

Arista Networks EOS 4.24.2.3F Release Notes

Version: 0.1

30 September 2020

Table of Contents

EOS 4.24.2.3F Release Highlights

New Platforms and Hardware

SNMP MIBs

SNMP Traps

Supported Hardware

Reference to MLAG ISSU Upgrade/Downgrade Table

Resolved Caveats in 4.24.2.3F

Layer 3

Multi-agent

DCS-7170 Series

Known Software Caveats in 4.24.2.3F

Accelerated System Update

Containerized EOS

CVX

General

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

MPLS

Multicast

vEOS Router

DCS-7170 Series

DCS-7170 Series

CCS-720XP Series

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

DCS-7300 and DCS-7500 Series

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

DCS-7280R and DCS-7500R Series

DCS-7280R2 and DCS-7500R2 Series

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

DCS-7020 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

Limitations and Restrictions in 4.24.2.3F

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

DCS-7170 Series

DCS-7170 Series

CCS-720XP Series

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

DCS-7300 and DCS-7500 Series

CCS-720XP Series

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

DCS-7280R and DCS-7500R Series

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

DCS-7020 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 4.24.2.3F 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
- IEEE8021-SECY-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-7170-64C
- DCS-7280SR2A-48YC6
- DCS-7050TX-64-F

- 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
- 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
- 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-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-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
- 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

- 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
- PWR-1011-AC-BLUE
- PWR-1011-DC-BLUE
- PWR-1011-DC-RED
- PWR-2411-AC-RED
- PWR-2411-DC-RED
- FAN-7012H-BLUE
- FAN-7012M-BLUE
- FAN-7012M-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-7280CR-48
- DCS-7280CR-48-F
- DCS-7280QR-C36
- 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-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-7280DR3K-24
- DCS-7280DR3K-24-F
- DCS-7280CR3MK-32D4
- DCS-7280CR3MK-32D4-F
- DCS-7280CR3MK-32D4-R
- DCS-7050TX-128
- DCS-7050TX-72
- DCS-7050TX-96
- 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-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
- 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
- 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-7050CX3M-32S
- DCS-7050CX3M-32S-F
- DCS-7050CX3M-32S-R
- CCS-720XP-24Y6-R

- DCS-7280QR-C36-F
- DCS-7280QR-C36-M-F
- DCS-7280QR-C36-M-R
- DCS-7280QR-C36-R
- DCS-7280SR-48C6
- DCS-7280SR-48C6-F
- 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-7250QX-64
- DCS-7250QX-64-F
- DCS-7250QX-64-R
- DCS-7250QX-64-SSD-F
- DCS-7250QX-64-SSD-R
- DCS-7050QX-32
- DCS-7050QX-32-F
- DCS-7050QX-32-R
- DCS-7050QX-32-SSD-F
- DCS-7050QX-32-SSD-R
- DCS-7050QX-32S
- DCS-7050QX-32S-F
- DCS-7050QX-32S-R
- DCS-7050QX-32S-SSD-F
- DCS-7050QX-32S-SSD-R
- DCS-7050SX-64
- DCS-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
- CCS-720XP-48Y6-R
- DCS-7050SX3-48C8
- DCS-7050SX3-48C8-F
- DCS-7050SX3-48C8-R
- DCS-7050TX3-48C8
- DCS-7050TX3-48C8-F
- DCS-7050TX3-48C8-R
- CCS-720XP-96ZC2
- CCS-720XP-96ZC2-F
- DCS-7050SX3-96YC8
- DCS-7050SX3-96YC8-F
- DCS-7050SX3-96YC8-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
- PWR-2700-DC-F
- PWR-2700-DC-R
- PWR-3K-DC-BLUE
- PWR-3K-DC-RED
- PWR-3KT-AC-BLUE
- 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
- DCS-7508
- DCS-7504
- 7300-SUP
- 7300-SUP-D
- 7300X-32Q-LC
- 7300X-64S-LC
- 7300X-64T-LC
- 7304X-FM
- 7308X-FM
- DCS-7304
- DCS-7308
- 7316X-FM
- DCS-7316
- 7320X-32C-LC
- 7324X-FM
- 7328X-FM

- PWR-3KT-AC-RED
- PWR-D1-3041-AC-BLUE
- PWR-2900AC
- DCS-7504N
- DCS-7508N
- DCS-7500-SUP2
- DCS-7500-SUP2-D
- 7500R-36CQ-LC
- 7500R-36Q-LC
- 7500R-48S2CQ-LC
- 7504R-FM
- 7508R-FM
- 7512R-FM
- DCS-7512N
- 7500R3K-36CQ-LC
- 7500R3-24P-LC
- 7800R3-36P-LC
- 7800R3-48CQ-LC
- 7800R3-48CQM-LC
- 7800R3K-48CQ-LC
- 7808R3-FM
- DCS-7800-SUP
- DCS-7808-CH
- 7804R3-FM
- DCS-7804-CH
- 7500R3-24D-LC
- 7800R3-36D-LC
- 7300X3-32C-LC
- 7304X3-FM
- 7308X3-FM
- 7300X3-48YC4-LC
- 7368
- 7368-16C
- 7368-SUP
- 7368-SUP-D
- 7368X4-SC
- 7300-SUP2
- 7300-SUP2-D
- 7368-4D
- 7368-4P

Transceiver

- 1000BASE-BX10
- 1000BASE-BX10-D
- 1000BASE-BX10-U
- 40GBASE-AOC
- 10GBASE-LRL
- 10GBASE-ZR
- 40GBASE-PLRL4
- 40GBASE-SR4
- 40GBASE-XSR4
- 40GBASE-LR4
- 40GBASE-LRL4
- 40GBASE-PLR4
- CFP2-100GBASE-LR4
- 10GBASE-AOC
- 40GBASE-UNIV
- 100GBASE-CR4
- 100GBASE-LR4
- 100GBASE-LRL4
- 40GBASE-ER4
- 10GBASE-DWDM
- 10GBASE-DWDM-T
- 10GBASE-DWDM-ZR
- CFP2-100GBASE-ER4
- 100GBASE-PSM4
- 100GBASE-SR4
- 40GBASE-SRBD
- 100GBASE-AOC
- CFP2-100GBASE-XSR10
- 100GBASE-CWDM4
- 25GBASE-AOC
- 25GBASE-SR
- 100GBASE-SRBD
- 100GBASE-SWDM4
- 100GBASE-ERL4
- 100GE-DWDM2
- CFP2-100G-DWDM-DCO
- CFPX-100G-DWDM
- CFPX-200G-DWDM
- 1000BASE-LX
- 1000BASE-SX
- 1000BASE-T
- 10GBASE-CR
- 10GBASE-LR
- 10GBASE-SR
- 10GBASE-SRL
- 100GBASE-XSR4
- 100GBASE-DR
- 10GBASE-ERBD
- 10GBASE-ERBD-D
- 10GBASE-ERBD-U
- 10GBASE-ERLBD
- 10GBASE-ERLBD-D
- 10GBASE-ERLBD-U
- ADPT-O-Q-100G
- AOC-O-O-400G
- CAB-D-2Q-200G
- CAB-D-2Q-400G
- CAB-D-4Q-200G
- CAB-D-4Q-400G
- CAB-D-8S-200G
- CAB-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
- 100GBASE-LR
- SFP-10G-T
- ACC-D-4QN-400
- ACC-O-4QN-400
- OSFP-400G-FR4
- OSFP-400G-XDR4
- QDD-400G-DR4
- QDD-400G-FR4
- QDD-400G-SR8
- QDD-400G-XDR4
- 100GBASE-XCWDM4
- SFP-25G-MR-LR
- SFP-25G-MR-XSR
- SFP-25GR-LR
- SFP-25GR-XSR
- 200GBASE-FR4
- 200GBASE-SR4
- 10GBASE-ER

- 25GBASE-CR
- 25GBASE-LR

- CAB-D-D-400G
- CAB-O-2Q-200G

- 40GBASE-CR4

Reference to MLAG ISSU Upgrade/Downgrade Table

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

Resolved Caveats in 4.24.2.3F

Layer 3

Multi-agent

- If aggregates are configured and the number of contributing routes to the aggregate changes from a high number to a low number, then the shared memory partition of the filesystem may fill up, leading to a low memory condition. Agents may unexpectedly restart, and the system may experience unexpected reload.

The workaround is to unconfigure aggregates. (521460)

DCS-7170 Series

- Configured non-default VXLAN UDP destination port is not effective on reload. The workaround is to unconfigure and re-configure the desired VXLAN UDP port after reload of the system. (524060)

Known Software Caveats in 4.24.2.3F

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 running 'reload fast-boot' or 'reload hitless', ports with optical transceivers may flap, causing traffic disruption (348172)
SKUs Affected: DCS-7050SX-128-F, DCS-7050SX-128-R, DCS-7050SX-128-SSD-F, DCS-7050SX-128-SSD-R, DCS-7050SX2-128-F, DCS-7050SX2-128-R, DCS-7050TX-48-F, DCS-7050TX-48-R, DCS-7050TX-64-F, DCS-7050TX-64-R, DCS-7050TX-72Q-F and DCS-7050TX-72Q-R
- When performing SSU on an MLAG setup, SSU can cause peer links on port-channel to flap, resulting in large traffic loss. (405788)

- sFlow may not work properly after ASU.
Restarting SandFap agent after ASU will fix the issue. (490542)
SKUs Affected: DCS-7280SR-48C6-F, DCS-7280SR-48C6-M-F, DCS-7280SR-48C6-M-R, DCS-7280SR-48C6-R, DCS-7280TR-48C6-F, DCS-7280TR-48C6-M-F, DCS-7280TR-48C6-M-R and DCS-7280TR-48C6-R

Containerized EOS

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

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)
- The OpenStack agent may use the wrong region's endpoints for name resolution in multi-region deployments

A workaround when using CVX to manage only a single region is to redirect queries to the correct endpoint with 'ip host <wrong endpoint> <correct endpoint ip>'. (503014)

- In a Macro-Segmentation Service (MSS) deployment, a firewall REST API providing a malformed response may cause enforced policy to be withdrawn and re-applied. (512505)

General

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

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

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

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

- 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-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-7280SR-48C6-F, DCS-7280SR-48C6-M-F, DCS-7280SR-48C6-M-R, DCS-7280SR-48C6-R, DCS-7280SR2-48YC6-F, DCS-7280SR2-48YC6-M-F, DCS-7280SR2-48YC6-M-R, DCS-7280SR2-48YC6-R, DCS-7280SR2A-48YC6-F, DCS-7280SR2A-48YC6-M-F, DCS-7280SR2A-48YC6-M-R, DCS-7280SR2A-48YC6-R, DCS-7280SR2K-48C6-M-F, DCS-7280SR2K-48C6-M-R, DCS-7280SRA-48C6-F, DCS-7280SRA-48C6-M-F, DCS-7280SRA-48C6-M-R, DCS-7280SRA-48C6-R, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F, DCS-7280SRM-40CX2-R, DCS-7280TR-48C6-F, DCS-7280TR-48C6-M-F, DCS-7280TR-48C6-M-R, DCS-7280TR-48C6-R, DCS-7280TRA-48C6-F, DCS-7280TRA-48C6-M-F, DCS-7280TRA-48C6-M-R, DCS-7280TRA-48C6-R, DCS-7304, DCS-7308, DCS-7504N, DCS-7508N, DCS-7512N and DCS-7516N

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

Example:

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

```
route-target import auto 302
```

Suggested Usage

```
router bgp 301
```

```
  vlan 300
```

```
    # In this case 301 is used as the AS number for generation of route  
target
```

```
  route-target import auto (255062)
```

- The MacroSegmentation Service (MSS) feature doesn't work in conjunction with the VARP VTEP IP configuration required for VXLAN routing where a subset of VTEPs are L2 VTEPs. (263532)
- 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)
- 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-7280R, DCS-7280R2, 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)
- 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)
- 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-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- Cli may throw an exception if interface group name or unit name starts with letter 'L'. Workaround is to use any other name not starting with letter 'L'. (486400)
- 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)
- Large PTP Offset From Master value fluctuations occur following PTP link failover resulting in disruption to downstream PTP clients. (499397)
- On a device with a high rate of matches against configured STRE policies, the resulting frequent counter update rate can cause OpenConfig to restart unexpectedly. (504297)
- CLI sessions may restart unexpectedly. (506350)
- PTP Sync messages may be buffered and delayed between 6 to 12 seconds. This may cause downstream clients to lose PTP time synchronization. (508551)
- ConnectivityMonitor reports packet loss as 0 instead of 100 in case of a probe error which includes unreachable hosts.
 A workaround for detecting a probe error is to examine the latency statistic which is zero in the case of an error. (512126)
- PTP agent may unexpectedly restart when port-channel links change its status from up to down and vice versa. (519578)

Layer 1

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

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

- 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

- 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)

- Changing interface speed from forced to auto-negotiated or vice versa may cause the forwarding agent to restart and all interfaces to flap. (491629)

Series Affected: DCS-7060X4 Series

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

- QSFP ports with CRT50216 PHY do not support link-debounce intervals shorter than 5 seconds. (498871)

SKUs Affected: 7368-16C, 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-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32P4-F and DCS-7280CR3MK-32P4-R

- Changing interface speed configuration from "speed 50G-1" to "speed auto 50G-2" or "speed 100G-2" to "speed auto 100G-4" causes an unexpected forwarding agent restart. The workaround is to configure "speed 50G-2" or "speed 100G-4" before executing "speed auto 50G-2" or "speed auto 100G-4"

accordingly. (500511)

Series Affected: DCS-7060X4 Series

- Interfaces configured with 50G-1, 100G-2, 200G-4 or 400G-8 auto-negotiation enabled may fail to link up following a link flap or an interface configuration change. The workaround is to configure 50G-1, 100G-2, 200G-4 or 400G auto-negotiation disabled, followed by auto-negotiation enabled of the same speed. (522165)

Series Affected: DCS-7060X4 Series

- Configuring "shutdown" or changing speed on an interface with MAC loopback ("traffic-loopback source system device mac") enabled may cause the forwarding agent to unexpectedly restart. This causes all interfaces to flap and a brief traffic outage to occur. (523006)

Series Affected: 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 a SSO switchover may result in STP topology reconverging. (470962)
- Restarting STP agent with a high number of interfaces and vlans in PVST mode may result in STP topology reconverging. (472675)
- In VXLAN deployment, remapping the VNI from one VLAN to another VLAN may cause the destination VTEP to drop ARP requests received over VXLAN tunnel in that VNI. (482403)

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)
- The Macro-Segmentation Service (MSS) for Layer 3 firewalls may not correctly program traffic redirection/offload flows if the service is enabled prior to DirectFlow being enabled on all switches.

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

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

- 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. (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 crash. (414702)
- 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)
- 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)
- 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)
- In scale situations, if a static ARP entry is added immediately before a route is deleted and re-added, the entry may not be properly added. Workarounds include removing and re-adding the static entry, flapping the interface the entry end up on, or restarting the Arp agent. (486998)
- TI-LFA backup paths will be computed slowly when there are large number of node-segments/anycast-segments/adjacency-segments that are protected. (489097)

- If a configuration change involving deletion of several VLAN interfaces along with anything which causes platform layer 3 agent to restart is applied, then the layer 3 agent may see multiple restarts following the first restart. (492891)
- 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)

- If an IP route resolves over ECMP of SR-TE policies containing multiple segment-lists using TI-LFA backup paths, forwarding agents may not properly failover to the backup path. This can happen if link of a primary path goes down and can lead to traffic loss. (497790)
- If a configuration change involving deletion of several VLAN interfaces along with anything which causes platform layer 3 agent to restart is applied, then the layer 3 agent may see multiple restarts following the first restart. (498164)
- After clearing own LSP from IS-IS LSPDB using clear isis database <own LspID>, Isis will fail to export its own LSP information to IGP Topology Database. This can result in CSPF agent to not compute the paths for RSVP and the backup paths for TiLFA. Also setting overload bit after clearing the LSP will cause Isis agent to restart. (498227)
- 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)
- If a MPLS L3 EVPN route resolves over a single NexthopGroup tunnel, then the packet may get sent out without the MPLS L3 EVPN encapsulation label. The workaround is to issue "ip hardware fib next-hop proxy disabled" command. (501739)
- When a RSVP tunnel A which has bandwidth reserved gets preempted by a higher priority tunnel B, sometimes Cspf may not compute an alternate path for A and so the tunnel will stay down. One workaround is to change any TE link's attribute like TE

metric,
admin group or SRLG temporarily and revert it back. An unused TE attribute
is better,
because it won't affect already established tunnels. (501974)

- 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)
- Isis agent may restart on a level1-2 IS-IS node when TI-LFA fast-reroute configuration is present and anycast prefix segment is configured on the nodes in different levels such that one node is SR capable and the other node is non SR capable. (508664)
- In rare conditions, Cspf agent may unexpectedly restart when a path's configuration is changed by Rsvp agent. (515795)
- Ira agent may restart continuously after an agent-restart or switchover when LDP is configured.

As a workaround, device should be reloaded. (520036)

- 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)

Multi-agent

- If a VRF is deleted while the Bgp agent is in the middle of processing paths in BRIB, it may restart unexpectedly. (209143)
- Performing a stateful switchover in a switch with a large number of VRFs configured, with BGP config in all VRFs and IGPs configured in a large number of VRFs, may cause the system to restart unexpectedly, leading to traffic loss. (376773)
- 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)
- BGP agent may unexpectedly restart if 'bgp skip rib-install bestpath-selection' is configured.

Workaround is to disable in-place fec update using 'bgp fec skip in-place update event all' in router bgp config mode (495110)

- EVPN A-D routes used for VPWS that are received or changed via inbound policy to contain color extended communities can cause patch panels to flap leading to brief traffic loss. This can occur when adding and/or removing color extended communities from a route that already contain color extended communities. (498735)
- If multiple dynamic BGP peers are present in a peer-group and the peer-group has a prefix-list in the configured inbound policy, the inbound policy may not get re-evaluated for all the dynamic peers on modification to the prefix-list.
"clear ip bgp *" or BGP agent restart can be used to re-evaluate the input policy correctly when this happens. (504936)
- BGP Fallback AS is not supported for dynamic BGP peers (507024)
- At large route scale, the restart of either the Bgp agent or the IpRib agent may cause repeated restart of the IpRib agent (508763)
- When BGP additional-path install is configured, the usage of ECMP FEC resource increases. (514388)

Rib

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

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

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

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

EOS 4.21.3F or newer releases allow much more granular control of what community types (standard, extended, large communities etc.) are included in the outbound route advertisements. (384629)

- Rib agent may restart with BGP PIC enabled (bgp additional-paths install) under route churn (512675)

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

- Performing SSO, or restarting L3 forwarding agent, on platforms that support MPLS non stop forwarding, might prevent L3 forwarding agent to warmup.
The workaround is to perform a hitful restart of L3 forwarding agent. (499349)

Multicast

- An S,G route may not properly cleanup previously sent S,G joins. This would prevent S,Gs from being deleted in the upstream router. A work around is to flap the interface in which the join is sent on. (291756)
- In an environment with a very large number of MSDP peers, during cleanup the Msdp Agent may crash. (354364)
- If large number of IGMP reports are received in a very short interval, some of them might get dropped resulting in the multicast route not being created. The routes will be re-learned on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIgmp", consider increasing the bandwidth for this Copp class. (356675)
- 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 crash. (418190)
- High CPU utilization might be observed for the forwarding agent when PIM Bidir routes are present. (479119)

- 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
- At switch restart, if PIM BSR configuration exists, PimBsr might crash and could potentially affect existing traffic. (495416)
- BessMgr agent could restart unexpectedly when resolving a large number of unresolved packets at once. (495863)

vEOS Router

- 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 PhyEthtool agent crash. 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)

- When CloudEOS is deployed on the AWS platform in Sfe mode, the Sfe agent crashes. (487773)
- If IPv6 or multicast packets are sent to vEOS on two or more interfaces vEOS forwarding engine can crash. Disable IPv6 and mcast configuration on the interfaces of the neighboring routers connected to vEOS. (497553)
- DPS (dynamic path selection) using Ipsec connections may become corrupted in the event of a Sfe agent restart.

The workaround is to administratively disable and enable the Vxlan1 interface.

Run "shutdown" then "no shutdown" on the Vxlan1 interface. (503017)

- An improperly formatted TCP MSS option can trigger a Sfe dataplane restart. (504565)

DCS-7170 Series

- If a LAG was created with member-ports in shutdown state, outgoing traffic can incorrectly get hashed to those member links leading to traffic loss. (250543)
- On the 7170, when dynamic NAT is configured on the layer 3 interfaces used for ECMP, return traffic does not get NAT applied unless it comes in on the same ECMP member of the outbound NATted traffic. (398104)
Series Affected: DCS-7170 Series
- When a packet is mirrored to a GRE tunnel, the ECN value in the encapsulation IP header of the mirrored packet may be intermittently corrupted. There is no workaround for this. (442457)

DCS-7170 Series

- If a mac learnt on a vxlan interface on an MLAG peer is promoted due to an event such as a MLAG peer-link going down, it may cause the forwarding plane agents to restart. (261152)
- A MAP-T configuration that has a resolved exception offload nexthop, may transition to unresolved, even if the underlying route remains resolved in 2 cases:
 - 1) If the BfnSlice-FixedSystem agent restarts.
 - 2) If the MAP-T configuration is rapidly replaced with the same configuration, or a configuration using the same exception offload IP. (508539)
- When a monitor destination is continually congested due to fan-in/high rates of traffic, and LAG is used in the system, the slice agent may restart when there is LAG membership change or port status change. (516260)

CCS-720XP Series

- If link speed is changed on a 2.5G interface, Strata forwarding plane agent may unexpectedly restart with "Process died with signal 3 \ (SIGQUIT\)" signature. Subsequent restart of the Strata agent will be resolve the issue. (368424)
- 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)
- If a MBA supplicant is put in authentication failure VLAN on devices that support MAC based VLAN assignment, MAC move for the supplicant won't trigger MBA on the new port. (493172)
- 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)
- In case of WoL, with dot1x port in multi-host mode and dynamic vlan assignment, if an authenticated host is deleted with 'clear dot1x host mac ...', the host

could

no longer be waken up from the dynamic vlan. It could still be waken up from native vlan.

Port shutdown/no shutdown would make waking up the host from dynamic vlan work again. (505855)

- Front panel BASE-T ports might fail to link up due to PhyIsland agent restarting continuously.

Workaround is to reload or power cycle the switch. (523161)

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

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 continuously crash. A reboot or linecard reset is required to clear this condition. (486369)

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

- If the SCD agent unexpectedly restarts, FPGAs controlled by the agent could remain in an uninitialized state. The SandFap agent may then unexpectedly restart with a hardware initialization timeout error.

A workaround is to configure "hardware access-list mechanism tcam". Doing so will have the side effect of disabling AlgoMatch. (501810)

SKUs Affected: DCS-7280CR2A-30, DCS-7280CR2A-60, DCS-7280CR2K-30 and DCS-7280CR2K-60

DCS-7300 and DCS-7500 Series

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

- Hotswapping a card may cause a kernel panic and the switch will reboot unexpectedly. Not running traffic on the card before powering down will make this hit less frequently, but it may still hit. (387173)
- After linecard removal and re-insertion, lag member links may remain inactive due to "waiting for hardware to program port to RX (or receive)".

Restarting the StrataLag agent recovers the port channels. (501496)

Series Affected: DCS-7300X Series

- On modular systems with redundant supervisor cards, if the standby supervisor has connected management interfaces, ZeroTouch provisioning may fail to download a configuration from any connected DHCP servers. The workaround is to "shutdown" the management interfaces on the standby supervisor. (507185)

DCS-7020, DCS-7280R, DCS-7280R2, DCS-7280R3, 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)
- VTEP cannot learn a remote host's MAC address from VXLAN encapsulated IPv6 Neighbor Advertisement packets received from that host. This might be an issue in a pure IPv6 deployment where the VTEP does not receive any bridged traffic from the host. In such a scenario, the traffic toward the host will be flooded in the VLAN.

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

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

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

- VXLAN bridging traffic sent out of LAG interfaces may be temporarily dropped after an SSO event. (251398)
Series Affected: DCS-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)
- 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)
- While in Tap Aggregation mode, the switch will drop all traffic with VLAN tags 0 and 4095, even when allowed VLAN is set to all. Furthermore, it is not possible to explicitly configure VLAN tags 0 and 4095 as allowed. (385763)
- If VXLAN is configured, PBR and QOS policies won't apply to traffic ingressing on a MLAG peer-link. (390804)
- 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)

- 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)
- 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)
- If an access-list is used in a PBR policy and also in a QOS policy, editing this access-list may result in a failure to apply the QOS policy. The workaround is to use two separate access-lists for the PBR and QOS policies. (450595)
- The SandFap agent unexpectedly restarts when a PRBS type unsupported by the interface is configured; this will lead to traffic loss on all the interfaces. (450978)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- VXLAN over a MACsec proxy port as underlay may stop forwarding traffic on reload. (451061)
- 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)
- EventMgr Agent may crash after a manual linecard hotswap. A temporary workaround is made to disable the following functionalities:
 - VOQ Delete event management
 - Environment related health metrics like 'cooling', 'power', 'voltage' and 'temperature' under 'monitor device-health' configuration mode (464788)
- Internally duplicated CFM packet may cause double counting and duplicate replies towards the sender. (466833)
- 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)
- Once the maximum number of supported unique PBR policies per chip is exceeded, removing the PBR policy on an interface, converting the interface to a vlan member and applying a RACL on the SVI does not work. (477753)
- SandL3Unicast agent may leak memory when unshared nexthop group entries are modified repeatedly. (478158)
- When BGP Flowspec rules associated to a VRF are not programmed due to overflow in the number of supported VRFs, and the neighbor advertising those rules is brought down, the SandAcl agent may restart unexpectedly. (479587)
- Restarting the SandAcl agent while BGP Flowspec is applied on more than the supported number of VRFs can lead to unexpected traffic behavior. (481575)
- When a subinterface connector is used in a local patch and the subinterface dot1q encapsulation has both outer/inner tags, packets forwarded to the Tx port corresponding to subinterface will have incorrect tags. An extra tag with VID=0 is inserted. (485071)
- TCAM bank reservation may not work after reload.
The workaround is to remove and reapply the reservation config for any existing feature after all the SandFap agents are up. (488675)
- With a tcam profile configuration containing "feature tunnel vxlan routing", SandAcl agent restart may lead to another unexpected restart before the agent recovers. (492320)
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Once storm control is configured on an interface, policers configured via a QOS policy-map or cpu-traffic policy on this interface will police packets more aggressively than the configured rate.
The workaround is to restart the SandFap agent.
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R, DCS-7500R2 (494749)
- If the SandFap agent is busy, or starved of CPU bandwidth, for an extended period of time, PTP stability may be affected negatively. The message 'TIMESYNC_ERROR' will be logged in syslog when this happens. (495418)
- Following an SSO event on an MLAG peer device with VXLAN configured, VXLAN split horizon rules might not be properly enforced. This can potentially create loops in the VXLAN network.

One possible workaround is to restart the SandTunnel agent on the impacted MLAG peer device. (497551)

- Using a user defined TCAM profile that has the packet type `mpls ipv4 forwarding routed decap` may cause unpredictable behaviour in access control lists (ACLs) and IP TOS field corruption for ingress traffic that is not MPLS. (499458)
- "hardware next-hop fast-failover" is not supported with MPLS EVPN flood adjacency resolved over an MPLS SR tunnel that is in turn resolved over TI-LFA tunnel. (503263)
- In a 6PE configuration, packet ingressing on an L3 subinterface with explicit null label and destined to this switch is not supported. (504582)
- If the source IP is not in the routing table, the unicast RPF will not drop the packet and will forward a corrupted packet. (505285)
- On a modular system, if only some of the linecards are configured in Tap Aggregation mode, inserting or removing a linecard which is configured for Tap Aggregation would cause traffic disruption on other linecards in Tap Aggregation mode. (508024)
Series Affected: DCS-7500R, DCS-7500R2 and DCS-7500R3 Series
- When a DirectFlow flow is modified, the original flow entry might still be programmed in the TCAM along with the modified one. The original TCAM entry will not be visible using CLI. The workaround is to manually restart the SandAcl agent. (508034)
- Source MAC address is not learnt for packets that come in on interfaces with IPv6 uRPF configuration. (512737)
- With hardware accelerated SFlow enabled and more than 20 LAGs in the system and mirroring sessions configured, the Sand agent restarts continuously. (512973)
- Traffic for some routes may get mis-forwarded using default route under rare circumstances, that are influenced by some route distribution, and order of route updates. (513287)
- On a dut configured with subinterface with shape rate configuration, link flap may cause Sand agent to crash indefinitely until reload.
System reload should resolve the issue (519535)

DCS-7280R and DCS-7500R Series

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

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

SKUs Affected: 7500R-8CFPX-LC

- VRF-mapped IP over MPLS packets with TTL=1 in the IP headers may be dropped, if the installed system TCAM profile uses packet types 'mpls ipv4 forwarding routed decap' or 'mpls ipv6 forwarding routed decap'. This could be a predefined TCAM profile such as pbr-match-nextthop-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)
- In an EVPN MPLS configuration, extra invalid packets may be sent out during flooding of a BUM packet in an EVPN MPLS VLAN which may cause loops in the topology.

The workaround is to restart SandL3Unicast agent. (498762)

- EVPN MPLS with implicit null tunnel underlay is supported only when MPLS tunnel counters are enabled. (500298)

DCS-7280R2 and DCS-7500R2 Series

- If there is a parity error in the PPDB_B block of a JerichoPlus or QumranAx chip, the router MACs pointing to the CPU may get deleted from the hardware MAC table on that chip leading to loss of L3 connectivity.

As a workaround, restart the SandMact agent using the CLI "agent SandMact terminate". (521075)

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)
- The SandFapNi agent may restart continuously preventing ports coming up. (462680)
 SKUs Affected: DCS-7280CR3-96, DCS-7280CR3K-96, DCS-7280DR3-24, DCS-7280DR3K-24, DCS-7280PR3-24 and DCS-7280PR3K-24
- 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)
- PBR, QOS, and CPU policies applied to a VLAN interface will not function correctly on untagged packets ingressing trunk ports mapping to a native VLAN equaling that of the VLAN interface. (483364)
- Link state change of an interface can cause an adjacent interface to flap. (486473)
- 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)
- With sustained high MAC scale, hardware learning of L2 MACs may get disabled leading to flooding.

When this condition occurs, a workaround is to clear the MAC table (clear mac address-table dynamic). (488731)

- PBR, QOS, CPU policies applied to VLAN interfaces may not function correctly if VLAN interface memberships change during SandFapNi restart, Linecard insertions or removals, or during interface speed changes.

The workaround is to restart the SandTunnel agent, or reapply the policy configurations. (488958)

- If tap aggregation and a CPU traffic policy are configured together, either the CPU traffic policy will not apply to any traffic or the SandAcl agent will crash continuously. (492549)
- With large Nexthop group updates, in rare instances Packet can egress with wrong header data. (499845)
- On QSFP ports with the CRT50216 PHY, configuring 50G-2 on the /1 and /3 may result in the /1 flapping 1 or more time. (500018)
- 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-32D4-F, DCS-7280CR3MK-32D4-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)
- If same IP/IPv6/MAC ACL is applied on a front panel port or Port-channel and a subinterface of that port or port-channel, then the behaviour will depend on the order in which the ACL is applied on these interfaces. The ACL may not actually be applied if ACL is applied first on at least one of the subinterface and then on the parent interface. Also, it is not possible to remove ACL configuration from the parent interface using 'no <ip/ipv6/mac> access-list <aclname> <direction>'. Configuration can be removed using 'configure replace'. Workaround is to use different ACLs (may contain same rules) on parent and its subinterfaces. (503650)
- Replacing a linecard with a different type of linecard while the switch is live will cause the SandFapNi agent to crash. Rebooting the switch will recover the problem. (503896)
- When a Gratuitous ARP (GARP) that maps an IP address to a special multicast MAC is received with another packet, the system can restart unexpectedly due to a kernel crash. (504140)
- VXLAN transit packets received on the MLAG peer-link can be misforwarded or dropped. (512214)
- If an IP/IPV6 or MAC ACL is removed from an interface, physical or logical, with the exception of vlan interfaces, after a SandAcl restart, then traffic on that interface may continue to hit the ACL if the ACL is still configured

on any other interface or it may hit some other ACL even though that ACL is not configured on this interface. (516000)

- Packets that exceed the configured L3 MTU are still subject to TX mirroring. (516199)
- 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)

DCS-7020 Series

- When multiple Ipsec tunnel interfaces are shutdown, then "no shut" in a short period of time, routes associated with some tunnel interfaces may not be installed after "no shut". The workaround is to do shut/no-shut again on the individual tunnels where routes not installed. (377984)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- When traffic is flowing from a 10G interface to a 1G interface and PFC is enabled on the ingress 10G interface, if the administrator changes the COS to TC mapping all the incoming traffic on the 10G interface is dropped. The only way to recover is by restarting SandFap. (400653)

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

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

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

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

- 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)

- 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)
- Hitless Strata agent restart may cause a Port Channel interface to flap if it contains members with auto-negotiation speed configuration (419369)
- If the number of next-hop groups configured exceeds 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
- 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)
- Slice agent may restart unexpectedly on performing speed change and/or shutting/no shutting on a PFC enabled interface with traffic flowing out of the interface.
The issue will not be seen if PFC is not configured or traffic is not flowing. (462139)
Series Affected: 7368 Series
SKUs Affected: DCS-7060DX4-32-F and DCS-7060PX4-32-F

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

- Changing the speed of a subordinate port will cause the rest of the ports on a QSFP port to flap. (472767)

- 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)

- LANZ fabric thresholds may not be applied if the Strata slice forwarding agent restarts. Workaround is to remove and reapply the thresholds. (480422)
SKUs Affected: DCS-7250QX-64-F, DCS-7250QX-64-R, DCS-7250QX-64-SSD-F, DCS-7250QX-64-SSD-R, DCS-7260CX-64-F, DCS-7260CX-64-R, DCS-7260CX-64-SSD-F, DCS-7260CX-64-SSD-R, DCS-7304, DCS-7308 and DCS-7316

- L3 agent restart or "reload fastboot" may cause routed packets to be misforwarded for up to 1 second. There is no workaround for this issue. (481604)
- Aggregate storm control rate limit is not honored for broadcast and multicast traffic. (490511)

- 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)
- In some cases, post forwarding agent restart, egress ACLs may not work properly and can affect traffic egressing out of ports where this egress ACL is not applied (501088)
- 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)
- Configuring 25G on the first lane of a QSFP port and then inserting a 40G transceiver may lead to a forwarding agent restart. Workaround is to remove the configuration before inserting the transceiver. (506473)
Series Affected: DCS-7050CX3, DCS-7050SX3 and DCS-7260CX3 Series
- L3 routes with ECMP to multiple remote VTEPs may take longer time to converge when members in ECMP goes down. (513994)
- If a multi-ASIC linecard is hotswapped with a single-ASIC linecard in the same slot, the Snmp agent may continuously restart. The only workaround is to reload the system. (515667)
Series Affected: 7300X3 and DCS-7300X Series

DCS-7010 Series

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

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

DCS-7260X3 Series

- Ports with 1g transceivers will still count towards the limit of 5 speed classes even when disabled and may lead to continuous restart of the Strata

forwarding agent. It can cause traffic outage on other ports.

Workaround is to remove the lg transceiver. (338741)

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

7368 and DCS-7060X4 Series

- 802.3X pause is not supported (348758)
- 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)

- Link may be unstable at 400G with 400GBASE-DR4, 400GBASE-AR8 or 400GBASE-SR8 transceivers. (472927)
SKUs Affected: 7368-4D and 7368-4P
- Control packets destined to link local IPv6 addresses may go to glean queue and hence could get dropped resulting in protocol flaps. There is no workaround to this issue. (475260)
- 100G-2 is not supported. (505274)
SKUs Affected: 7368-16C
- OSPFV3 Control packets destined to the CPU may go to glean queue and hence could get dropped resulting in protocol flaps. There is no workaround to this issue. (520443)

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

- Configured as MLAG, link flap of one or more LAG peer-link member ports could cause brief flooding of packets over the peer-link and can also lead to double delivery of packets on the MLAG interfaces. (117870)
- (A1 silicon only) MPLS traffic gets incorrectly prioritized, causing other traffic on the same port to get dropped. (160269)
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-FixedSystem agent to restart following which the interface continues to transmit properly. (245501)

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)
- Externally calibrated SFP form factor transceivers will operate with system default high alarm threshold of 63C and system default high warning threshold of 58C and may be logged as overheating when operating within their thresholds as programmed in the module EEPROM. (489787)

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 an 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 an 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

- Performing SSU may result in extended traffic outage, leading to total traffic loss being slightly over 1 second. (477139)
- Some NAT rules may not work correctly if the StrataL2 forwarding agent happens to restart. Packets may not get translated as expected
The impacted rules need to be un-configured and then re-configured. (483821)
Series Affected: DCS-7300X Series
- Aggregate storm control rate limit is not honored for traffic that is routed and flooded in the egress VLAN. (485492)
- If EVPN VXLAN multihoming is configured, and link towards ethernet-segment goes down, then some macs that reside behind that link may get stuck in hardware. As a result, traffic destined to those macs may get dropped. The workaround is to reboot the switch. (495546)
- Having any ACLs with UDF qualifiers as well as any of the following LAG/ECMP hashing features present in the configuration: RRoce hashing or TTL/HopLimit hashing, will result in incorrect operation of the LAG/ECMP hashing where the switch incorrectly hashes on IP payload. To avoid this problem the conflicting features should not be configured at the same time. (497016)
Series Affected: DCS-7050CX3, DCS-7050SX3 and DCS-7300X Series
- If NAT configuration is modified or if a NAT-enabled interface is shutdown and enabled while traffic is flowing, then some of the NAT connections might not be setup properly leading to packet loss or untranslated traffic.

Workaround is to be quiesce traffic before modifying the NAT configuration and then resume traffic. (497045)

- With VXLAN enabled, during SSO, control plane traffic may cause a kernel panic and unexpected reload on the new primary supervisor. (498652)

DCS-7160 Series

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

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

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

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

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

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

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

A workaround is to reboot the switch. (255957)

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

Series Affected: DCS-7160 Series

- XPL3Unicast process may crash when the switch is reloaded if PBR policies with nexthop VRF actions are applied on an interface in the non-default VRF. (262713)
- 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)
- A sub-50 ms power outage may cause the links to go down permanently with the XpCentral agent crashing repeatedly. Reloading the system is required to recover from this failure. (457940)

SKUs Affected: DCS-7160-48TC6-F and DCS-7160-48TC6-R

Limitations and Restrictions in 4.24.2.3F

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)
- Mss does not support traffic interception for the hosts connected to Service VTEP. (354754)
- The MacroSegmentation Service (MSS) allows users to configure multiple tags on CVX that can be used in firewall policy definition, to identify policies MSS should work on. It is however not possible to delete one of the multiple tags configured.

A workaround would be to remove all tags and reconfigure the subset required. (372320)

- In Macro-Segmentation Service (MSS) for L3 Firewalls , Virtual SVI IP addresses cannot be treated as intercepted hosts and thereby cannot be added to any redirect policy. However, these IPs can still be part of offload policies, even though implicit redirect will not work for them. (372504)
- MLAG devices in an environment with the Macro-Segmentation service (MSS) configured (with L3 firewalls), should have a high-priority DirectFlow rule configured to ignore traffic over the MLAG peer link. This can be configured using the following flow definition on each switch (assuming Po1 is the MLAG peer link):
flow bypassPeerLink
 priority 65535
 table trident ifp
 match input interface Po1
 match ethertype ip (375185)
- The Macro-Segmentation service (MSS), running with L3 firewalls, requires the following flows to be configured on the service VTEP devices connected to the firewall manually in order to prevent return traffic from the firewall from being subject to MSS rules again. The flows required are:

 flow bypassFwIntf3

```

    priority 65535
    table trident ifp
    match input interface Pol103    >>>>>>> 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)
- 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)

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

- Enabling next-hop sharing (using "ip hardware fib next-hop sharing" command) could result in slower convergence in the case of link failures. (369305)
- 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)
- In the 'show interfaces interactions' CLI, the list of speeds an interface is limited to does not account for the limitations that result from speed-group configuration. (434363)
- Configuring both TACACS+ and LDAP with SSL profiles may cause the Aaa agent to continuously restart (439333)
- When we first configure a port-channel lag_type, we must set values for both aggregation/config/lag-type and ethernet/config/aggregate-id (i.e. member configuration) in the same transaction. The lag-type is populated in EOS once the port-channel has one member port and remains set. (445421)
- 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)
- 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)

```

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-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32P4-F and DCS-7280CR3MK-32P4-R

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)

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)
- BMP is not supported for router bgp instances configured in non-default VRFs. (210363)
- SSO is not supported when BFD is configured for ISIS IPv6 sessions. (239704)
- An LACP port channel configured with BFD per-link and echo may experience BFD session flap on existing member links if new member link(s) are added to the port channel while the BFD peer device undergoes Stateful Switchover (SSO). (255042)
- Eos does not support programming IPv6 ECMP paths to kernel (258387)
- The DSCP value of BGP control plane packets set by the command 'bgp transport qos dscp <value>' will be overwritten when the platform 'qos rewrite dscp' is enabled and the outgoing interface is a SVI based on traffic-class to dscp map present on the platform. (269764)
- In the output of the "show ipv6 dhcp relay counters" or "show ip dhcp relay counters" command, interface-specific DHCP relay counters might not sum up to "All Req" and "All Resp" counters. (278786)
- Tunnel interfaces flap when configuring ttl, tos and path mtu discovery on them. (281464)
- IPv6 addresses are not supported as tunnel source and destination. (283912)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289217)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289224)
- The command "[no] service routing configuration route-map no-set-community explicit-members"

affects the behaviour of the route-map command

"no set community [community-list] [member]+ [action]".

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

- When BFD is configured with per-link mode rfc-7130 for an IPv6 link-local address on a Port-Channel and the Port-Channel transitions to the Down state, the Port-channel may not transition back to the UP state.

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

- On Macro-Segmentation Service (MSS) deployment with L3 firewalls, a flow entry that fail to get programmed due to lack of TCAM space, is not programmed when TCAM space becomes available at a later time. The workaround is to "shut" and "no shut" on the direct flow config. (372232)
Series Affected: DCS-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)
- 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)

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)
- In multi-agent mode, extended Sflow BGP data is not supported for BGP LU routes. (411762)
- 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)
- An AS_PATH prepend operation in an RCF function applied towards an IBGP peer is ignored. (474006)
- RCF compiler does not accept all characters allowed in prefix list names, community list names, and AS path list names. (483676)

- When a router receives a BGP route with an AIGP metric attribute value approaching 64-bit unsigned integer maximum value, another BGP route with a next hop covered by the first BGP route may not have it's MED attribute correctly advertised should it match an outbound route-map policy with 'set metric igp next-hop-cost' configured (487359)

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)
- 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 crash on config replace with scaled setup over 30 VRFs and large number of PIM enabled interfaces. Work around is to reduce the scale. (398185)
- Pim "non-dr install-oifs" feature becomes disabled if the non-dr needs to route multicast traffic for any reason to the interface on which the feature was originally configured. This is how the feature is expected to work. The workaround is to setup the network in a way that the DR always becomes the assert-winner. However, if the non-DR becomes assert winner for some routes on an interface and feature becomes disabled, flap the pim sparse mode configuration on non-DR on that interface to attempt to get out assert winner state. (401672)

vEOS Router

- 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)

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)
- Accelerated Software Upgrade (ASU) is not supported in combination with the packet-filter profile. The hitless reload will time out and fall back to a normal reload. (473985)
Series Affected: DCS-7170 Series
- On the DCS-7170 series, ASU is not supported on the firewall profile. (475112)
Series Affected: DCS-7170 Series
- Hitless reload may incur several seconds of traffic outage if route scale is large (490931)

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

- 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)

- 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)
- WOL feature is currently only available for phone trunk port, but not for regular trunk port. (481468)
- Timeout value configured for MBA fallback feature can block MBA Auth for upto 60 seconds more than the configured value since first non-EAPOL packet is received. (483995)
- 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)
- When authentication caching is enabled, a link down event will continue to cache any attached hosts for up to two minutes. If a MBA host is removed from the port and the link is brought up by plugging in a new host before the two minutes have expired, the original host's credentials will remain cached even though it is no longer attached to the switch. Accounting information will show that the host is still connected.

One work around is to wait more than two minutes before plugging in a new host, however the accounting information will show that the host was connected during this time. If precise accounting is required, use the 'clear dot1x host mac' command to manually remove the MAC address as soon as it has been unplugged. (504881)

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)

DCS-7300 and DCS-7500 Series

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

Use LACP or disable the port-channel members on the inserting linecard before the linecard insertion (361336)

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

- When 'poe shutdown action power-off' CLI config is in use, PoE power will not be turned off when an interface is errdisabled (506677)

DCS-7020, DCS-7280R, DCS-7280R2, DCS-7280R3, 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-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, 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)
- show interfaces counters rates' for port channels may show a slightly inaccurate data rate (including above 100%). (67367)
- 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 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)
- After applying or removing shaping configuration on an L2 subinterface, the L2 subinterface must be flapped for the configuration to take effect. (278730)
- With TCP MSS Clamping feature configured, TCP SYN packets might get dropped under high CPU usage conditions if the incoming TCP SYN packet rate is high. (279370)
- Packets larger than 10196 bytes will be discarded by the forwarding ASICs. (280459)
- PFC on DSCP trusted ports is not supported. (281214)
- Loop protect feature is not supported when AlgoMatch HW is enabled. (290706)
- 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)
- Egress MAC ACLs are matched on the incoming packet's ethernet header and not on the rewritten ethernet headers. (458724)
- IPv4/IPv6 counters are not supported on port-channel interfaces. (479658)
- Egress mac acls do not work on interfaces that are not front-panel interfaces, even though they can be configured. (483583)

- When the feature "flow" in a TCAM profile is configured with only the "drop" action, the profile is rejected. The workaround is to add any other action to feature "flow". (486849)
- 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-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series

- Lag subinterfaces are not currently supported in pseudowire or local-local patches (i.e. a patch with two local interfaces). (490834)
- 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
- While a switch is in Tap Aggregation mode, if the total number of FAPs on all linecards in Tap Aggregation mode is more than 73, then some Tap Agg features, e.g. MPLS Pop, may stop working on some FAPs, due to insufficient hardware resources. When this happens, the message "Hardware resources are insufficient to program all entries in ISEM-A table" will be logged. (498192)
- For Packets with GRE tunnel encap header, MTU enforced on the hardware is 3 bytes more than configured value in Cli. (498470)
Series Affected: DCS-7280CR2, DCS-7280R, DCS-7280R2, 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)

DCS-7280R and DCS-7500R Series

- An IPv6MPLS packet terminated by L3 EVPN MPLS configuration will not be forwarded. The workaround is to disable IPv6 exact match host routes using 'no platform sand ipv6 host-route exact-match' and disable IPv6 route optimizations for prefix length 128. (201414)
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-7280R, DCS-7280R2, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series

- Egress Ipv4 ACLs and Egress MAC ACLs cannot be configured together (443671)
- Egress Ipv6 ACLs and Egress MAC ACLs cannot be configured together. (443673)
- With `counters per-entry` for egress mac acl, we can uniquely count max of 3967 rules per fap. (460819)
- MTU is not enforced for MPLS packets forwarded by a MPLS swap route, which have a length one byte more than MTU. (464742)
- The two rate three color (trTCM) QoS policy-map counters can show incorrect values for packets marked green or yellow if the ingress traffic is broadcast or multicast traffic. (479267)

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

- L2VPN or VPWS resolution through SR-TE tunnels consisting of 5 or 6 label stack isn't supported. (494953)

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)
- With "match mpls" PBR configuration, MPLS packets matched both by PBR and a LFIB route with a pop action would get forwarded correctly only if the MPLS route contains "access-list bypass" configuration. (370118)
- 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)
- Traffic drop may happen when we sent traffic to all 8 traffic-classes (aggregate rate nearing linerate) to port which has been configured with Strict-Priority scheduling. (406550)
- Traffic drop may happen when we sent traffic to all 8 traffic-classes (aggregate rate nearing linerate) to port which has been configured with WRR scheduling. (406561)
- Only port-channel interfaces are supported as peer links. Single member port-channel interfaces are supported. (409958)

- Per TC based VOQ tail-drop thresholds isn't supported in software. (413344)
- Set dcsp will not apply to mpls routed packets that egress with mpls labels (418791)
- Unidirectional links are not supported at 40G speed on QSFP100 interfaces using the CRT50216 PHY. (467672)
- Certain speed groups are limited to one active serdes rate configuration. These speed groups are those that show only one serdes rate in the Active column of the output of "show hardware speed-group status". (486659)
- For VXLAN decapsulation + routing of overlay packet, the ECN bits may not be correctly copied from the incoming overlay packet. (489998)
- VXLAN VARP VTEP with IPv6 overlay routing requires vxlan-routing TCAM profile to be applied on the switch. (503496)
- The "phy link serdes mapping direct" CLI command requires the 25g serdes rate to be configured, unlike what is displayed in "show hardware speed-group configuration". (523850)
 SKUs Affected: 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3K-48CQ-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32P4, DCS-7280CR3K-96, DCS-7280CR3MK-32D4 and DCS-7280CR3MK-32P4

DCS-7020 Series

- The interface bin counter for both in and out packets with sizes 1523-Max may be inaccurate. Packets with sizes between 9217-9236 are not counted by the switch for interfaces Ethernet 1-48. (188121)
- Mirroring and sFlow can not be supported on the same source interface at the same time due to a hardware limitation. (209060)
- AES128/SHA-1 and AES256/SHA-256 are the only supported encryption/authentication algorithm pairs. (342005)
 SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- The current implementation of IPSec supports only encryption/authentication algorithm combinations of AES128/SHA1 or AES256/SHA256. No other combinations are currently supported. (359986)
 SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Using an Ipvsec tunnel interface for policy-based routing nexthop is not supported. (378861)
 SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Ipvsec 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-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)
- 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)

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

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

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

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

SKUs Affected: DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050TX3-48C8-F and DCS-7050TX3-48C8-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 'hardware nat static access-list resource sharing' is in the switch's config,
and a static NAT filter ACL and a security ACL have overlapping rules, the security
ACL's hit counter might not increment in all cases. Regardless of the hit
counter,
both the security ACL rule and the NAT filter ACL rule will be applied
correctly.

Workaround is to disable static NAT ACL resource sharing with:

'no hardware nat static access-list resource sharing'

Alternative workaround is to disable port ACL in VFP with:

'platform trident tcam port-acl vfp disabled' (450080)

- 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 the incoming packet is VLAN tagged, then the truncated mirrored packet
will be 4 bytes shorter than expected. (463986)
Series Affected: 7300X3, 7368, CCS-720XP, DCS-7050X3 and DCS-7060X4 Series
- 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)
- By default, configuring a static NAT rule with an empty access-list is
treated the same as no access-list: for source NAT, the switch will ignore
the destination IP, and vice versa for destination NAT.
With "hardware nat static access-list resource sharing" enabled, configuring
a static NAT rule with an empty access-list is treated the same as an
undefined access-list: some hardware resources may be consumed, but the
traffic will be forwarded untranslated. (473783)
Series Affected: CCS-720XP, DCS-7050CX3 and DCS-7050X3 Series

- 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)
- When persistent port security is enabled on an interface, MAC addresses will persist when an interface is down. While the interface remains down, traffic destined to these MAC addresses will be dropped.

Persistent port security is enabled by default, and can be disabled with `switchport port-security persistence disabled`. (485828)

- If dynamic load balancing is enabled for an ECMP group, it's member link flap might cause poor traffic distribution. Traffic distribution may improve over time. (487595)
Series Affected: DCS-7060X4 and DCS-7260X3 Series
- If dynamic load balancing is enabled for ECMP groups while the traffic is flowing, then the load on the ECMP members may not be distributed uniformly. Traffic distribution may improve over time. (490418)
Series Affected: DCS-7060X4 and DCS-7260X3 Series
- When dynamic load balancing is enabled, If multiple traffic flows destined to the same ECMP group arrive at the same time, then the load on the ECMP members may not be distributed uniformly. Traffic distribution may improve over time. (490716)
Series Affected: DCS-7060X4 and DCS-7260X3 Series
- Packets sent from the CPU cannot be egress mirrored unless they are being sent out of an SVI. (495146)
- If the speed of an interface changes with traffic running, it is possible for "show interface counters" and "show interface counters fast-poll" to diverge on that interface. (502252)
- Egress mirroring of multicast packets that are going to be encapsulated on the switch independent of mirroring is not supported. (512880)

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

- LANS 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)
- 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)
- When 'hardware counter feature gre tunnel interface in' feature is enabled, the gre tunnel interface decap counters are not incremented when subinterface is used as the ingress interface. (464972)

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)
- L2 protocol forwarding is not supported (470496)

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)

- 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)

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 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
- Dynamic NAT requires the forwarding table partition to be configured in flexible mode and entries allocated for exact-match table.
Eg: platform trident forwarding-table partition flexible exact-match 49152 12-shared 32768 13-shared 131072
If not configured, no logs or notification will be generated, and NAT might not work. (494088)
- NAT is not supported with non-default VRF. Configuring NAT over non-default VRF might cause Nat agent to restart continuously.

Workaround is to remove the unsupported configuration. (499723)

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)

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.