAN.

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

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

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

- VXLAN bridging traffic sent out of LAG interfaces may be temporarily dropped after an SSO event. (251398)  
Series Affected: DCS-7500E and DCS-7500R Series
- When linecards are powered on/off, there may be a traffic loss for a short duration on the interfaces where vlan mapping feature is configured. (254476)
- Unknown unicast and multicast packets are unnecessarily copied to CPU after VXLAN decapsulation when there is an SVI configured on the VXLAN VLAN. (258990)
- When a bidirectional PIM multicast receiver is connected to only one MLAG peer switch, the IGMP reports from the receiver may get dropped resulting in multicast traffic drops towards the receiver. (259308)
- Egress acl deny logging does not work when the packet is destined to a RFC5549 next-hop. (264045)
- When a DirectFlow flow is created with an action to drop the outer tag, the packet is still leaving the interface tagged. (264336)  
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- When a physical port (e.g. Et1) is removed from a Port-Channel (e.g. Po1) that has subinterfaces configured (e.g. Po1.1), and there are subinterfaces configured under that physical port (e.g. Et1.1), the subinterfaces under the physical port can be incorrectly programmed. This results in traffic being dropped for the incorrectly programmed subinterfaces.

The workaround is to flap the subinterfaces affected. (269823)

- If one or more linecards in a modular system is running in tap aggregation mode, a small percentage of multicast traffic on linecards running in normal mode may be dropped, following a change to membership of some tap aggregation destination groups. (275286)
- After applying or removing shaping configuration on an L2 subinterface, the L2 subinterface must be flapped for the configuration to take effect. (278730)
- 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)
- Ethernet over MPLS packets being terminated and decapped on a LDP pseudowire will be hashed as if the inner header is IPv4 rather than Ethernet (382763)
- In the presence of virtual MAC with mask configuration, packets that undergo recirculation may not be properly forwarded. The workaround is to eliminate the configuration that causes the recirculation to happen. (387535)
- If VXLAN is configured, PBR and QOS policies won't apply to traffic ingressing on a MLAG peer-link. (390804)
- Traffic subject to CPU traffic policy `_policer_` action will not be `_forwarded_`. This includes multicast control traffic which is replicated to both the CPU but also forwarded (e.g. PIM, OSPF HELLOs, etc.). Unroutable unicast traffic is also flooded, and so will not be forwarded if matched by a CPU traffic policy policer action. An example policy of concern is

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

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

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

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

To workaround the issue configure `errdisable max flap value`; i.e.,

"errdisable flap-setting cause link-flap max-flaps 24 time 10". Note the setting is global and affects non Macsec Proxy interfaces as well. (396589)

- Insertion of a fabric module may cause brief traffic loss on the device (402508)  
SKUs Affected: 7504R3-FM and 7508R3-FM
- With scaled Acl configurations, SandAcl agent restart may result in additional agent restarts. (408260)
- When the DirectFlow 'priority' keyword is set greater than 65534 the configured priority will not be set to the correct value and the SandAcl agent may continually restart. (412206)
- With "ip pim spt-threshold infinity" configuration and having 10K or more Ipv4 (\*,G) multicast routes, in addition to what hardware can accommodate, SandMcast agent may restart repeatedly. Work-around is to reduce the scale of (\*,G) multicast routes to a number that can fit in hardware. (415164)
- Configuring 4094 different Egress IPv4 RACL on 4094 SVI will cause repeated SandAcl Agent crash.  
No workaround is available except for scaling down Egress IPv4 RACL. (415577)
- A switch reload with qos policy configured may result in SandAcl agent restart. (417935)
- Storm control maximum burst size is fixed to 10kB which can lead to storm control drops for bursty incoming traffic. (423231)
- With WRR scheduling on many ports and linerate across multiple traffic classes there can be a small amount of egress buffer drops. (425123)
- Sometimes on reload MACsec sessions on MACsec proxy ports stay down when FIPS restriction is enabled. (430728)
- When the outgoing interface for a Vxlan underlay route changes from one forwarding chip to another, it can impact VXLAN overlay route convergence. (440233)
- Performing many updates to a PBR policy in a short amount of time may cause the SandAcl agent to restart unexpectedly. (441912)
- With VXLAN configuration, if the route to a remote VTEP churns while the tunnel hardware resource usage is closer to the hardware limits, then the SandL3Unicast agent might restart unexpectedly. (445834)
- V6 unknown multicast packets are not headend replicated to remote VTEPs. (448950)

- The SandFap agent unexpectedly restarts when a PRBS type unsupported by the interface is configured; this will lead to traffic loss on all the interfaces. (450978)

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

- VXLAN routing over MACsec proxy port stops working when a proxy port that is not bound to an external interface is updated to be bound to an external interface. (455759)
- Sflow samples of VXLAN packets terminated by the switch, will have their VXLAN headers stripped. (457096)
- "match nexthop-group" filter is not supported in PBR rules when the corresponding nexthop-group is referenced by ECMP routes. (458189)
- On DCS-When a port-based connector is used for a pseudowire or local cross-connect patch, dot1q-tagged packets will have the outer tag incorrectly removed on ingress/encap if any one of several other features are configured on the switch (on any interface) including: L2 subinterfaces, Dot1q Tunnel, Vlan Mapping, and Vxlan. The only workaround is to unconfigure these other features. (474093)

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

- SflowAccel does not send flow samples to collectors routed via BGP routes with additional backup paths. (476553)
- If the number of BGP Flowspec matches exceed the hardware TCAM banks and entries available, and then Flowspec is enabled on additional interfaces while in the overflow state, the matches may fail to be programmed on the additional interfaces, even if the published Flowspec matches later fall back under the hardware limits. (477150)
- Once the maximum number of supported unique PBR policies per chip is exceeded, removing the PBR policy on an interface, converting the interface to a vlan member and applying a RACL on the SVI does not work. (477753)
- The SandAcl agent may crash if PBR is applied on SVIs during a linecard hotswap (486402)
- When BGP neighbors advertise Flowspec rules containing the police action and the system's policer resources are exhausted, any further change to Flowspec rules may cause the SandAcl agent to restart unexpectedly. Upon a SandAcl agent restart, existing Flowspec rules are removed from hardware and the system will attempt to program the new set of advertised rules, which will succeed if there are enough hardware resources.

As a workaround, restart the SandAcl agent after freeing enough policer

resources, e.g., remove rules from features such as QoS policies that use policer resources. (486688)

- Configuring 'port qualifier size N bits' in 'feature flow-spec port ipv4|ipv6' causes BGP Flowspec to stop working on non-Algomatch switches. (487271)
- TCAM bank reservation may not work after reload.  
The workaround is to remove and reapply the reservation config for any existing feature after all the SandFap agents are up. (488675)
- In a 6PE configuration, packet ingressing on an L3 subinterface with explicit null label and destined to this switch is not supported. (504582)
- Removing a field-set configuration (i.e., via "no field-set ipv4 prefix <field-set-name>") does not cause traffic-policies that match on the field-set to be updated in the switch. Packets will still match on the IP prefixes or L4-ports of the removed field-set configuration. The removal of the field-set from hardware will take effect in a subsequent change to the policy configuration. (512073)
- TCAM entries may become corrupted when two CLI commands which read TCAM entries from hardware are executed simultaneously. The commands which cause a TCAM hardware read are 'show platform fap acl tcam diff', and 'show platform fap acl tcam hw hits'. (521356)
- SandFap agent restart or a "config replace" operation may cause some subinterfaces to transition to "notconnect" state.

"shut/no shut" the parent interface or the affected subinterface to bring the subinterface back up. (535787)

- Ipv6 multicast routes with IIF as port-channel subinterface may remain unprogrammed after reloading the peer switch to which the IIF is connected.

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

- Routed traffic may be continuously dropped after all SandFap agents restart. The workaround is to restart SandL3Unicast agent. (540586)
- While in Tap Aggregation mode, changing the composition of a tool group may trigger a cascade of events which could lead to the group not able to replicate and transmit packets.

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

- 'Persistent Internal Drop DeqDeletePktCnt' error messages are seen due to incorrect/missing system ACEs. Hitful restart of SandAcl agent is needed to resolve the issue. The trigger for this issue is unknown. (548524)

## DCS-7500E Series

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

The workaround is to always enable autonegotiation on these ports. (70667)

SKUs Affected: 7500E-72S-LC

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

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

- In multi-agent mode, configuring multiple VRFs and Vxlan routing may lead to continuous OOM failures. (399674)
- AradLanz agent will restart if an invalid VOQ ID is read from the hardware. (557570)

## DCS-7280R and DCS-7500R Series

- In MPLS L2EVPN configuration, traffic requiring 3 or more LU labels to be pushed, may not be forwarded correctly. (275834)
- On the 7500R-8CFPX-LC linecard modules, rapid removal and reinsertion of the transceiver may result in the Denali agent restarting.

Inserting transceiver 10 seconds after removal from the same slot is the workaround. (293594)

SKUs Affected: 7500R-8CFPX-LC

- PFC pause frames, if received on /3 subport of Sesto2 master ports, will not be honored on the affected platforms. (411352)  
SKUs Affected: DCS-7280QR-C72-F, DCS-7280QR-C72-M-F, DCS-7280QR-C72-M-R,

DCS-7280QR-C72-R, DCS-7280QRA-C36S, DCS-7280QRA-C36S-F, DCS-7280QRA-C36S-M-F and DCS-7280QRA-C36S-M-R

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

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

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

- Configuring more than 6 MPLS labels for the pseudowire path (including VC label and flow label), will cause the SandTunnel Agent to crash and not recover. The only workaround is to reduce the number of transport labels and/or to remove flow label such that the total is 6 or fewer. (443140)
- When an ingress Port Acl includes TCP or UDP port range matches and vxlan is configured on the switch, any vxlan transit packet ingressing into the switch may incorrectly match on such rules. To prevent this, key field l4ops must be excluded from the tcam profile. (477195)
- Optimized routes may not be programmed properly if TCAM profile is changed from <profile-1> to <profile-2> with both profiles having 'feature flex-route' in it.

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

- With VXLAN, and PIM configuration on overlay SVI, VXLAN packets with inner IPv6 multicast packet consisting of Hop-by-Hop option may not be terminated and forwarded correctly. (484825)
- If the CPU is included in monitor sessions with multiple destinations, it will only work if the device mirror0 is selected. If multiple mirroring to CPU sessions with multiple destinations are configured, all packets will be sent to mirror0. (525919)
- When an SVI is configured with an MTU value and a new interface is created with breakout cable and configured to be in that vlan, the new interface is not programmed with the correct MTU value. The workaround is to re-apply the MTU configuration on the SVI. (536927)
- Reloading a switch in Tap Aggregation mode with multiple large tool groups (35 members or more) and continuous ingress traffic may result in the

forwarding chip's multicast buffer being locked up; this will causing some traffic to not be forwarded.

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

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

- During restart of SandAcl agent, packets may not be subjected to ACL rules for a short period of time. (372971)
- Neighboring interfaces may flap if speeds or speed groups are configured on an interface. (394764)
- A statically configured MAC entry may be replaced by a dynamically learned MAC entry leading to transient mis forwarding of packets destined to that MAC address. This transient condition gets corrected in the next poller cycle. (428580)
- SandAegis agent may crash when the last IPv6 rule is deleted from a traffic-policy which has both IPv4 and IPv6 rules (435733)
- VXLAN IPv6 overlay configuration along with VARP and Virtual VTEP configuration is not supported. (442914)
- On Switch with Jericho2 FAP(s), if Pimsm spt threshold is configured as infinity, Pimsm multicast traffic forwarding may not work properly. (447366)
- Speed change on a port with running traffic may lead to stuck queues for packets egressing on that port. If this occurs, shut the port and restart SandFapNi agent to recover. To avoid this, shut the port before changing speeds. (479144)
- Outer DSCP is not derived from Traffic Class in VXLAN encapsulated bridged packet. (487274)
- DSCP of VXLAN de-capsulated IP packet is set to DSCP value of outer IP packet. (488045)
- Forwarding of IPv4 and IPv6 multicast packets over MPLS Pseudowires is not supported. (493142)
- Links with 100GBASE-LR4/LRL4 transceivers starting with serial number prefix ADP may flap. (500081)  
    SKUs Affected: DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-R, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3MK-32P4-F and DCS-7280CR3MK-32P4-R

- Configuring 513 different Egress IPv6 SubInterface ACLs will cause repeated SandAcl Agent restarts. No workaround is available except for scaling down Egress IPv6 SubInterface ACLs. (502185)
- In interface traffic-policies, when configuring a destination port after configuring source port field-set, the ports are treated as field-set names (506962)
- If a MLAG LAG member is unconfigured while the SandAcl agent is not running, BUM traffic from the peerlink to the specific member may be dropped.

The work around is to reconfigure and unconfigure the MLAG LAG member while the SandAcl agent is running. (516712)

- Parity Interrupts for the first bank of first Fap of TCAM sometimes cause corruption in MLAG entries in the Large Exact Match table. This can be detected using the TCAM interrupt counter increase. Restarting of DUT fixes the problem. Alternatively, one can change the TCAM profile and change it back. (519945)
- When two or more of the following features: Port ACL, PBR, QoS, Flowspec, Aegis, and CPU Policy are applied on the same interface with a non-default TCAM profile, one or more of the features may not work (526913)
- If the "acl subintf ip/ipv6/mac" feature is added to a TCAM profile without the corresponding "acl port ip/ipv6/mac" feature being added to the same profile, subinterface ACLs will not work. (532393)
- Some queues on a port may exhibit a condition where enqueued packets may get stuck completely or drain very slowly. (532780)
- MTU violating VXLAN routed packets with either IPv6 overlay or an Ethernet tag on the underlay ethernet header can cause the egress port to lock up. When in this state, no packets will egress the port and traffic will get deleted on the ingress VOQ. Shutting and unshutting the port will restore the traffic. (537323)
- MPLS tunnel terminated GRE-over-IPv6-over-MPLS packets are forwarded incorrectly. (550842)

## 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)

- The SandIpsec agent does not automatically restart after a TCAM profile change, which means that IPsec-related changes to the profile will not take effect.

The agent must be manually restarted after the TCAM profile is changed using "agent SandIpsec terminate". (567183)

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

## **DCS-7500R3 and DCS-7800R3 Series**

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

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

- Traffic on a subinterface will be subject to BGP Flowspec matches applied to the parent interface, unless the TCAM profile has the flowspec feature configured with 'port qualifier size N', where  $N > 1$ . (541950)
- Configuring multiple DirectFlow flows of different priorities at once may restart the SandAcl agent continuously. A workaround would be to configure the DirectFlow flows one after the other. (543622)
- The flexroute feature may override the forwarding decisions from PBR policies and BGP Flowspec matches. (544872)
- Routed packets may not be forwarded correctly when L2 subinterfaces are configured. (559712)

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

- When using 64-bit versions of EOS, enabling auto-negotiation on any non-BASE-T interface will cause the forwarding agent to restart.

The workaround is to configure forced speed. (522773)

- Under sustained heavy load, when more than one interface is receiving traffic at line rate capacity, the SandCounters process may restart repeatedly until traffic is reduced. (543931)
- In the presence of virtual MAC with mask configuration, packets that undergo recirculation may not be properly forwarded. (559242)

## DCS-7280R and DCS-7280R2 Series

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

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

- When ACLs are applied on a large number of interfaces (more than 1000 SVI interfaces), the SandHalo agent may restart with a SOCKET\_CLOSE\_LOCAL syslog message. (253015)
- IPv6 routing with virtual VTEP configuration will not work on platforms that run in AlgoMatch mode. Workaround is to change the access-list mechanism configuration to tcam (472142)
- Traffic egressing out of a LAG interface may not be subject to egress ACL on reload (473227)
- RACLs are not supported for VXLAN traffic ingressing on a MLAG peer-link when AlgoMatch is enabled. (485677)

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

- In multi-chip systems with L2 MTU greater than 9236, packets that are larger than 9236 bytes that flow between the switch ASICs will be dropped. (102837)

Series Affected: 7300X3, DCS-7250X, DCS-7260CX and DCS-7300X Series

- Traffic streams matching DirectFlow flows with VLAN or MAC rewrite actions could face brief disruption across a StrataL3 agent restart. (111833)

Series Affected: DCS-7050X Series

- When performing reload hitless, link flaps on BASE-T ports might be observed. (152459)

Series Affected: DCS-7050TX Series

- MAC ACLs and PACLs cannot co-exist on the same interface if QoS policy ACLs exist in the switch (186949)

- The StrataVlanTopo agent may unexpectedly restart after many linecard online insertion and removal operations. (193787)  
Series Affected: DCS-7300X Series
- Ethernet65, Ethernet66 may experience extra link flap when performing fast boot. (221059)  
SKUs Affected: DCS-7260QX-64-F, DCS-7260QX-64-R, DCS-7260QX-64-SSD-F and DCS-7260QX-64-SSD-R
- The Strata fabric or linecard agent may restart due to fabric link flaps leading to traffic outage. (225302)  
Series Affected: 7320X and DCS-7300X Series
- DirectFlow flow in table type vfp will not match malformed Arp requests. (233922)
- When traffic is redirected using PBR and when the FIB lookup based on destination IP address indicates ARP miss, the PBR redirected traffic may be rate-policed.

The workaround is to use count action for the ARP-needed class. (238084)

- When the utilization of the MAC address table grows close to the full hardware capacity, the StrataL2 agent may encounter an out of memory (OOM) condition in the system, causing the agent to get restarted. (238115)  
Series Affected: 7320X, DCS-7010, DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7260X, DCS-7260X3 and DCS-7300X Series
- If nexthop resources are exhausted as indicated by syslog `VXLAN_HWHOP_NEXTHOP_RESOURCE_FULL`, forwarding of BUM traffic to that VTEP will be affected. (246391)  
Series Affected: DCS-7050X and DCS-7060X Series
- During an SSO failover, Strata Fabric agent may not detect hardware, resulting in a fatal error reboot with increased traffic outage. (250129)  
Series Affected: DCS-7300X Series
- In an L3 EVPN setup, traffic that is sourced by a VTEP that is a pure VXLAN router (i.e., it performs no VXLAN bridging and has no hosts attached to it) might incur some loss at the receiving VTEP, where the packets are decapsulated and forwarded.

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

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

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

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

- IS-IS configured over LAGs may see neighbor adjacency flapping during a hitless restart. (276232)
- A PCI error can break counter collection for an ASIC.

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

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

Workaround is to perform hitful restart of forwarding agent. (283175)

SKUs Affected: DCS-7050CX3-32S, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050SX3-48YC12, DCS-7050SX3-48YC12-F, DCS-7260CX3-64, DCS-7260CX3-64-F and DCS-7260CX3-64-R

- Applying a large configuration containing many interface changes may result in Strata forwarding agent restart, which could cause traffic loss during the reconfiguration. (330840)
- Some ports (esp. Et73-Et96 and Et103-Et104) might blackhole traffic or fail to linkup. (358361)  
SKUs Affected: DCS-7050TX-128-F, DCS-7050TX-128-R, DCS-7050TX-128-SSD-F and DCS-7050TX-128-SSD-R
- When VXLAN is configured and when a route or ARP points to an ecmp group with both local and remote nexthops, then a packet destined to this ecmp group may get looped between the linecards and fabric cards resulting in link congestion and packet drops. The work around is to disable the source port based hashing by removing "ingress-interface" in "ip load-sharing trident fields" command. (368073)
- Irrespective of IGMP snooping being enabled or disabled, IGMP packets will not be flooded to remote VTEPs. (368106)  
Series Affected: 7300X3, DCS-7050X2 and DCS-7050X3 Series
- DirectFlow entries that redirect traffic out of a port-channel may incorrectly egress packets on none or multiple members of the port-channel. (378467)  
SKUs Affected: DCS-7250QX-64-F, DCS-7250QX-64-R, DCS-7250QX-64-SSD-F,

DCS-7250QX-64-SSD-R, DCS-7260CX-64-F, DCS-7260CX-64-R, DCS-7260CX-64-SSD-F, DCS-7260CX-64-SSD-R, DCS-7304, DCS-7308 and DCS-7316

- After switching scheduler mode, the StrataCounters agent might be unable to complete a hitless reload shutdown stage, which will cause a hitless reload to be hitful. This can be pre-identified by seeing "SBus ack error in schan command" in the StrataCounters agent log. Workaround is to reset the ASIC before running reload hitless. (388694)

Series Affected: DCS-7050X3, DCS-7060X, DCS-7060X2 and DCS-7260X3 Series

- Having any ACLs with UDF qualifiers as well as any of the following LAG/ECMP hashing features present in the configuration: RRoce hashing, flexible hashing or TTL/HopLimit hashing, will result in incorrect operation of the LAG/ECMP hashing where the switch incorrectly hashes on IP payload. To avoid this problem the conflicting features should not be configured at the same time. (403182)

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

- Changing PFC Watchdog configuration while traffic is flowing can result in sub-second traffic outage when PFC Watchdog is in non-disruptive mode and multi-second traffic outage when PFC Watchdog is in disruptive mode. (403790)
- Switch may drop ingress UDP traffic destined to the switch MAC address when destination UDP port of the ingress packet matches VXLAN protocol default UDP port even if the destination IP address of the packet does not match any of the configured local VTEP IP addresses. (404138)

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

- Restart of StrataL2 may lead to continuous StrataL2 restarts. Workaround is to reload the system. (408138)
  - DirectFlow flows do not take precedence over conflicting router ACLs. (408734)
- Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX2, DCS-7050SX3, DCS-7050X2 and DCS-7050X3 Series

- DirectFlow flows can be uninstalled from a linecard on reaching the idle timeout for the flow even though matching traffic continues to match the flow on another linecard. The flow can be restored by re-programming it or disabling the idle timeout.

Series Affected: 7300X3 (411851)

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

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

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

- Configuring more than 16K MPLS pop/swap labels will cause continuous restart of forwarding agent. (419243)
- On ECMP shrinkage due to link down events, VXLAN tunnels reachable via ECMP may experience extended outage until routes reconverge. (419753)
- If the number of next-hop groups configured exceeds maximum supported hardware NHGs, then STRATA-3-MPLS\_MAC\_DA\_PROFILE\_TABLE\_FULL syslog message is logged and L3 forwarding agent may continuously restart.

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

- Ethernet49-54 may not link up with peers over long copper cables. Workaround is to enable autoneg on both ends of the link with 'speed auto <speed>' command. (421121)  
 SKUs Affected: DCS-7060SX2-48YC6-F and DCS-7060SX2-48YC6-R
- After a speed change, flows with mirroring actions on affected interfaces may not be mirrored as expected. Strata Slice agent can be restarted to restore expected behavior. (422221)  
 Series Affected: CCS-720XP, DCS-7050SX3 and DCS-7260CX3 Series
- The forwarding agent may continuously restart if the speed is changed from 100G to 50G when a 50G QSFP transceiver is inserted. This only affects 2-lane 50G transceivers, and does not affect 100G transceivers operating in 50G mode. A workaround is to configure the speed before inserting the transceiver. The forwarding agent restarts can be stopped by temporarily configuring the affected interfaces in 25G mode. (424990)
- In scale policy-map configuration, in some cases, a config replace can lead to restart of Strata slice agent. (430654)  
 Series Affected: DCS-7260CX Series
- During an ASU upgrade, interfaces with 100GBASE-SRBD transceivers may experience a single link flap. (440451)  
 Series Affected: DCS-7260CX3 Series
- Hitless upgrade may fail if there are errdisabled interfaces.

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

- Changing Acl configuration rapidly and multiple times may cause the forwarding agent to repeatedly restart unexpectedly. (447584)
- Special filter is installed on MLAG port channel member interfaces to block packets ingress from MLAG peer via peer link at egress to avoid loop. When

the filter is not installed on a member interface before it is added to MLAG port channel, there might be a momentary loop when interfaces belonging to an MLAG port channel link up. (471268)

- With PFC pause stream for lossless priorities on egress interface and data traffic destined to the egress interface for lossy and lossless priorities, few seconds of loss may be seen on lossy priority traffic on toggling PFC globally.  
Workaround: Stopping PFC pause frames before toggling PFC globally. (477941)
- Slice agent might restart repeatedly in scale QoS policy-map configuration with policers, which causes hardware policers to exhaust. The workaround is to reduce number of policy-maps with policers attached to them. (489846)
- Aggregate storm control rate limit is not honored for broadcast and multicast traffic. (490511)  
Series Affected: DCS-7010, DCS-7050X, DCS-7050X2, DCS-7060X, DCS-7060X2, DCS-7250X, DCS-7260X3 and DCS-7300X Series
- If an MBA supplicant is not assigned a dynamic VLAN when it is authenticated on a port, and if the supplicant later moves to a new port which is in unauthorized state, MBA authentication on the new port will not succeed and the port will continue to stay in unauthorized state. (491230)
- Programming or updating a Router ACL may rarely result in forwarding plane slice agent restart and can leave the corresponding vlan permanently blocked. (498502)
- MLAG connectivity between the peers might go down after linecard OIR. Flapping MLAG peer-link interfaces will bring the MLAG connectivity back up (500162)
- It is possible that storm-control configuration will still be applied in hardware after removing the configuration. Workaround is to restart the Strata agent. (504901)
- Interface default mtu may not apply properly on ports post reboot and cause IP fragmentation issues depending on packet mtu size. The workaround is to toggle the interface default mtu configuration. (515239)
- 'deny any any log` MAC ACL rule will not drop DHCP traffic and DHCP traffic will still be snooped. (523037)
- Enabling per VLAN routing with an L3 interface that only has an IPV6 address will result in no routing of packets received on this L3 interface. Workaround is to flap the interface. (528490)

- During an SSO failover, interfaces with VXLAN configured may result in an increased traffic outage on the associated interface (528749)

Series Affected: DCS-7300X Series

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

On fixed systems, the workaround is to run `platform trident reset`. On modular systems, the workaround is to run `platform trident <linecard> reset` on the affected linecard. (529148)

SKUs Affected: CCS-720XP-24Y6 and CCS-720XP-48Y6

- A mirror to GRE packet may incorrectly be tagged with a VLAN tag corresponding to the native VLAN if the next hop for the GRE tunnel is through a trunk port that is in an SVI. A workaround to this issue is making the GRE tunnel reachable through an access port if in an SVI or a routed port. (533903)
- Attaching a policer to class-default of an egress policy-map consisting of both IPV4 and IPV6 ACLs, is not supported and may lead to forwarding agent restart (534152)
- Enabling sFlow while in cut-through mode can cause corruption of data-transfer to CPU. The corruption may affect all traffic to CPU until the forwarding ASIC is reset. Workaround is to run in store-and-forward mode. (541940)

Series Affected: 7320X, DCS-7060X, DCS-7060X2, DCS-7260X and DCS-7260X3 Series

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

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

- Loop protect fails to detect loops on MLAG port channels. (559665)  
Series Affected: 7300X3, 7320X, CCS-720XP, DCS-7050CX3, DCS-7050SX3, DCS-7060CX, DCS-7060CX2, DCS-7060DX4, DCS-7060PX4, DCS-7060SX2, DCS-7260CX, DCS-7260CX3 and DCS-7260QX Series
- VXLAN rule in ACL applied on SVI will cause Strata slice agent to continuously restart. Workaround is to remove VXLAN rule from ACL applied on SVI. (559747)
- TCAM resource exhaustion can lead to VXLAN forwarding to fail.  
Workaround: Reduce TCAM resource utilization to allow VXLAN entries in TCAM by unconfiguring other features that occupy TCAM. (561231)

- When "show platform trident vxlan multihoming groups non-peering" is executed, StrataL3 agent may restart unexpectedly. There is no workaround to this issue. (561420)
- VXLAN BUM traffic received on a MLAG peer link may violate split horizon checks and result in sending traffic back to the VXLAN fabric resulting in MAC flaps in remote VTEPs.  
Workaround is to flap the MLAG peer link port channel member that is affected. (567667)  
Series Affected: 7300X3, CCS-720XP, DCS-7050QX2, DCS-7050SX2, DCS-7050SX3, DCS-7050TX2, DCS-7050X2 and DCS-7050X3 Series

## DCS-7010 Series

- 802.1x pause and guaranteed bandwidth cannot be configured at the same time. (91141)
- IPV6 traffic received over IPV4 tunnel interface may be dropped.

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

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

- If RPR or SSO is configured, Strata fabric agent restarts may cause switchover to occur. (368082)  
SKUs Affected: DCS-7304 and DCS-7308
- Packet buffer memory corruption can result in packets getting discarded without visibility in counters.  
When this condition occurs, the buffer usage reported by the "show platform trident mmu queue status" command exceeds the total number of buffers reported by the "show platform trident mmu buffers" command.  
A workaround is to run the "platform trident reset" command, which results in forwarding agent restart causing the links of the system to flap momentarily. (553299)

## DCS-7260X3 Series

- Forwarding agent may restart when a 40GBASE-SRBD transceiver is inserted after the hitless reload. To work around, configure "speed forced 40gfull" on the desired interface prior to inserting the transceiver. (417311)

## 7368 and DCS-7060X4 Series

- 802.3X pause is not supported (348758)
- L3 flooded packets do not preserve VLAN priority values. (365998)

- MPLS traffic will be dropped if the nexthop MAC address ages out or not known.

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

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

Series Affected: DCS-7060X4 Series

- Lossy traffic could be discarded at ingress interface instead of egress interface, if the ingress interface has "priority-flow-control mode on" configured. Because of this, the lossy priorities on the ingress port are vulnerable to head-of-line blocking. Buffer build-up on the lossy priorities is a necessary condition for the problem. (514254)
- Speed change on a port-channel member can cause a flap in BFD and BGP sessions for the entire port-channel (520638)

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

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

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

- Configured as MLAG, link flap of one or more LAG peer-link member ports could cause brief flooding of packets over the peer-link and can also lead to double delivery of packets on the MLAG interfaces. (117870)
- (A1 silicon only) MPLS traffic gets incorrectly prioritized, causing other traffic on the same port to get dropped. (160269)  
SKUs Affected: DCS-7050QX-32

- Altering the priority of transmit queues of a port will result in some traffic loss on that port. (194460)
- A single-event upset (SEU) on the tables REPLICATION\_FIFO\_BANK\* may not be corrected, which may lead to traffic loss. (391625)
- Sometimes, an interface can stop transmitting despite there being queued packets on the interface's egress queues. Doing a "shutdown" followed by "no shutdown" on the interface can resolve the problem. (400610)

## DCS-7050X2 Series

- When an interface which is transmitting at line rate is disabled and re-enabled, it might end up in a condition where it is unable to transmit any more packets. Disabling the interface again causes the Strata-Fix System agent to restart following which the interface continues to transmit properly. (245501)
- Packets sourced from an authenticated MAC on a Dot1X MBA authenticated port get punted to CPU. Unicast packets destined to authenticated MACs egress out tagged on a Dot1x MBA authenticated access port. Restarting StrataL2 will fix the issue. (532463)
- Vxlan multicast decapsulation will not work for packets that are sourced from unknown source VTEP. (538407)

## CCS-720XP Series

- IPv6 BGP peering may not be formed due to HLI=1 packets not reaching to CPU if
  - there is a speed change on any port,
  - there is a transceiver hot swap
  - IPv4/6 multicast routing is enabled/disabled.Workaround in case of first two cases is to toggle IPv4/6 unicast routing. For the third case, if IPv4 multicast routing is enabled, toggle IPv6 unicast routing, else if IPv6 multicast routing is enabled toggle IPv4 unicast routing. There is no workaround in case both IPv4/v6 unicast routing is enabled. (563045)

## DCS-7050X3 Series

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

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

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

The workaround is to reduce the number of 4x10G or 4x25G breakout connections to 2 (one breakout per following interface groups Ethernet49-52 and Ethernet53-56) and to configure 40G-only transceivers with the "speed

40g" interface configuration. After the configuration is applied, the forwarding agent must be restarted using the CLI command "platform trident reset". This will cause all interfaces to flap and a brief traffic outage to occur. (546409)

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

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

## Transceiver Series

- When performing an accelerated switch upgrade with the command reload fast-boot, some transceivers may experience a link interruption. This can occur on dual-speed QSFP100 optical modules that support both 100G and 40G, specifically when the module is configured to operate at 2x50G. Examples include the 100GBASE-PSM4 and the 100GBASE-CWDM4. (343684)
- Third party QSFP twinax cables may fail to initialize. This will result in them not showing up in 'show inventory'. Susceptible cables will have lower page 0 byte 2 bit 2 set to 0. This can be confirmed with the 'show idprom transceiver ethernet <intf>'. There is no workaround. (556437)

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

- Egress ACL on SVI for BUM traffic will not match ACL rules for packets being sent across fabric interfaces. (284543)
- When performing SSU, there may be failure in table programming resulting in fatal error cold reboot with extended traffic outage. (399774)
- Tagged packets with any VLAN priority received through MLAG peer links are always mapped to the same egress queue, and not complying with "show qos map". (407056)
- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (416694)
- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (417223)

- If a hitless forwarding agent restart is attempted while the hardware configuration does not match the running configuration, then it may cause an additional forwarding agent restart and loss of traffic. This may also happen if the forwarding agent is restarted while some interfaces have been disabled due to insufficient logical ports. (417835)  
SKUs Affected: DCS-7050SX3-48YC8-F and DCS-7050SX3-48YC8-R
- When "platform trident forwarding-table flexible" is configured, L3 unicast forwarding to hosts / L3 multicast forwarding may not work resulting in unexpected behavior. The workaround is to remove the CLI and use "platform trident forwarding-table partition" instead. (421084)
- Performing SSU may result in extended traffic outage, leading to total traffic loss being slightly over 1 second. (477139)
- Having any ACLs with UDF qualifiers as well as any of the following LAG/ECMP hashing features present in the configuration: RRoce hashing or TTL/HopLimit hashing, will result in incorrect operation of the LAG/ECMP hashing where the switch incorrectly hashes on IP payload. To avoid this problem the conflicting features should not be configured at the same time. (497016)  
Series Affected: DCS-7050CX3, DCS-7050SX3 and DCS-7300X Series
- Egress ACL rules can not be used to match ip protocol 50. (535104)
- Disabling static NAT access-list sharing with "no hardware nat static access-list resource sharing" may in some cases leak hardware VFP resources.  
  
Workaround is to reload the switch. (539951)
- Dynamic NAT access-list requires the subnets/hosts to be specified. It does not support ACL "permit ip any any".  
  
Possible workaround is to use deny ACL at the beginning to filter out traffic between private subnets and gateway that do not need translation. (546206)
- When Egress ip counter feature is enabled then running 'clear platform trident counters' command will cause slice agent to restart. System will recover after restart (546577)
- When the switch is reloaded with EVPN All Active config, the BUM traffic from some of the VTEPs may not be pruned correctly resulting in duplicate delivery of packets on some of the shared ethernet segments. The workaround is to restart the StrataL3 forwarding agent. (560034)
- VXLAN encapsulated packets received with inner payload having more than one dot1q tag will result in outer dot1q tag being stripped when egressing out on a port. There is no workaround to this. (567587)  
Series Affected: 7300X3 Series

- During bootup or L3 agent restart, one or more linecards may not be programmed with all the routes. This will cause traffic loss or misforwarding. There is no workaround for this issue. (568757)  
SKUs Affected: 7300X3-32C-LC

## DCS-7160 Series

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

In order to recover, remove the ACL configuration on ports/SVIs, and reattach them. (251114)

- On XP platforms, if XpMact agent restarts, and we have a large number of learned mac addresses and a large number of host routes, some of the host routes may get removed from hardware.

To recover, clear all learned mac addresses and restart XpMact. (255279)

- VXLAN routing & bridging traffic may fail to converge after Sysdb agent is restarted.

A workaround is to reboot the switch. (255957)

- In cases of high TCAM utilisation by Policy-based Routing, a 'configure replace' can result in Incorrect TCAM programming.  
In such cases, remove and re-apply the same policy-map which will result in programming the correct TCAM rules. (260906)  
Series Affected: DCS-7160 Series

- XPL3Unicast process may unexpectedly restart when the switch is reloaded if PBR policies with nexthop VRF actions are applied on an interface in the non-default VRF. (262713)
- IGMP IPv4 multicast control packets may create loop over a Vxlan Tunnel , if IGMP snooping is turned off in a VLAN. Workaround is to enable IGMP snooping in a VLAN ( which is enabled by default in EOS) (429933)
- Unexpected restart of the XpMact forwarding agent while there are unprogrammed MAC entries can cause some IPv4 and IPv6 unicast host routes to become unprogrammed. The fact that the routes are not programmed is only visible via "show platform xp ip(v6) route".

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

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

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

## Limitations and Restrictions in 4.23.7.1M

### Accelerated System Update

- During a hitless reload upgrading from a release before 4.21.3, more than 3 LACP packets could be transmitted within a one second interval. This should not cause problems. (338048)  
    SKUs Affected: DCS-7050SX-128-F, DCS-7050SX-128-R, DCS-7050SX-128-SSD-F, DCS-7050SX-128-SSD-R, DCS-7050SX-64-F, DCS-7050SX-64-R, DCS-7050SX-64-SSD-F, DCS-7050SX-64-SSD-R, DCS-7260CX3-64, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F and DCS-7260CX3-64E-R
- Aborting 'reload fast-boot' command with Ctrl-C is unsupported (344574)
- SSU of VRRP Master and Backup routers at the same time is unsupported and may cause extended traffic loss. Successfully SSU one router then SSU the other router. (345657)
- On systems with 4GB of /mnt/flash space, there may be insufficient space to store both an old swi image and a new swi image to perform an SSU upgrade. The workaround is to delete the old swi image from /mnt/flash prior to performing the upgrade. (385364)

### BGP EVPN L2/L3 VXLAN/MPLS

- On a BGP router, EVPN creates dynamic VLANs for L3 VNIs in receiving EVPN RT2 or RT5 advertisements even when the router does not have any L3 VRF or import route-target configurations. (294328)
- On a BGP EVPN device, if a VLAN-aware bundle is configured to include VLANs from multiple IP VRFs and if there are any IP addresses reused across multiple VRFs, then only one of the corresponding MAC/VLAN ARP bindings is distributed over EVPN. As a workaround, configure VLAN-aware bundles such that all VLANs in the bundle are in the same IP VRF. (514980)

### CVX

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

If a switch has more than 128 directly connected TaskTracker nodes with

MapReduce Tracer deployed, then the MapReduceTracer Agent can increase CPU utilization significantly. It is recommended to keep the directly connected TaskTracker count to below 128 when MapReduce Tracer feature is deployed. (80403)

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

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

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

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

- Intercept hosts on trunk ports with native vlan are only supported on 7050X2 platforms (257580)
- The Macro-Segmentation Service (MSS), requires MLAG fast redirection be enabled on the service MLAG TOR pair, when configured to work with a L2 transparent FW. This configuration minimizes traffic loss when any individual interface in the MLAG port channel goes down. (268291)
- The MssClient agent, can restart unexpectedly, when the MacroSegmentation Service, deployed with a L2-transparent firewall and is configured to intercept more than 10,000 end-points on a single switch (a number that's far greater than the total number of intercepts possible on the switch). (283221)
- The Arista OpenStack plugin has been enhanced to create a default route for the routers created in OpenStack, when VRF creation and vrf\_default\_route option is enabled in the networking-arista plugin configuration (295274)
- 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)
- If the Macro-Segmentation Service (MSS), working with a L3 firewall, is configured to intercept traffic from a host generating multicast traffic, the traffic can be black-holed.

To work around this issue, configure the following DirectFlow rule on all TOR switches MSS is programming rules on and assign it the highest priority:

```
flow bypassMulticast
    priority 65535
    table trident ifp
    match ethertype ip
    match destination ip 224.0.0.0 240.0.0.0
! (375156)
```

- MLAG devices in an environment with the Macro-Segmentation service (MSS) configured (with L3 firewalls), should have a high-priority DirectFlow rule configured to ignore traffic over the MLAG peer link. This can be configured using the following flow definition on each switch (assuming Po1 is the MLAG peer link):

```
flow bypassPeerLink
    priority 65535
    table trident ifp
    match input interface Po1
    match ethertype ip (375185)
```

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

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

- The MssPolicyMonitor agent might unexpectedly start running and eventually exit when the CVX functionality is disabled on the master CVX instance.

This does not have any side-effect on the operation of the Macro-Segmentation Service (MSS). (378077)

- When the Macro-Segmentation Service (MSS) is enabled and works with a L3 firewall, there can be a momentary traffic outage on reloading a MLAG peer if the CVX to VTEP (MLAG switch) connection is not re-established before the MLAG reload delay expires at the reloaded peer. (418560)
- If MSS policy enforcement rules configuration is changed from verbatim to group verbatim, the SandAcl agent can restart unexpectedly (433592)
- When Macro-Segmentation Service (MSS) is enabled on the CVX in an environment where no switch is configured as CVX client, multiple "excessive warmup delay" syslog messages for MssPolicyMonitor agent might be noticed. (497132)
- In a Macro-Segmentation Service (MSS) deployment with a PaloAlto firewall, enforcement policy configured with application using dynamic port assignment is not supported. (515749)
- In a Macro-Segmentation Service (MSS) deployment with a Palo-Alto firewall, a policy with multiple source or destination zones will be enforced over only one source zone and/or one destination zone.

A workaround is to split the policy into multiple policies with single source and/or destination zone. An alternative is to add the source and destination subnets to the policy configuration. (567376)

## General

- CLI does not allow access to files whose names contain spaces. (539)
- When VmTracer is used in a network that includes Cisco devices, ESX hosts connected to the Arista switch will see CDP frames from the Arista switch as well as from the other Cisco devices. VmTracer will display the VM info when the ESX host reports Arista's CDP info, but then will delete it when the ESX host reports other CDP info.

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

- EOS SDK-based applications must be run via the 'daemon' command configuration. They cannot be run directly from bash. (158431)
- EOS extension scripts and agents which do not use EOS-SDK will need to be modified to work in 4.16.6M and beyond. (158467)
- EOS swix extensions containing RPMs which were based out of or procured from Fedora distributions prior to Fedora 18 may fail to install due to

dependency issues. Suggested workaround is to recreate the swix files with the corresponding RPMs from Fedora 18 distribution. (162325)

- Packets with size greater than MTU of egress interface will not honor directflow flows and may not be forwarded (180927)
- SSO is hitful with subinterfaces. Protocols running on subinterfaces may experience session flaps during SSO. (188070)
- ACL drop counters cannot be cleared. (202905)
- when ACLs are configured with stats over port-channels, these stats are kept at the physical port level not at the port-channel level. (203131)
- With international and HW-REV-01 images and the introduction of "ip access-group <access-list name>" other configuration commands beginning with 'ip' that are configurable in the global config mode are not accessible in "router bgp" mode.

To access those commands exit the "router bgp" config mode to switch to the global config mode. (219390)

- Changes made to running-config after a config session is created might be reverted when the config session is committed which contains changes in similar or related areas. Use 'show session-config diff' inside the config session before committing to verify what changes will be made. (226850)
- The uptime in the output of "show module" may not match the uptime as shown by "show version" or "show uptime". The uptime in "show module" represents the time the module has been up since the last change in the "Status" column in the output of "show module", which differs from the meaning of uptime in "show version" and "show uptime". (248230)
- To avoid interoperability problems when using strong SNMPv3 encryption algorithms, follow the following minimums:
  - When using AES192 for privacy, use a minimum of SHA224 for authentication.
  - When using AES256 for privacy, use a minimum of SHA256 for authentication. (268290)
- Copy commands with sftp: or scp: URL do not work via eAPI by passing password through the "input" parameter. SSH keys have to be setup to run the commands without a password. (283716)
- EOS support alphanumeric VLAN names, however ODL gives an error when reading the leaf /interfaces/interface/routed-vlan/config/vlan if the name contains alphabetic characters.

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

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

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

- Enabling next-hop sharing (using "ip hardware fib next-hop sharing" command) could result in slower convergence in the case of link failures. (369305)
- For static tunnel interfaces, duplicate or conflicting GRE tunnels are tunnels that have the same source, destination, tunnel mode (GRE) and key configurations as an existing tunnel. A duplicate tunnel is always in the down state. When there are duplicate static tunnel interfaces and a config-replace is executed, it is non-deterministic which one of the conflicting tunnels will be in the up state. (373135)
- When an interface profile is used to configure an interface's access VLAN, the VLAN is not created automatically. The VLAN must be manually created with the vlan command. (376621)
- When using interface profiles, configuring interfaces may take a few minutes, depending on the number of interfaces being configured. It may be faster to apply an empty interface profile to the interfaces, and then update the profile. (382120)
- The CLI commands "locator-led [ multifan | powersupply | chassis ]" will not cause multifan, powersupply and chassis LEDs on the front panel of the system to flash, as they did on previous generations of fixed systems. The beacon LED on the back panel of the chassis is purple by default. The CLI commands "locator-led chassis" causes it glow red instead of expected blue. (395892)  
 SKUs Affected: DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3K-32P4-F, DCS-7280PR3-24-F and DCS-7280PR3-24-M-F
- The configuration lock feature does not prevent changes to the configuration via SNMP. (402556)
- If an EOS CLI command is used to create a reference to a VRF that does not exist, a subsequent OpenConfig YANG operation may cause the VRF to be created in EOS. (411908)
- Large amounts of monitoring traffic processed at the control plane via aggressive sflow sampling or mirroring to CPU could affect other control plane functions due to CPU bandwidth contention. (415992)

- IP Locking may send out additional queries when STP changes state (even possible in portfast config which is common for the IPLocking use-case). (417806)
- When hardware or software flow tracking is enabled and tracking flow is a VXLAN flow, the exported IPFIX packet may contain incorrect inner VLAN ID. (421437)  
Series Affected: CCS-720XP Series
- When "logging format hostname ipv4" is configured, log messages may not display the correct IP address if the IP address is updated. (422666)
- A range of 0.01 sec to 255 seconds is allowed for the YANG VRRP advertisement-interval value (range is 1..25500 as advertisement-interval is specified in hundredths of a second), but EOS only accepts values in the range 1 sec to 255 secs. (426391)
- In the 'show interfaces interactions' CLI, the list of speeds an interface is limited to does not account for the limitations that result from speed-group configuration. (434363)
- When we first configure a port-channel lag\_type, we must set values for both aggregation/config/lag-type and ethernet/config/aggregate-id (i.e. member configuration) in the same transaction. The lag-type is populated in EOS once the port-channel has one member port and remains set. (445421)
- When eAPI is configured with invalid SSL profile, nginx will stop running and all http/s requests will be refused (474422)
- Currently EOS does not support parsing and separating multiple ACL rules in the same NAS-Filter-Rule AVP using zero byte delimiter. (475210)
- Installed licenses without corresponding licensing features configured might not be visible from the CLI (504793)
- The YANG path '/network-instances/network-instance/mppls/global/config/null-label' is unsupported, but it is incorrectly listed as supported, and accepted by the OpenConfig agent. (514041)
- The following YANG paths are unsupported, but are incorrectly listed as supported and accepted by the OpenConfig agent.

```

/network-instances/network-instance/policy-forwarding/policies/policy/rules/
rule/action/config/decapsulate-gre
/network-instances/network-instance/policy-forwarding/policies/policy/rules/
rule/action/state/decapsulate-gre
/network-instances/network-instance/policy-forwarding/policies/policy/rules/
rule/action/config/discard
/network-instances/network-instance/policy-forwarding/policies/policy/rules/
rule/action/state/discard (515061)

```

- If Octa, OpenConfig or TerminAttr are started under low memory conditions, they can cause a spike in CPU usage before unexpectedly restarting. (534841)
- When the "tunnel destination" address is configured as a broadcast address for static tunnel interfaces then there could be packets drops during decapsulation at the tunnel endpoint. (544032)
- For a given interface, multiple simultaneous MAC address flush requests on different VLANs may result in clearing all MAC addresses on that interface. This is an optimization to support a higher scale of VLANs and interfaces. To disable this optimization, use the "mac address-table flushing interface vlan aggregation disabled" CLI configuration. (544712)

## Layer 1

- The MACSec outPktsEncrypted and outOctetsEncrypted counters are not supported with the MACSEC Proxy feature (281715)
- On Vxlan MACsec proxy, disable learning is not supported on proxy patch vlan. (284348)  
 SKUs Affected: 7500R2M-36CQ-LC, DCS-7280CR2M-30-F, DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R
- Ports configured at PAM4 speeds (50G-1, 100G-2, 200G-4, 400G-8) do not support link-debounce intervals shorter than 2 seconds. (500674)  
 Series Affected: 7368 and DCS-7060X4 Series  
 SKUs Affected: DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-96-F, DCS-7280CR3MK-32P4-F and DCS-7280CR3MK-32P4-R
- Interfaces with BCM82391 PHY do not support auto-negotiation. (537530)  
 SKUs Affected: 7500RM-36CQ-LC and DCS-7050CX3M-32S

## Layer 2

- LACP PDU fast rate ("lacp rate fast" on an Arista switch) should not be configured on any port in an MLAG pair, or any port connected to an MLAG pair. (13950)
- Static ARP inspection may not work on modular platforms after SSO switchover. The workaround is to disable and re-enable static ARP inspection after switchover. (244477)
- The command "lacp rate fast" is not supported on modular systems with stateful supervisor switchover (SSO) enabled. Neighboring devices should

not have "lacp rate fast" configured on ports connected to the said system. (248121)

- In a MACsec profile, when a primary key's CKN is changed, then operating primary continues to be principal actor even when a fallback key is configured till the time the new primary establishes a successful MKA session. There is zero traffic disruption. (275678)
- Port Security Protect does not support trunk ports. If a port is secured to source MAC <A> in VLAN <X>, traffic from source MAC <A> in other VLANs of the trunk port will also be allowed. (344773)
- While performing SSO, a device's peers might not receive LLDP packets within default TTL (120 seconds) and those peer devices may not maintain LLDP status information for the device undergoing SSO. After SSO is complete, the peers will again have LLDP state available.

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

- If 'traffic unprotected allow' is configured in a Macsec profile, agent restart or Stateful Switchover may result in traffic loss (362766)
- An MLAG switch cannot be used as an EVPN multi-homing provider edge. When a switch has MLAG enabled, MAC addresses from ports with an ethernet segment configured are advertised through BGP as if they were single-homing, and local ethernet segments are not advertised through BGP. (397867)
- Even with fallback to unprotected traffic configured, traffic loss in the order of milliseconds may be seen during initialization (491232)
- Upon programming a large number (~250) of VTEPs in a flood set at once, the L2Rib agent may restart. Workaround is to program them in batches, waiting a few seconds between each batch. (544051)
- 802.1X is not supported on LAG member ports. (544556)

## Layer 3

- ECMP is not supported for routes learned via RIP. (34773)
- DHCP relay is not supported when running virtual machines on EOS (101754)
- CLI does not detect TCAM exhaustion and automatically revert the change when a PBR policy is applied to an L3 interface in the "shutdown" state, since the actual hardware programming happens when the interface comes up. (122651)
- OSPFv3 support for multiple address families (RFC 5838) in EOS introduces support for IPv4 routing with OSPFv3. OSPFv3 for IPv4 AF in EOS requires

configuring neighboring OSPFv3 IPv4 routers on the same link with primary IPv4 addresses in the same subnet. Adjacency is established in OSPFv3 IPv4 AF even if the neighbor's primary IPv4 address is in a different subnet but routes through the neighbor on a different subnet will however not be installed. This limitation may be removed in a future release. (140234)

- When BFD is unconfigured with Fhrp as the only client and/or BFD is administratively shut, both VRRP routers will become Master. This transition is temporary and Master-Backup relationship will be restored. (159647)
- BFD sessions established between Arista switches and Cisco peers over IPv6 link-local addresses may flap when BFD echo mode is enabled. (159803)
- vEOS does not support BFD echo sessions. (160770)
- If "bfd per-link rfc-7130" is configured on port-channel(s) with primary and secondary IPv4 addresses configured in the same subnet, and the "bfd neighbor" is configured with peer switch port-channel's secondary IP address and there is a BFD session using primary IP address, then the port-channel might flap continuously.

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

- BFD echo functionality is not available on L2 Port-Channels with BFD RFC7130 mode enabled. (190418)
- If a port channel is configured with "bfd echo" and "bfd per-link rfc-7130" and is used for OSPFv3 or PIM, BFD sessions will not come up with echo functionality. BFD with SSO for OSPFv3 or PIM is not supported. (190997)
- SSO is not supported when BFD is configured for ISIS IPv6 sessions. (239704)
- An LACP port channel configured with BFD per-link and echo may experience BFD session flap on existing member links if new member link(s) are added to the port channel while the BFD peer device undergoes Stateful Switchover (SSO). (255042)
- Eos does not support programming IPv6 ECMP paths to kernel (258387)
- The DSCP value of BGP control plane packets set by the command 'bgp transport qos dscp <value>' will be overwritten when the platform 'qos rewrite dscp' is enabled and the outgoing interface is a SVI based on traffic-class to dscp map

present on  
the platform. (269764)

- In the output of the "show ipv6 dhcp relay counters" or "show ip dhcp relay counters" command, interface-specific DHCP relay counters might not sum up to "All Req" and "All Resp" counters. (278786)
- Tunnel interfaces flap when configuring ttl, tos and path mtu discovery on them. (281464)
- IPv6 addresses are not supported as tunnel source and destination. (283912)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289217)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289224)
- When BFD is configured with per-link mode rfc-7130 for an IPv6 link-local address on a Port-Channel and the Port-Channel transitions to the Down state, the Port-channel may not transition back to the UP state.

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

- On Macro-Segmentation Service (MSS) deployment with L3 firewalls, a flow entry that fail to get programmed due to lack of TCAM space, is not programmed when TCAM space becomes available at a later time. The workaround is to "shut" and "no shut" on the direct flow config. (372232)  
Series Affected: DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X and DCS-7060X2 Series
- The update wait-install feature may not work as expected if there are routes which fail programming due to hardware resource exhaustion. The routes which failed programming may still get advertised out to peers. (372765)
- DHCP packets will not be processed by EOS DHCP Server when DHCP Relay is relaying packet from one SVI where DHCP Relay is configured to another SVI where DHCP Server is configured. (389017)
- When a sampled packet egresses a trunk port and the path is ECMP, sFlow will arbitrarily pick one of the ECMP next hops for the next hop IP of the router extension. (415747)
- DHCP clients will fail to obtain IPv6 leases when "ipv6 dhcp relay destination" is set to IPv6 multicast address. (426796)
- IPv6 Egress ACLs may not be applied to traffic if 'ip hardware fib next-hop sharing' is also configured (446000)

- VXLAN features are not compatible the 'ip hardware fib next-hop sharing' configuration. Configuring VXLAN routes may result in identical FECs not being shared and potential misforwarding after changes to FECs are made. (477471)
- The config 'ip hardware fib next-hop multi-path maximum' does not restrict the size of installed next-hop group ECMP FECs (498484)
- CLI configuration of a static route with nexthop pointing to self ip address is not rejected. (514894)
- New NAT dynamic profiles are not supported on multi-chip and modular system. (525213)  
Series Affected: 7300X3, CCS-720XP and DCS-7300X Series
- BFD over Vxlan does not work for IPv6 underlay. (526897)

## Multi-agent

- If a BGP peer is configured for BFD fallover, the BFD session is not cleaned up even after the BGP peer is unconfigured. (217199)
- Traffic destined to an L3EVPN route whose underlay has ECMP BGP-LU routes, may experience temporary traffic loss when the underlay goes from n way to n-1 way or lower. The duration of loss can vary depending on the scale of BGP-LU routes or L3EVPN routes. (251692)
- When the multi-agent routing protocol model is configured, the maximum number of next hops for an OSPF ECMP route is 128. (274888)
- With very a very low multi-hop BFD timer configured, a BFD session may flap when a path (belonging to an ECMP set), over which the BFD session is established, goes down. (287571)
- Route-map specified as match-map for dynamic prefix-list is evaluated only against BGP routes. (345444)
- The IpRib agent runs out of addressable memory when a large number of routes are leaked to a large number of VRFs. (346575)
- Routes which are imported into a target VRF using a "leak routes source-vrf" policy cannot be matched on the basis of the source-protocol from the source-VRF as part of the "rib fib policy" (403739)
- When using 64-bit versions of EOS, Bgp, BgpCliHelper and Bmp agents may report a virtual memory size up to 50 GB larger than actual memory used. The actual memory usage is better represented by the sum of "RES" and "SHR" from "show processes top". (427364)

- When BGP router has advertised a large scale of routes (e.g. several million) and it reloads, there may be traffic drops while it is coming back up. This happens when the reloading router comes back up before the directly connected peering routers (which are receiving routes via an intermediate Route Reflector) get a chance to process millions of withdrawals and delete all the routes from FIB. (458127)
- IP ACLs are not supported for control plane route filtering for VRF leaking (536329)
- Match dynamic prefix-list clauses in route maps are only supported for BGP inbound and outbound route maps. (564552)

## Rib

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

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

## MLAG

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

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

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

## MPLS

- The MPLS agent may continuously restart on a router with high BGP LU Tunnel scale. (377939)
- BGP LU Tx MPLS routes corresponding to BGP LU advertisements with multiple labels are not supported. (395029)

## Multicast

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

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

- After ASU reload, if the first PIM hello packet received from the neighbor gets lost and If the neighbor has a high hello query-interval, it will result in traffic loss. Workaround is to use default PIM hello query-interval and increase the PIM hello query-count. (341447)
- With scale greater than 5000 multicast IPV6 routes, show platform sfe mroute ipv6 <vrf> commands might not display the right number of routes when performing multiple addition/deletion operations across VRFs. This is applicable only if user level multicast forwarding agent is configured to be used instead of kernel. (378841)
- Bessd might unexpectedly restart on config replace with scaled setup over 30 VRFs and large number of PIM enabled interfaces. Work around is to reduce the scale. (398185)
- Pim "non-dr install-oifs" feature becomes disabled if the non-dr needs to route multicast traffic for any reason to the interface on which the feature was originally configured. This is how the feature is expected to work. The workaround is to setup the network in a way that the DR always becomes the assert-winner. However, if the non-DR becomes assert winner for some routes on an interface and feature becomes disabled, flap the pim sparse mode configuration on non-DR on that interface to attempt to get out assert winner state. (401672)

## vEOS Router / CloudEOS

- The `/mnt/flash/cloud_ha_config.json` based configuration for the CloudHa agent is not supported from this release. If you are upgrading from an older image (< 4.20.5) please reconfigure HA using EOS CLI. Please see information on converting json file based config to CLI commands in the vEOS Router Configuration Guide. (264660)
- While vEOS instances in an AWS cloud can be created using either a Bring Your Own License (BYOL) or Pay As You Go (PAYG) format, there currently is no show command in vEOS to easily help the user or support determine which mode this instance uses. (272649)
- Tx/Rx ring parameters of an "Elastic Network Adapter (ENA)" interface cannot be set through config.json in vEOS. (276382)
- vEOS supports up a maximum of 8 physical cores. With hyperthreading enabled it is 16 vCPUs (278286)
- In dpdk based vEOS(MODE=sfe in `/mnt/flash/veos-config`), SR-IOV mode doesn't work with Intel x520/82599 on ESXi. As a workaround use mixed mode instead of the SRIOV only mode in ESXi (296718)
- Hot detaching network interface instances is unsupported. Reboot the instance after deattching a network interface. (306132)
- Site-to-site path monitored under "router path-selection" may experience flapping, when a lot of packets on the path (>10%) are dropped due to too much traffic injection into the virtual router. Workaround is to use QoS service policy to assign Inband Telemetry control packets (default UDP destination port is 4789) with a higher priority. (434087)
- CloudEOS does not support more than 10k /30 routes. The SFE agent will restart if more than 10k /30 routes are added to the system. The workaround is to reduce the number of /30 prefixes. (540204)

## DCS-7150 Series

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

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

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

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

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

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

- Systems with A0 silicon, packets are not forwarded to the egress port when egress truncation and timestamping are enabled together. Removing the configurations resumes packet forwarding. Truncation and timestamping are not supported together. (72756)
- The VXLAN flood VTEP list cannot be more than 128k VLAN-VTEPs, where VLAN-VTEPs is the number of VLANs multiplied by the number of VTEPs. (73228)
- Ingress port access control list (PACL) on a VXLAN-enabled interface is not supported. If ingress PACL is configured on a VXLAN-enabled interface, in-hardware VXLAN encapsulation of incoming packets is not expected to work correctly. VXLAN encapsulation of packets in software is not affected. Removing ingress PACL configuration on the affected interface restores correct VXLAN hardware encapsulation behavior on that interface. (76085)
- Multicast-routed frames get tx-mirrored even when they do not egress the tx-mirrored multicast destination due to VLAN suppression. This is a limitation of the switch ASIC. (87933)
- Enabling counters for Static or Twice NAT connection can cause FocalPointV2 agent restarts, when the number of connections exceed 275. (109048)
- On a DCS-7150 series switch, VXLAN encapsulating a packet post routing may result in the outgoing packet being corrupted if the packet was permitted by an ACL with "statistics per-entry" enabled. (294588)

- IPV6 hop-limit based filter in access-list is not supported even if accepted by CLI. (383400)  
Series Affected: DCS-7150 Series

## 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)
- 10/25/50 Gbps interfaces only support MTU up to 7Kbytes. (258823)
- After applying or removing shaping configuration on an L2 sub-interface, the L2 sub-interface must be flapped for the configuration to take effect. (281312)  
Series Affected: DCS-7170 Series
- After setting CoPP rate = 0, a few packets (<10) might be let into the CPU, but everything will be dropped afterwards. (350676)
- Full mroute scale may not be achievable under certain scenarios due to hash collision for the multicast route table. (380545)
- When egress NAT and MLAG are both used, 'show ip mroute' may show NAT'ed addresses in addition to the receiving mroute addresses. (394571)
- firewall profile on DCS-7170 platform does not support IP multicast (401134)
- VXLAN encapsulated ieee reserved mac addresses are not trapped to CPU and hence any control protocol that is VXLAN encapsulated will not work. (423963)

## DCS-7170 Series

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

## CCS-720XP Series

- "show environment power" does not show insertion status or telemetry information for PWR-1021-AC-RED or PWR-621-AC-RED when there is no input

power. (369368)

SKUs Affected: PWR-1021-AC-RED and PWR-621-AC-RED

- When hardware or software flow tracking is enabled, Vxlan encapsulated TCP flows may get exported with incorrect tcpFlags. The tcpFlags field track outer header's TCP flags hence for Vxlan packets it is undefined. There is no workaround to this issue. (416079)
- On receiving COA-REQ for changing VLAN for EAPOL supplicant, the NAS deliberately fails authorization for the first time in the new vlan. (449890)
- The per-interface configuration to ignore reauthorization timeouts from the AAA server ("dot1x timeout reauth-timeout-ignore") is not designed to work if manual reauthorization is requested from the command line. The reason for this is that manual reauthentication is designed to force the supplicant to contact the AAA server to retrieve the latest credentials. (453044)
- CoA request for VLAN change for EAPOL supplicant returns CoA-NAK. The VLAN does get changed when device is reauthenticated and the AAA server sends the new VLAN in an Access-Accept response. (459726)
- Hitless FPGA upgrades are not supported due to hardware limitations. Therefore FPGA images will not be upgraded during SSU. (474978)  
SKUs Affected: CCS-720XP-24Y6-F, CCS-720XP-24Y6-R, CCS-720XP-24ZY4-F, CCS-720XP-48Y6-F, CCS-720XP-48Y6-R and CCS-720XP-48ZC2-F
- Negation operator (!) and 'assigned' keyword are not supported in Radius NAS-Filter-rule( attribute-id 92) Attribute. (475207)
- Ip options qualifiers are not supported in the NAS-FILTER-RULE attribute if present in the Access-Accept or COA packet from the AAA server. If Ip options are present in NAS-FILTER-RULE, it may cause supplicant to fail authorization. (480659)
- When authentication caching is enabled and the link on a switch port is brought down due to a device being unplugged, it may take up to 2 minutes for the static mac address associated with the port to be removed from the mac address table. This is a grace period provided to give time for the device to be re-inserted into the same port. (489930)
- If a switchport loses power, the credentials of the attached phone and PC will remain cached when authentication caching is enabled. However, when power is

restored, the EAPOL enabled PC may not be reachable within the default reauthorization limit of 3 attempts and the supplicant will be removed as a result.

The solution in this situation is to increase the reauthorization limit using the

per-interface dot1x command: `dot1x reauthorization request limit 10` (500488)

- Limitation - 802.1X Web Authentication Implicit ACL feature doesn't work when IP Locking is configured on the interface. The workaround is to use 802.1X ACL based Web Authentication. (557986)

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

- On devices with Algomatch enabled, removing an ACL on an interface might fail because the remaining ACLs on the system may not fit in the hardware tables. (192811)

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

- 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)
- Mixed supervisor types are not supported on dual-supervisor modular systems. (450053)
- With scaled VRF, SVI and Arp entries, typically beyond claimed support, the Arp agent may experience a SIGQUIT during SSO switchover, on the new active supervisor. There is no workaround for this issue and switchover is ultimately successful. (495528)
- When 1000 VRFs are configured with 4000 SVIs, the Ira agent may experience a SIGQUIT during SSO switchover. There is no workaround for this issue and switchover is ultimately successful. (495529)

## CCS-720XP Series

- Polycom SoundStation IP 7000 phone might stop sending traffic after a switch power cycle if the phone stays on PoE power without an L1 link for some

time.

Workaround is to toggle PoE power on the affected interfaces with "poe disabled" followed by "no poe disabled". (502688)

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

- The L3 MTU on interfaces is not enforced on SVIs. (11659)  
Series Affected: DCS-7020, DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series
- Disabling IPv6 routing on an L3 interface is not supported when IPv4 routing is also enabled on that L3 interface. (43093)
- Linecards with HW revision < 02.00, when configured for 100G and a physical layer problem is present, a one-way link may result. The link experiencing the physical layer problem will be down. However, the peer link may still be up.

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

SKUs Affected: 7500E-12CM-LC and 7500E-72S-LC

- show interfaces counters rates' for port channels may show a slightly inaccurate data rate (including above 100%). (67367)
- Linecards with HW revision < 02.00, under extreme scenarios where we receive linerate small packets on all interfaces connected to the same forwarding ASIC and ingress mirror or sflow dangerous is turned on many of these interfaces, some packets (both the original and the mirror/sFlow copy) can be dropped. (67942)  
SKUs Affected: 7500E-12CM-LC, 7500E-36Q-LC, 7500E-48S-LC and 7500E-72S-LC
- A VXLAN encapsulated packet after de-encapsulation is not forwarded back to the port on which the original packet was received. Similarly, a native ethernet packet after VXLAN encapsulation is not forwarded back to the original receive port towards the remote VTEP. (75075)
- VARP and VRRP cannot be enabled at the same time. If VRRP is configured, there can be up to 128 combinations of interface and Virtual Router ID (VRID), but there can only be at most 8 different VRIDs and the VRIDs must differ only in the 3 least significant bits. VRRP for IPv4 and IPv6 cannot be configured at the same time. (75752)

- 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 running PTP with 4X10G break-out cables, a variable delay is observed with the PTP timestamps that can cause instability to the PTP time calculations. (186891)  
SKUs Affected: 7500E-36Q-LC
- When the 'vxlan-routing' hardware TCAM profile is configured, VXLAN traffic flows in overlay may not work alongside IPv6 traffic flows in underlay. The workaround is to use the CLI command 'no platform sand ipv6 host-route exact-match'. (190017)
- In Tap Aggregation mode, Tap Aggregation features may not work for 802.1BR/VN tagged packets when 802.1BR/VN tag stripping is not configured. (198798)
- Links are not compatible with systems running EOS 4.18.1F

The workaround is to upgrade both sides of the link. (199152)

SKUs Affected: 7500R-8CFPX-LC

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

Command "hardware phy cmx42550 <intf > diag read64 mac line 0x138" can be used to dump the counter for the purpose of debugging". (218429)

SKUs Affected: 7500R2M-36CQ-LC, 7800R3-48CQM-LC, DCS-7280CR2M-30-F, DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R

- Vxlan decapped packets will not match rules that match on 'tcp established' in egress acls (226582)
- Native VLAN setting on Tap interfaces does not take effect. (227770)
- While in Tap Aggregation mode, ingress VLAN membership filtering does not work when packets are destined to a tool interface configured with egress packet truncation. (227841)

- While in Tap Aggregation mode, interfaces configured for egress truncation will continue to consume hardware resources even when they are down. (228384)
- LLDP functionality is not supported on interfaces with ingress truncation enabled. (229983)
- Sflow agent CPU utilization is around 10% even when hardware acceleration is enabled and the agent is not processing any flow samples. (251674)
- With MTU increased beyond 9100B, there can be some multicast packets of size greater than 9100B that are dropped as a port gets close to linerate. (261446)
- Changing PMF profile on one linecard may cause a brief disruption of PMF-based features on other linecards which should not be affected by the same change. (270849)
- A QoS policy, with a class map which has either set-tc or set-dscp on the ingress port and which has DSCP rewrite map on the egress port, results in indeterministic behavior with respect to DSCP remarking of the outgoing packet. (273320)
- If a QOS and PDP Policy both having police action are applied on an interface and the traffic hits both the policies then QOS Policer Counter is messed up. The "show policy-map type qos pmap-name" will give the PDP Policer Counter, not the QOS Policer Counter. (275994)
- Rate values for subinterfaces shown by "show interface counter rates" may deviate up to 15% from corresponding rate on parent interface. (278700)
- With TCP MSS Clamping feature configured, TCP SYN packets might get dropped under high CPU usage conditions if the incoming TCP SYN packet rate is high. (279370)
- Packets larger than 10196 bytes will be discarded by the forwarding ASICs. (280459)
- PFC on DSCP trusted ports is not supported. (281214)
- Loop protect feature is not supported when AlgoMatch HW is enabled. (290706)
- L3 forwarding through SVIs is not support on L2 sub-interfaces. (293410)  
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- When a port becomes a LAG member, there may be a very brief moment where a few packets may be duplicated on the LAG. (295260)
- Fragmented IP packets will not match on flow-spec rules that include a match on TCP/UDP port. (297309)

- When all the L2 subinterfaces belonging to a parent interface are shutdown, packets are bridged and mac addresses are learnt on parent interface (300202)  
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- VXLAN routing is not supported on a modular system on which one of the linecards is configured for TAP aggregation. (306180)
- Show interfaces counters ip only show IPv4/Ipv6/MPLS counters per physical interface and not per port-channel (306209)
- When a BGP peer session is cleared, traffic which hits existing flow-spec rules may increment the wrong counter for a brief period of time. (339523)
- If non shaped sub interfaces are configured under same interface having shaped sub interfaces, traffic will egress out of the non shaped sub interfaces in a round robin fashion. It is recommended to not have shaped and unshaped sub interfaces under the same parent. (341851)  
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Hardware accelerated sflow with profile sflowaccel-v2 doesn't work with 7500E linecards. (342130)
- With uRPF configuration, packets sent by the Cpu that match a route forwarding into MPLS, GRE, or a Vxlan tunnel are incorrectly dropped in the hardware due to uRPF violation. (343891)
- Stale telemetry flow-spec statistics may be seen until the next "show flow-spec" command is issued. (345195)
- Filtered mirroring is not supported on linecards or fixed-systems operating in TAP Aggregation mode when traffic steering is enabled. (349883)
- Hardware accelerated Sflow does not support collectors reachable over arbitrary tunnel types (354749)
- IPv4 packets with options in the IP header will not be subjected to Unicast Reverse Path Forwarding (uRPF) checks, and hence will not be dropped even if the source IP is unknown. (355542)
- Software forwarding is not supported for GRE or IP-in-IP tunnel terminated packets using decap-groups. (358369)
- The egress DSCP value in the Outer IP header for packets that undergo headend replication after routing is derived from the Traffic Class (and not copied from the ingress DSCP value). (366189)
- While in Tap Aggregation mode, if the identity tag configuration of a Tap port is changed, traffic received on this port may be sent to an invalid

multicast destination momentarily. This will cause the DchUnreachables counter to increment. (372757)

- TX mirroring shows VXLAN encapsulated packets that are actually dropped by split horizon. (375795)
- If more than maximum supported Ipsec tunnels are configured, it is random as to which ones are assigned flow entries within on-chip encrypt/decrypt engine. The set of Ipsec tunnel interfaces assigned flow entries may not be the same as the set of tunnel interfaces which have IKE connections established. As a result, some tunnel interfaces with IKE connection established may not get link up. (376431)
- Different streams being rate-limited by the same policer might exhibit bias towards some streams and fair policing across all the streams is not guaranteed. (380599)
- Configuring more than 250 BFD sessions on a physical interface may result in many of those sessions flapping because of hardware drops on the interface. (389387)
- Drop Precedence based tail drop thresholds can be ignored when multiple queues are congested and buffering resources are low. (395693)
- If a security-ACL is configured on a sub-interface beyond the supported limit of acls, SandAcl might restart continuously until the number of configured acls is brought back within the supported limits (398815)
- Queue build-up doesn't show up for 64 byte packets in case of congestion via following commands:-
  - 1) platform fap diag diag cosq non\_empty\_queues
  - 2) show interfaces queue length

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

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

offset and vlan-tag-format qualifier doesn't differentiate between single and double tagged packets. (415587)

- When MPLS label is pushed on the traffic ingressing CoS trust port then EXP bits are not derived from TC based on packet's COS. (417108)
- Byte counts are not supported for IPSec tunnels. (419031)
- SflowAccel supports up to 200 VRFs. (423892)
- PAUSE frames with correct DMAC=01:80:C2:00:00:01, Ethertype=0x8808, opcode=0x0001 are never forwarded and are always consumed by the hardware. (426047)
- Header truncation for mirrored packets can not be used when the destination of a monitor session is a GRE tunnel. (428547)
- A port transmitting packets smaller than 100B that were tunnel terminated on the switch might be limited to ~90% linerate (429797)
- Priority Flow Control pause watchdog actions "errdisable" and "notify-only" actions are not supported. (450857)
- Interface-level Priority Flow Control (PFC) watchdog commands are not supported. (450860)
- IPv4/IPv6 counters are not supported on port-channel interfaces. (479658)
- When a LAG is used as a destination for a monitor session, a given set of packets mirrored from a Rx source will hash differently compared to the same set of packets captured from a Tx source, due to hardware limitations. Consequently, symmetric hashing will not work as expected with a monitor session which captures both Rx and Tx traffic from a single port, and using a LAG as a destination.

A workaround is to only capture traffic on the ingress (Rx direction). Another workaround is to use two different monitor sessions, and destinations, for the two directions of the target traffic, and implement symmetric hashing on a second stage. (486994)

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

- Egress MPLS tunnel counters are not supported for implicit null tunnels. (496375)
- When "ip hardware fib optimize" is enabled for a non-nibble aligned prefix-length, local-remote MPLS Pseudowires must have Control Word enabled for MPLS decap forwarding to work correctly. If Control Word is not enabled, then ETHoMPLSoETH traffic that ingress with the first nibble of the inner DMAC

starting with "0x4" will be speculative parsed as IPv4oMPLSoETH and incorrectly dropped. (496624)

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

- For Packets with GRE tunnel encaps header, MTU enforced on the hardware is 3 bytes more than configured value in Cli. (498470)

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

- With 'hardware next-hop fast-failover' configuration, the maximum ECMP size supported is reduced to 1023 in R, R2 series and 127 in R3 series switches respectively. (502940)
- VRF label termination into non-default VRF is not supported for multi-label MPLS packets. (504763)
- When BGP neighbors advertise Flowspec rules with police actions where the number of rules exceed what can be programmed in TCAM, policers will be allocated in hardware for all requested rules even though best-effort programming will only program the subset of rules that fit in TCAM. As a result, hardware policer resources will be unnecessarily consumed for rules which are not installed in hardware. (506688)
- With flex-route compress option enabled or the flex-route internet profile along with a hardware tcam profile including the flex-route feature, if the default route resolved over tunnel, it may lead to SandL3Unicast agent to restart continuously. (525611)
- Counter information flows from the forwarding chips to the tables that support TerminAttr in a very bursty manner. When averaged over a several minute window the rates will be accurate, but for shorter intervals calculated rates may oscillate significantly above and below the nominal rate.

For better accuracy, please use a longer interval to generate rate information. (540116)

- `show forwarding destination` may give incorrect results when the traffic would hit an egress IPv4 ACL using bridged IPv4 traffic or hit an egress IPv4 ACL using one of the following qualifiers: TCP established, fragment, and L4 port range. Note that TCP and UDP egress IPv4 ACL rules use the fragment qualifier. (554480)

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

- "shape rate" configuration is not supported on LAG subinterfaces used for MPLS pseudowires. (555799)

- Route churn under hardware FIB exhaustion may cause BFD sessions to flap. (556089)
- Egress packets greater than interface MTU (fragmented packets) are not counted in "show int counters ip".

Packet size should be limited to MTU to avoid fragmentation and loss of packet counts. (557785)

- For products with paired QSFP100 (odd numbered) and QSFP+ (even numbered) ports, the even numbered port may experience link flaps when a transceiver with a different media type from the previous installed transceiver is inserted in the odd numbered port. (565969)

SKUs Affected: DCS-7280QR-C72-F, DCS-7280QR-C72-M-F, DCS-7280QR-C72-M-R, DCS-7280QR-C72-R, DCS-7280QRA-C36S-F and DCS-7280QRA-C36S-R

- VRRP IPv6 using VRRP IPv4 MAC prefix may have unexpected packet forwarding behavior when the VRRP IPv6 and IPv4 have different states and one of the states is Master.

There is no workaround. (567504)

## DCS-7500E Series

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

## DCS-7280R and DCS-7500R Series

- An IPv6oMPLS packet terminated by L3 EVPN MPLS configuration will not be forwarded. The workaround is to disable IPv6 exact match host routes using 'no platform sand ipv6 host-route exact-match' and disable IPv6 route optimizations for prefix length 128. (201414)  
Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- When a linecard is inserted in a chassis there can be a temporary marginal drop in fabric throughput for other cards. (205880)
- On the system with DCS-7500R linecards, if DCS-7500R2 linecards are inserted, routes may consume more resources in the hardware routing table. (215393)
- Jericho chip buffering DRAM test at agent startup does not catch bad parts while it should fail to force RMA of bad parts. (218971)

- 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)
- The CPU traffic policy feature does not support filtering on IPv6 extension headers. (364996)
- Ingress ACLs and PBR are not supported for MPLS tunnel terminated traffic using VRF label. (366695)
- VRF terminated traffic is not sampled with sFlow. (369010)
- Drop-precedence thresholds are not supported on mirror session ports. (388932)
- We do not support incoming multicast LDP packets with a destination multicast MAC address.  
RFC 5332 doesn't mandate the use of sending labeled multicast traffic in a multicast ethernet frame. Multicast MAC destination address could be used in upstream label assignment. (393406)  
Series Affected: DCS-7020, DCS-7280CR2, DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500E, DCS-7500R and DCS-7500R2 Series
- MTU is not enforced for MPLS packets forwarded by a MPLS swap route, which have a length one byte more than MTU. (464742)

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

- "Drop" MAC entries only drop packets whose destination MAC address matches the "drop" MAC entry. Packets whose source MAC address matches the "drop" MAC entry are not dropped by the "drop" MAC entry. (341123)
- IP egress ACLs are not supported for MPLS terminated packets. (383263)
- Strict Priority scheduling for 400G interfaces don't work properly for 64 byte packets when traffic rate is above PPS limit of the chip. (388517)
- Cli to change tail-drop threshold for 200/400G ports is not supported by software. (396029)

- Applying a Qos policy on Vlan interfaces may result in SandAcl restart, since this functionality isn't supported. (403280)
- Applying a Qos policy on sub-interfaces may result in SandAcl restart, since this functionality isn't supported. (403281)
- Traffic drop may happen when we sent traffic to all 8 traffic-classes (aggregate rate nearing line rate) to port which has been configured with Strict-Priority scheduling. (406550)
- Traffic drop may happen when we sent traffic to all 8 traffic-classes (aggregate rate nearing line rate) to port which has been configured with WRR scheduling. (406561)
- Only port-channel interfaces are supported as peer links. Single member port-channel interfaces are supported. (409958)
- Per TC based VOQ tail-drop thresholds isn't supported in software. (413344)
- Application of low rate shaping on multiple interfaces with line-rate traffic into the box will create a backlog in DRAM. Due to this majority of traffic will be served to egress from slower DRAM instead of SRAM. In the congestion state, even after the removal of shaping with line-rate traffic ingress from multiple ports, switch won't be able to deliver full throughput. The recovery from DRAM to SRAM will happen over a period of time once the congestion is removed. (414481)
- Set dcsp will not apply to mpls routed packets that egress with mpls labels (418791)
- "ip tunnel <ip> nexthop-group" configuration for nexthop-groups with type mpls consisting of direct IP nexthops without label encapsulation is not supported. (458065)
- Unidirectional links are not supported at 40G speed on QSFP100 interfaces using the CRT50216 PHY. (467672)
- For VXLAN decapsulation + routing of overlay packet, the ECN bits may not be correctly copied from the incoming overlay packet. (489998)
- The "phy link serdes mapping direct" CLI command requires the 25g serdes rate to be configured, unlike what is displayed in "show hardware speed-group configuration". (523850)  
 SKUs Affected: 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3K-48CQ-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32P4, DCS-7280CR3K-96 and DCS-7280CR3MK-32P4
- Encapsulating PTP packets over VXLAN tunnel is not supported. (526461)

- When SandCounter agent is down, datagram counters are not accounted. Hence in show sflow output, the "Number of Datagrams" entry displays lesser value than expected. (536213)
- If all front-panel ports are populated, and traffic load is high, an improperly programmed current limit could be exceeded generating a spurious power-off. (536331)  
SKUs Affected: DCS-7280CR3-96 and DCS-7280CR3K-32D4
- Symmetric hashing is always on for IPv6 addresses, when both source and destination addresses are included in the applied load-balance profile. (557714)

## DCS-7020 Series

- The interface bin counter for both in and out packets with sizes 1523-Max may be inaccurate. Packets with sizes between 9217-9236 are not counted by the switch for interfaces Ethernet 1-48. (188121)
- Mirroring and sFlow can not be supported on the same source interface at the same time due to a hardware limitation. (209060)
- AES128/SHA-1 and AES256/SHA-256 are the only supported encryption/authentication algorithm pairs. (342005)  
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Using an Ipsec tunnel interface for policy-based routing nexthop is not supported. (378861)  
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Ipsec tunnel interfaces are not supported as part of nexthop-group. (378890)  
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Dynamic NAT, AKA Port Address Translation, is not supported for IPsec tunnels. (384624)

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

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

## DCS-7280R and DCS-7280R2 Series

- Flex-route is not supported with security ACLs (Port ACLs, Router ACLs & MAC ACLs) when AlgoMatch is enabled. The workaround is to configure 'hardware access-list mechanism tcam'. (348149)

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

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

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

- System does not stop transmitting traffic for the exact amount of time quanta received in PAUSE or PFC frame. This can lead to packet loss during congestion if the peer switch does not have enough buffers. The workaround is to use XON/XOFF instead of time based flow control, or to increase the time to account for one MTU sized packet. (27190)
- QoS shaping and guaranteed bandwidth will only work in store-and-forward mode. They are not supported in cut-through mode. This includes 'shape rate X' and the 'bandwidth guaranteed X' commands. (56790)
- Accelerated Software Upgrade (ASU) from 4.13.x (x < 6) to 4.13.6 and above requires a special upgrade procedure due to a software image format change. (90097)
- 10 Gbps packet replication on all SFP+ ports at the same time may lead to packet discard. (91149)
- Egress L2 MTU violations in cut-through mode may result in unexpected discards of other frames. (95577)
- VXLAN encapsulated packets cannot be TX mirrored at the egress properly. The mirrored packets will have an incorrect MAC header. (97272)
- Forwarding-table partition mode 4 is not supported. (100862)  
Series Affected: DCS-7010 Series
- When IGMP snooping is enabled a single unknown multicast floodset is shared across all VLANs with IGMP snooping enabled. This can cause unexpected flooding to a trunk port that has no multicast routers on a given VLAN, but has multicast router attached to the same trunk port on another VLAN. (105188)
- When "hardware access-list resource sharing vlan in" is configured, systems will start sharing resources that are required for programming access-group/

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

This affects ip access-list, ip standard access-list, and ipv6 access-list.  
ipv6 standard access-list  
is not impacted. (105498)

Series Affected: DCS-7010, DCS-7050X and DCS-7050X2 Series

- "hardware access-list resource sharing vlan in" configuration will not share resources required by IPv6 standard access-lists. (105502)  
Series Affected: DCS-7010 and DCS-7050X Series
- When "hardware access-list resource sharing vlan in" is configured, Link Local Multicast traffic will not be intercepted by the ACL attached to the VLAN interface. (105504)
- A large ACL and PBR policy configuration that consumes all TCAM resources may not fit after doing a config-replace, ACL agent restart or switch reload. After performing a config-replace, ACL agent restart or a switch reload, it is possible that only the ACL or the PBR policy will be able to fit. (105667)  
Series Affected: DCS-7010 and DCS-7050X Series
- There may be a momentary traffic disruption every time the nexthop specified in a PBR policy resolves to a new designation: affected packets may be forwarded by normal L3 lookup, or be dropped.

The system will recover automatically; the PBR policy will return to route packets as per the new nexthop resolution momentarily. (105670)

- When any PBR policy is applied, packets that violate the egress MTU and are destined to flood the VLAN will end up being flooded instead of being sent to the CPU. (105680)
- Subinterfaces cannot be used to send or receive VXLAN encapsulated packets. (105767)
- When "hardware access-list resource sharing vlan in" is configured, changing the ACL on a VLAN interface:
  1. From an ACL that is not shared with other VLAN interfaces to one that is shared with other VLAN interfaces can cause deny traffic to get permitted until the new ACL takes effect.
  2. To another ACL that doesn't fit in the TCAM can cause deny traffic to get

permitted until the system reverts back to the previously applied ACL.  
(106646)

- If IP-in-IP decap groups are configured, "qos trust dscp" configuration is not supported on traffic matching the decap groups. (107579)
- Vxlan and MPLS features may not be configured at the same time. (109330)
- When "hardware access-list resource sharing vlan in" is configured and TCAM is close to being full, restart of the forwarding agent or a reboot of the box, can cause some ACLs to not get programmed properly. (109876)
- The Accelerated Server Upgrade feature is not supported with the "hardware access-list resource sharing vlan in" configuration. (110114)
- IGMP snooping is not supported on Vxlan enabled VLANs. The multicast packets will always be flooded locally and to remote VTEPs. (110479)
- A front-panel or fabric port may get stuck in an errored state and drop traffic egressing the port. When this happens, the forwarding agent log will contain "port <internal port #> start error detected". This condition will persist until it is manually cleared. A forwarding agent restart is required to clear the error and forward traffic normally. (112051)
- MPLS is not supported on subinterfaces. (113110)
- Packets being VXLAN encapsulated are constrained to a single underlay physical nexthop per physical egress port.

Topologies involving SVIs over trunk ports, or non point-to-point routed ports towards the core will not work optimally, the encapsulated packets will be sent to the wrong next-hop for all but one of the VTEPs.

Topologies involving virtual VTEPs connected via a downstream L2 switch will not work, as the receiving vswitch will not have the capability to route the packet to the right VTEP. (113722)

- Port ACLs are not supported on VXLAN terminated traffic. (113726)
- If the configured RACLs do not fit in the TCAM, then disabling the RACL sharing feature might cause removing a RACL from an individual VLAN interface to fail.

To recover from this state, delete/remove the access-list and then reconfigure the access-list. (114028)

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

- L2 flooding of BUM packets on Vxlan vlans uses L3 replication tables, which are also used by IP multicast routing. The replication tables have an entry for each vlan, physical port combination (lag or non lag). We will run out of this table resources if we need more than 64k entries. (114517)
- Toggling "hardware access-list resource sharing vlan in" configuration could cause ACLs that were already programmed on VLAN interfaces to not get programmed. This could traffic loss on the VLANs where the ACLs were originally applied. (115348)  
Series Affected: DCS-7010 and DCS-7050X Series

- If the switch is in cut-through mode, L2 MTU violations will not be counted if the packet headers are changed while forwarding. (116760)
- QoS policies are not supported for L3 Unicast flooded packets (routed packets with nexthop MAC not learned) (123883)
- When a PBR is attached to any interface, Layer-3 packets which are getting flooded in the egress VLAN will not be treated with Egress Router ACL if one is configured on that egress VLAN. (133144)
- VXLAN routing is not supported on systems with BCM56750 fabric chips. Such switches can be identified using the CLI command "show platform trident l3 summary", which will show bcm56750 for "Fabric chip model". (134625)
- Mirroring of incoming multicast traffic to a GRE tunnel may result in data traffic loss on the egress interface in an oversubscription scenario. (139591)
- Egress IP/IPv6 router ACL is not supported for VXLAN encapsulated traffic. (148642)

Series Affected: DCS-7050X2, DCS-7050X3 and DCS-7300X Series

- IP ACLs configured on SVIs to match routed IP multicast traffic may also match bridged IP multicast traffic in that SVI. (152523)
- If hardware tables for VLAN translation overflows, the entries need to be un-configured and re-configured after the overflow condition is removed due to appropriate configuration changes. (157196)
- Up to 4 mirroring sessions are supported.

Each direction mirrored on a source port counts towards that limit of 4. For example, if a port is configured with rx and tx mirroring (the default), this counts as 2 consumed sessions. (158490)

- 1) Multicast traffic matching reserved IPv4 multicast address range 224.0.0.X is not counted by L3 interface counters.
- 2) L3 interface counters not supported for IPv6 multicast traffic.

- 3) Unicast traffic to CPU is counted by L3 interface counters only if routing is enabled (i.e. ip routing for IPv4 traffic and ipv6 unicast-routing for IPv6 traffic) (160930)
- Mirroring of Vxlan encapsulated packets in the egress direction is not supported. (167189)
  - If a trunk port is configured as a member of both the underlay and overlay VLANs, the following limitations apply:
    - If packets received on the trunk port are flooded in the overlay VLAN, the VxLAN encapsulated copy is not forwarded on the ingress trunk port.
    - Vxlan encapsulated packets received on the trunk port, if flooded in the overlay VLAN after decapsulation, are not forwarded on the ingress trunk port. (185321)Series Affected: DCS-7050CX3, DCS-7050SX3, DCS-7050X2 and DCS-7300X Series
  - If ingress ACL is applied on an interface, traffic sent to CPU from that interface is falsely counted as InAclDrop in "show int counters ingress acl drop". (203465)
  - Clause 37 auto-negotiation on SFP28 and QSFP28 ports is not supported. A workaround is to use 1G forced speed configuration on the affected ports. (206778)  
Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX3, DCS-7050X3, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7260CX3 Series
  - When "reload fast-boot" is performed from a release prior to 4.17.2, switches with BGP peering connections with the switch undergoing "reload fast-boot" could experience BGP KeepAlive timeouts before they see port flaps. This could result in longer than expected data plane downtime during the upgrade. (214700)  
Series Affected: DCS-7050X, DCS-7060X and DCS-7060X2 Series
  - A multi-bit ECC error in forwarding ASIC packet memory cannot be automatically corrected. A hitful restart of the slice agent managing the forwarding ASIC or a reboot of the system is required to clear the condition. (214950)
  - If ip-ttl is the only enabled field for port channel hashing, after reload hitless, packets may get hashed to a different link than before the reload. (223568)  
Series Affected: DCS-7050TX, DCS-7050X and DCS-7060X Series
  - A configured mirroring ACL may not be applied in hardware. This can be confirmed by checking the output of 'show platform trident tcam mirror'. (236908)

- Sflow output interface determination will not work when the output interface is a subinterface. (250424)
- A switch may not be able to bridge or route VXLAN traffic if it hits one or all of following conditions:
  - Nexthop resources are exhausted as indicated by syslog `VXLAN_HWHOP_NEXTHOP_RESOURCE_FULL`.
  - MAC table overflow as indicated by `show platform trident mac-address-table` missing.

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

- Flexible port-channel hashing may not hash accurately when the inner L4 header option is configured for IPV6 GRE packets. (267487)
- If a packet stream is policed by multiple policers, the policed rate may be lower than any of the configured policer rates. (271046)
- To perform filtered mirroring of a packet based on VLAN from an interface configured to perform VLAN translation, the ACL must be programmed with the translated VLAN. The mirrored packet will be the same as the packet at the source interface. (278599)
- BGP neighbor sessions may flap when URPF is configured on an interface. (279113)
- To perform filtered mirroring of a packet based on VLAN from a subinterface source, the ACL must be programmed with the internal vlan of the subinterface. The mirrored packet will be the same as the packet at the source interface. (280943)
- Matching on inner VLANs is not supported in mirror ACLs. If an entry is specified matching on inner VLAN, it will be ignored and permit all packets. (284371)
- SSU is supported when upgrading from the associated release in each release train: 4.18.8, 4.19.8, 4.20.7, 4.21.0. (289548)
- Issuing "no shutdown" on a member link of port-channel may send duplicate packets for sub-second on that member link (291514)
- DirectFlow (and features such as the MacroSegmentation Service that use DirectFlow) are unsupported on the 7260, 7300 series of switches. (296715)
- When primary vlan and secondary vlan are part of different MST instances, the hardware programming of lag will not happen. (300119)
- Not all interfaces in the groups Ethernet49-Ethernet52 and Ethernet53-Ethernet56 can be broken out in to 4x25G, 4x10G or 2x50G mode due to MAC

resource limitation on the system. Interfaces without MAC resource will be marked inactive and 'STRATA-4-HW\_PORT\_RESOURCE\_FULL' will be logged in syslog.

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

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

- When MLAG peer MAC routing is enabled, OSPF neighborship over VXLAN overlay may never get established. (349242)
- Ip-length based rules are not supported on Egress ACLs. (353247)
- Packets dropped in the egress pipeline may still be egress mirrored successfully. (357609)
- Some ports (esp. Et73-Et96 and Et103-Et104) might blackhole traffic or fail to linkup. (358362)

SKUs Affected: DCS-7050TX2-128-F and DCS-7050TX2-128-R

- SSU is not supported with Vxlan and MLAG config. (388669)
- Ports with ingress MAC ACL still learn MAC addresses although traffic may be dropped (400211)
- Reload hitless from EOS-4.22 or earlier results in /2 and /4 ports on Ethernet 1-12, 21-44 and 53-64 to remain inactive when /1 and /3 are configured in 10G or 25G. In order to activate /2 and /4 ports, configure 40G, 50G or 100G /1, then configure 10G or 25G back. These ports are also activated following a forwarding agent restart. (409057)

Series Affected: DCS-7260CX3 Series

- "reload fast-boot" is not supported with virtual IP address config. (411323)
- Speed change on a port with macsec enabled is not supported and could cause the port to get into a persistently un-authenticated state. Workaround is to disable macsec on the port, change the speed and then re-enable macsec on the port. (415737)

SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R

- Performing SSU from any release prior to 4.22.2 can lead to extended outage and mis-forwarding. (424812)

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

- IP ACL match over MPLS Encapsulated Packets is not supported. (426413)
- During Mpls PHP operation inner payload type gets ignored. (428546)
- When both "hardware counter feature vni decap" and "hardware counter feature vni encap" are enabled, the egress flexcounter pools will only be

released after both features are disabled. (441887)

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

- When an interface undergoes a hitless speed change, sFlow sample numbers and sequence counts for that interface may be reset. (452949)  
Series Affected: 7300X3, 7368, DCS-7050X3, DCS-7060X4 and DCS-7260CX3 Series
- With VXLAN dataplane learning disabled, locally learned MAC addresses will not age out if VXLAN encapsulated packets from remote VTEPs are received with inner packet's source mac as the locally learned MAC. The locally learned MAC entry needs to be cleared manually by issuing 'clear mac address-table dynamic vlan <VLAN> address <ADDRESS>'. (460607)
- If virtual mac address of the switch is learnt against a front panel port, it may not age out. The workaround is to clear the mac address manually by issuing 'clear mac address-table dynamic vlan <VLAN> address <ADDRESS>' command. (470965)
- When IPV4 and IPV6 PBR are both programmed in TCAM, forwarding plane agent restart may result in IPV6 PBR rules not being programmed in hardware. There is no workaround to this. (473582)
- The total length field in the IP header is not populated correctly when operating in cut-through mode. The switch should be operated in switch-and-forward mode when mirroring to a GRE tunnel. (475273)
- Mirrored packets that violate the MTU requirements will not be fragmented. (481513)
- Packets sent from the CPU cannot be egress mirrored unless they are being sent out of an SVI. (495146)
- Egress mirroring of multicast packets that are going to be encapsulated on the switch, independent of mirroring, is not supported. (512880)
- The mirroring logic can only capture and mirror packets seen by the ingress pipeline. Packets generated by the memory-management unit do not traverse the ingress pipeline and hence cannot be mirrored. This applies to all types of mirroring. (527668)
- When policer is attached to the class-default of an egress policy-map consisting of both IPV4 and IPV6 ACLs, IPV4 and IPV6 traffic hitting class-default would be policed individually with rate configured on default class-map (541843)
- The Strata agent restarts after running "platform trident diag psr <port>" command. (546588)

## DCS-7010 Series

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

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

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

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

- Cut-through mode is not supported on the SFP+ interfaces. (182728)
- Cut-through mode is not supported at 1G speed. (219427)

## DCS-7260X3 Series

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

## 7368 and DCS-7060X4 Series

- Multicast replication resources in TH3 are only 10% of their unicast counterparts. Due to this multicast replication in TH3 is limited to 1.2 Tbps. (297317)
- show running-config incorrectly displays interfaces for non-inserted cards. Interface statuses for missing or powered down cards are displayed as inactive. (306105)  
Series Affected: 7368 Series
- Layer2 source MAC learning may occur on packets received with CRC errors. (325507)
- Routed unicast traffic egressing a SVI for which destination MAC is not learnt are sent to CPU and then flooded on the egress VLAN. (368228)
- Per vlan routing and sub interface are not supported together. (486439)

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

- Packets sourced from the CPU to be VXLAN encapsulated will not have the right DSCP marking in the outer header based on the global QoS map configuration. (139248)

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

## DCS-7050X2 Series

- When IEEE 1588/PTP is configured on interfaces running at 100Mb/s speeds, the master offset and mean path delay calculations are inaccurate due to a limitation in the internal PHY (122816)
- Switch tries to perform VXLAN decapsulation of a packet if the destination IP of the packet matches the local Vtep IP address and the destination UDP port matches the configured VXLAN Port, even though the L2 Destination Mac address does not match any of its local Mac addresses. This can result in the packet getting dropped. This issue is more likely to be encountered when the switch happens to be an MLAG peer and the MLAG peers have the same

VTEP IP configured. The MLAG switch may drop the received VXLAN encapsulated packets addressed to its MLAG peer's Mac address when the destination IP matches the shared VTEP IP configured on the MLAG Peers. (145099)

Series Affected: DCS-7050SX2 Series

- Egress VLAN translation will not work on routed multicast packets when the VLAN to be translated is part of the outgoing interface list. (151093)
- VLAN translation is not supported on vxlan core ports (154309)

Series Affected: DCS-7050X2 Series

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

## Transceiver Series

- SFP-1G-T transceivers can only advertise one speed at a time. Desired speed and negotiation must be explicitly configured. Recommended commands are:

```
* "speed auto" / "speed auto 1gfull" / "speed auto 1000full" : enable autoneg, advertise only 1G
```

```
* "speed auto 100full" : enable autoneg, advertise only 100M
```

```
* "speed 100full" : disable autoneg, force speed to 100M (159401)
```

SKUs Affected: 1000BASE-T

- 25G transceivers could suffer signal degradation when using the MAM1Q00A-QSA QSFP-to-SFP adapter.

Use a MAM1Q00A-QSA28 QSFP28-to-SFP28 adapter instead. (253324)

- For SKUs CAB-O-2Q-400G, CAB-O-4Q-400G, CAB-D-2Q-400G and CAB-D-4Q-400G, when operating at 25G or 50G NRZ speeds, forward error-correction should be manually configured (enabled or disabled) on both sides of the link as the default error-correction on both ends may not match. (392124)
- Revision 20 OSFP-400G-2FR4 transceivers may not correctly report the operational value of the configured Tx Input Equalization. This can be observed in 'show interfaces transceiver tuning detail', in which the Operational Tx Input Equalization may not match the Configured value. If the 'Tx Input Equalization' Configured column shows auto, automatic Tx Input Equalization is enabled, despite the Operational value shown. (423045)

SKUs Affected: OSFP-400G-2FR4

- There is a hardware incompatibility between 100G AOCs with serial numbers starting with "AEB" and certain Mellanox adapters.

The recommendation is to RMA for AOCs with serial numbers that don't start with "AEB". (457261)

- Interfaces with BCM82391 PHY do not support 1000BASE-T transceivers. (535121)
- QSFP200 cables and optics encoded using the CMIS specification are not recognized in QSFP100 ports.

There is no workaround. (536318)

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

- The DirectFlow "set vlan action" command is not supported. (237046)
- When configuring a DirectFlow action, specifying "interface hardware-loopback" as the destination of "action output" is not supported. (243031)
- IPv6 packets with WESP payload (next header 141) bypass the IPv6 security ACLs and service policies. (243387)
- Cut-through mode is not supported on SFP+ interfaces. (252731)
- Directflow configuration with "set vlan", but no "action set output interface" configured is not supported. (345852)
- NAT is not supported on subinterfaces (353236)  
SKUs Affected: 7300X3-32C-LC, 7300X3-48YC4-LC, DCS-7050CX3-32S and DCS-7050SX3-48YC12
- 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)
- MPLS packet undergoing swap operation does not honor the MTU configured on ingress/egress interface. (422668)
- Only 2K MPLS routes are supported (425691)  
Series Affected: CCS-720XP Series
- If the number of MPLS next-hop groups configured exceeds 512, then STRATA-3-MPLS\_MAC\_DA\_PROFILE\_TABLE\_FULL syslog message is logged and L3 forwarding agent may continuously restart.

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

- If the packet ingresses on one line card and egresses on a different line card, then a truncated egress mirrored packets will be 4 bytes shorter than expected. (445200)  
SKUs Affected: DCS-7304 and DCS-7308
- 50g port speed is not supported in this release (445999)  
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- When IPV6 VXLAN underlay core interface is an access port on a SVI, VXLAN encapsulated packets may be sent out with VLAN tag in the outer ethernet header. The work around is to configure the port as a trunk port. (533929)

## DCS-7160 Series

- If there is a port ACL (PACL) or router ACL (RACL) that matches on L4 information for TCP/UDP/ICMP, and if the incoming packet has IP options, then those packets are dropped. (172932)
- No support for matching on DSCP field on IP packets for security ACLs. (174114)
- If the MAC address of the next-hop to reach a remote VTEP is not learned, and if the remote VTEP is configured for Head End Replication (HER) for a VXLAN VLAN, then, the the VTEP is not included in the list of HER VTEPs and no VXLAN BUM packets will be sent to that VTEP. The VTEP is included in the list of HER VTEPs once the underlay MAC address of the next-hop to reach that VTEP is learned.

The workaround is to ensure that the underlay MAC of the next-hop to reach remote VTEPs are always present (which is the normal operation). (178395)

- Control plane Policy Map counters are not supported. (180303)

- A maximum of 3840 remote VTEPS are supported when the switch is configured as a Vxlan Vtep. (188553)
- ACLs with Vxlan is not supported, and may result in unexpected forwarding behavior depending on the configuration and pipeline profile. (217888)
- Any IPv6 ACL containing a rule with the first 96 bits set as 0 cannot be configured when the algomatch profile is being used. (218453)
- Customer may witness around 100 mililiseconds of packet loss when VxLAN ECMP nexthop changes. (219888)
- There may be traffic loss of up to 100 mililiseconds when a VxLAN nexthop (as part of ECMP) undergoes modifications. (268600)
- The DSCP value of BGP control plane packets set by the command 'bgp transport qos dscp <value>' is overwritten when 'qos rewrite dscp' is enabled and the outgoing routed interface is a SVI. The rewritten value is based on traffic-class to DSCP map configuration. (288617)
- If a different XP profile than the one that is currently running is configured, and ACL configuration/ACL viewing commands without running a chip reset/reboot (via platform xp xp0 reset), the command can cause XpAcl to restart. (418083)
- VLAN ID match in IP ACLs is not supported on DCS-7160 platforms. (514366)
- If an incoming non-routable VXLAN-encapsulated packet is decap-eligible, but the encapsulation does not conform to VARP VTEP - VARP MAC binding, the packet will not be dropped and will continue to be switched based on the inner packet header. The impact is limited to extra or duplicate packets in the destination bridging domain, and does not affect packet forwarding correctness. However, it indicates an issue with the sender switch which is potentially generating illegal VXLAN-encapsulated packets. (562335)

## Conditions and Impacts

The release notes listed above use shorthand terminology to refer to conditions that arise frequently in connection with bugs. This section explains each condition and its impact in more detail.

**Condition:** Rib agent restart

**Impact:** When the Rib agent goes down, all routing sessions are terminated; all BGP sessions drop, all OSPF adjacencies are lost, and neighbors stop forwarding

traffic to us. When the Rib agent restarts, adjacencies are re-formed, and, after protocol convergence, traffic flows through us again. In most cases where there is adequate network-level redundancy, application-visible impact should be minimal

**Condition:** Rib agent hang

**Impact:** When the Rib agent hangs, routing protocols stop running. Routing peers will generally time out their session with the hung Rib agent, withdrawing routes accordingly. Meanwhile, the device continues to forward packets based on existing routing state. During this time, the device will not respond to routing topology changes. After a heartbeat timer expires (typically 10 minutes), the Rib agent restarts. In networks with sufficient redundancy, application impact of a Rib agent hang is usually low, because there is no data path disruption. However, in conjunction with other network changes (such as link failure or introduction), routing loops may occur.

**Condition:** kernel crash

**Impact:** A kernel crash has impact similar to issuing the "reload now" command. All links go down and all forwarding stops. Then, the system reloads, which may take up to 15 minutes. In a network with adequate redundancy, there should be minimal traffic loss associated with this period. After reload, links come up and protocols (LACP, STP, BGP, etc) reconverge. Protocol reconvergence is not necessarily hitless, depending on topology and configuration. For example, a switch may advertise a prefix to a connected subnet before STP has converged. However, any associated outage should be short lived, typically lasting around 30 seconds. The MLAG-SSO feature can dramatically reduce the window of disruption.

**Condition:** forwarding agent restart

**Impact:** A forwarding agent restart typically results in the switch ASIC being reset. It clears packet memory, dropping all packets currently in the switch, and causes all links to go down. The restart can take up to 45 seconds, during which time all links are down. Once the restart completes, all links come back up automatically, followed by protocol reconvergence (MLAG, LACP, STP, BGP/OSPF/IS-IS, etc). Depending on which protocols and options are in use, the reconvergence may take a few seconds up through several minutes. On modular switches, the impact of forwarding agent restart is limited to the affected line card; that is, if the forwarding agent for line card 3 restarts, then all ports on line card 3 bounce and protocols on those ports reconverge, but ports on other line cards are unaffected.