

Arista Networks EOS 4.29.9.2M Release Notes

Version: 0.1

21 May 2025

Table of Contents

EOS 4.29.9.2M Release Highlights

New Platforms and Hardware

SNMP MIBs

SNMP Traps

Supported Hardware

Reference to MLAG ISSU Upgrade/Downgrade Table

Resolved Caveats in 4.29.9.2M

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7260X, DCS-7260X3 and DCS-7300X Series

DCS-7060X5 Series

Known Software Caveats in 4.29.9.2M

AristaSAI

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

vEOS Router / CloudEOS

CVP

Network Provisioning

CVX

DANZ Monitoring Fabric

Enterprise routing

General

IPSEC

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

MPLS

Multicast

NAT

SwitchApp

vEOS Router / CloudEOS

DCS-7170 Series

DCS-7170 Series

CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM and CCS-750 Series

CCS-750 and DCS-7060X4 Series

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

DCS-7130 Series

7358, 7368, 7388, CCS-750, DCS-7300X, DCS-7500R, DCS-7800 and DCS-7800R3 Series

7500R3, DCS-7020, DCS-7130, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280R3A, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800, DCS-7800R3 and DCS-7800R3A Series

DCS-7280R and DCS-7500R Series

DCS-7280R2 and DCS-7500R2 Series

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

7500R3, DCS-7130, DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series

DCS-7280R3A and DCS-7800R3A Series

DCS-7020 Series

DCS-7280R3 Series

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

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

7500R3, DCS-7130, DCS-7280R3, DCS-7280R3A, DCS-7500R3, DCS-7800R3 and DCS-7800R3A Series

DCS-7280R and DCS-7280R2 Series

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7260X, DCS-7260X3 and DCS-7300X Series

CCS-750 Series

7300X3, 7320X, DCS-7060X and DCS-7260X Series

DCS-7060X2 Series

DCS-7260X3 Series

7368 and DCS-7060X4 Series

DCS-7060X5 Series

7388 and DCS-7060X5 Series

CCS-710P, CCS-720DP and CCS-720DT Series

CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series

CCS-720DF and CCS-720XP Series

CCS-750 Series

DCS-7050X3 Series

7300X3, AS7326, AS7726 and DCS-7050X3 Series

7358 and DCS-7050X4 Series

7388 and DCS-7060X5 Series

Transceiver Series

7300X3, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX and DCS-7050X3 Series

7358 and DCS-7050X4 Series

Limitations and Restrictions in 4.29.9.2M

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

Containerized EOS

vEOS Router / CloudEOS

CVX

Enterprise routing

EOS SDK

General

IPSEC

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MetaMux

MetaWatch

MLAG

MPLS

Multi-domain Segmentation Services

Multicast

NAT

SwitchApp

DCS-7170 Series

DCS-7170 Series

CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM and CCS-750 Series

CCS-750 and DCS-7060X4 Series

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

DCS-7130 Series

7358, 7368, 7388, CCS-750, DCS-7300X, DCS-7500R, DCS-7800 and DCS-7800R3 Series

CCS-710P, CCS-720DP, CCS-720XP, CCS-722XPM and CCS-750 Series

7500R3, DCS-7020, DCS-7130, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280R3A, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800, DCS-7800R3 and DCS-7800R3A Series

DCS-7280R and DCS-7500R Series

DCS-7280R2 and DCS-7500R2 Series

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

7500R3, DCS-7130, DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series

DCS-7280R3A and DCS-7800R3A Series

DCS-7020 Series

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

7500R3, DCS-7130, DCS-7280R3, DCS-7280R3A, DCS-7500R3, DCS-7800R3 and DCS-7800R3A Series

DCS-7280R and DCS-7280R2 Series

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7260X, DCS-7260X3 and DCS-7300X Series

CCS-750 Series

7300X3, 7320X, DCS-7060X and DCS-7260X Series

DCS-7260X3 Series

7368 and DCS-7060X4 Series

CCS-710P, CCS-720DP and CCS-720DT Series

CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series

CCS-720DF and CCS-720XP Series

CCS-750 Series

7300X3, AS7326, AS7726 and DCS-7050X3 Series

7358 and DCS-7050X4 Series

7388 and DCS-7060X5 Series

Transceiver Series

7300X3, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM,
CCS-750, DCS-7010TX and DCS-7050X3 Series

7358 and DCS-7050X4 Series

Conditions and Impacts

EOS 4.29.9.2M Release Highlights

New Platforms and Hardware

- None

SNMP MIBs

- ARISTA-ACL-MIB
- ARISTA-ASIC-COUNTERS-MIB
- ARISTA-BGP4V2-MIB
- ARISTA-BRIDGE-EXT-MIB
- ARISTA-CONFIG-COPY-MIB
- ARISTA-CONFIG-MAN-MIB
- ARISTA-DAEMON-MIB
- ARISTA-ENTITY-SENSOR-MIB
- ARISTA-FIB-STATS-MIB
- ARISTA-GENERAL-MIB
- ARISTA-HARDWARE-UTILIZATION-MIB
- ARISTA-IF-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-MAU-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-NEXTHOP-GROUP-MIB
- ARISTA-PFC-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-PRODUCTS-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-QOS-MIB
- ARISTA-QUEUE-MIB
 - Excluding 7150
- ARISTA-REDUNDANCY-MIB
 - 7500R, 7500E, 7300X, 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-SMI-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-SNMP-TRANSPORTS-MIB
 - 7050X2, 7060X, 7260X, 7320X, 7060X2, 7160, 7020R
- ARISTA-SW-IP-FORWARD-MIB
 - IPv4 only
- ARISTA-SW-IP-FORWARDING-MIB
- ARISTA-TEST-MIB
- ARISTA-VRF-MIB
- ARISTA-VXLAN-MIB
- ARISTA-XCVR-DWDM-MIB
- ARISTA-XGS-MIB

- ENTITY-MIB
- ENTITY-SENSOR-MIB
- ENTITY-STATE-MIB
- HC-RMON-MIB
 - etherStatsHighCapacityGroup
- HOST-RESOURCES-MIB
- IEEE8021-PFC-MIB
- IEEE8021-SECY-MIB
- IEEE8023-LAG-MIB
- IGMP-MIB
- IPMCAST-MIB
- IPMROUTE-STD-MIB
- LLDP-EXT-DOT1-MIB
- LLDP-EXT-DOT3-MIB
- LLDP-MIB
- MPLS-LDP-STD-MIB
- MPLS-TE-STD-MIB
- MSDP-MIB
- NOTIFICATION-LOG-MIB
- OSPF-TRAP-MIB
- OSPFV3-MIB
- PIM-MIB
- PTPBASE-MIB
- RFC 2863 IF-MIB
 - obsoletes RFCs 1229, 1573, 2233
- RFC 2864 IF-INVERTED-STACK-MIB
- RFC 3418 SNMPv2-MIB
 - obsoletes RFCs 1450, 1907
- RFC 3635 EtherLike-MIB
 - obsoletes RFCs 1650, 2358, 2665
- RFC 3636 MAU-MIB
 - ifMauDefaultType and ifMauAutoNegStatus are writeable
- RFC 4022 TCP-MIB
 - obsoletes RFC 1213
- RFC 4113 UDP-MIB
 - obsoletes RFC 1213
- RFC 4188 BRIDGE-MIB
- RFC 4273 BGP4-MIB
- RFC 4292 IP-FORWARD-MIB
 - obsoletes RFC 1354
- RFC 4293 IP-MIB
 - obsoletes RFC 1213
- RFC 4363 Q-BRIDGE-MIB
- RFC 4750 OSPF-MIB
- RMON-MIB
 - rmonEtherStatsGroup

- RMON2-MIB
 - rmon1EthernetEnhancementGroup
- SNMP-TLS-TM-MIB
 - RFC 6353
- SNMP-TSM-MIB
 - RFC 5591
- SNMPv2, SNMPv3
- VRRPV2-MIB

SNMP Traps

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

- VRRPv2-MIB
 - vrrpTrapNewMaster

Supported Hardware

Fixed System

* 7326-56X-O-AC	* DCS-7130-48LS-F (HwRev:B2)	* DCS-7280SR-48C6
* 7326-56X-O-AC-F	* DCS-7130-48LS-F (HwRev:B3)	* DCS-7280SR-48C6-F
* 7726-32X-O-AC	* DCS-7130-48LS-R (HwRev:B2)	* DCS-7280SR-48C6-M-F
* 7726-32X-O-AC-F	* DCS-7130-48LS-R (HwRev:B3)	* DCS-7280SR-48C6-M-R
* CCS-710P-12	* DCS-7130-96	* DCS-7280SR-48C6-R
* CCS-710P-16P	* DCS-7130-96-F (HwRev:A7)	* DCS-7280SR2-48YC6
* CCS-720DF-48Y	* DCS-7130-96-F (HwRev:A8)	* DCS-7280SR2-48YC6-F
* CCS-720DF-48Y-2	* DCS-7130-96-R (HwRev:A7)	* DCS-7280SR2-48YC6-M-F
* CCS-720DF-48Y-2F	* DCS-7130-96-R (HwRev:A8)	* DCS-7280SR2-48YC6-M-R
* CCS-720DF-48Y-F	* DCS-7130-96E	* DCS-7280SR2-48YC6-R
* CCS-720DP-24S	* DCS-7130-96E-F (HwRev:A7)	* DCS-7280SR2A-48YC6
* CCS-720DP-24S-2	* DCS-7130-96E-R (HwRev:A7)	* DCS-7280SR2A-48YC6-F
* CCS-720DP-24S-2F	* DCS-7130-96L	* DCS-7280SR2A-48YC6-M-F
* CCS-720DP-24S-F	* DCS-7130-96L-F (HwRev:A2)	* DCS-7280SR2A-48YC6-M-R
* CCS-720DP-48S	* DCS-7130-96L-F (HwRev:A3)	* DCS-7280SR2A-48YC6-R
* CCS-720DP-48S-2	* DCS-7130-96L-R (HwRev:A2)	* DCS-7280SR2K-48C6
* CCS-720DP-48S-2F	* DCS-7130-96L-R (HwRev:A3)	* DCS-7280SR2K-48C6-M-F
* CCS-720DP-48S-F	* DCS-7130-96LA	* DCS-7280SR2K-48C6-M-R
* CCS-720DT-24S	* DCS-7130-96LA-F (HwRev:A2)	* DCS-7280SR3-40YC6
* CCS-720DT-24S-2	* DCS-7130-96LA-F (HwRev:A3)	* DCS-7280SR3-40YC6-F
* CCS-720DT-24S-2F	* DCS-7130-96LA-R (HwRev:A2)	* DCS-7280SR3-40YC6-R
* CCS-720DT-24S-F	* DCS-7130-96LA-R (HwRev:A3)	* DCS-7280SR3-48YC8
* CCS-720DT-48S-2	* DCS-7130-96LAS	* DCS-7280SR3-48YC8-F
* CCS-720DT-48S-2F	* DCS-7130-96LAS-F (HwRev:B1)	* DCS-7280SR3-48YC8-R
* CCS-720DT-48S-2R	* DCS-7130-96LAS-F (HwRev:B2)	* DCS-7280SR3E-40YC6
* CCS-720XP-24Y6	* DCS-7130-96LAS-R (HwRev:B1)	* DCS-7280SR3E-40YC6-F
* CCS-720XP-24Y6-F	* DCS-7130-96LAS-R (HwRev:B2)	* DCS-7280SR3E-40YC6-R
* CCS-720XP-24Y6-R	* DCS-7130-96LB	* DCS-7280SR3K-48YC8
* CCS-720XP-24ZY4	* DCS-7130-96LB-F (HwRev:A2)	* DCS-7280SR3K-48YC8-F
* CCS-720XP-24ZY4-F	* DCS-7130-96LB-F (HwRev:A3)	* DCS-7280SR3K-48YC8-R
* CCS-720XP-48Y6	* DCS-7130-96LB-R (HwRev:A2)	* DCS-7280SR3K-48YC8A
* CCS-720XP-48Y6-F	* DCS-7130-96LB-R (HwRev:A3)	* DCS-7280SR3K-48YC8A-F
* CCS-720XP-48Y6-R	* DCS-7130-96LBA	* DCS-7280SR3K-48YC8A-R
* CCS-720XP-48ZC2	* DCS-7130-96LBA-F (HwRev:A2)	* DCS-7280SR3M-48YC8
* CCS-720XP-48ZC2-F	* DCS-7130-96LBA-F (HwRev:A3)	* DCS-7280SR3M-48YC8-F

* CCS-720XP-96ZC2	* DCS-7130-96LBA-R (HwRev:A2)	* DCS-7280SR3M-48YC8-R
* CCS-720XP-96ZC2-F	* DCS-7130-96LBA-R (HwRev:A3)	* DCS-7280SRA-48C6
* CCS-722XPM-48Y4	* DCS-7130-96LBAS	* DCS-7280SRA-48C6-F
* CCS-722XPM-48Y4-F	* DCS-7130-96LBAS-F (HwRev:B1)	* DCS-7280SRA-48C6-M-F
* CCS-722XPM-48ZY8	* DCS-7130-96LBAS-R (HwRev:B1)	* DCS-7280SRA-48C6-M-R
* CCS-722XPM-48ZY8-F	* DCS-7130-96LBS	* DCS-7280SRA-48C6-R
* DCS-7010TX-48	* DCS-7130-96LBS-F (HwRev:B1)	* DCS-7280SRAM-48C6
* DCS-7010TX-48-DC	* DCS-7130-96LBS-F (HwRev:B2)	* DCS-7280SRAM-48C6-F
* DCS-7010TX-48-DC-F	* DCS-7130-96LBS-R (HwRev:B1)	* DCS-7280SRAM-48C6-R
* DCS-7010TX-48-DC-R	* DCS-7130-96LBS-R (HwRev:B2)	* DCS-7280SRM-40CX2
* DCS-7010TX-48-F	* DCS-7130-96LS	* DCS-7280SRM-40CX2-F
* DCS-7010TX-48-R	* DCS-7130-96LS-F (HwRev:B1)	* DCS-7280SRM-40CX2-R
* DCS-7020SR-24C2	* DCS-7130-96LS-F (HwRev:B2)	* DCS-7280TR-48C6
* DCS-7020SR-24C2-F	* DCS-7130-96LS-R (HwRev:B1)	* DCS-7280TR-48C6-F
* DCS-7020SR-24C2-R	* DCS-7130-96LS-R (HwRev:B2)	* DCS-7280TR-48C6-M-F
* DCS-7020SR-32C2	* DCS-7130-96S	* DCS-7280TR-48C6-M-R
* DCS-7020SR-32C2-F	* DCS-7130-96S-F (HwRev:B1)	* DCS-7280TR-48C6-R
* DCS-7020SR-32C2-R	* DCS-7130-96S-R (HwRev:B1)	* DCS-7280TRA-48C6
* DCS-7020SRG-24C2	* DCS-7130LBR-48S6QD	* DCS-7280TRA-48C6-F
* DCS-7020SRG-24C2-F	* DCS-7130LBR-48S6QD-F	* DCS-7280TRA-48C6-M-F
* DCS-7020SRG-24C2-R	* DCS-7130LBR-48S6QD-R	* DCS-7280TRA-48C6-M-R
* DCS-7020TR-48	* DCS-7170-32C	* DCS-7280TRA-48C6-R
* DCS-7020TR-48-F	* DCS-7170-32C-F	* FAN-7000
* DCS-7020TR-48-R	* DCS-7170-32C-M-F	* FAN-7000-F
* DCS-7020TRA-48	* DCS-7170-32C-M-R	* FAN-7000-R
* DCS-7020TRA-48-F	* DCS-7170-32C-R	* FAN-7000H-R
* DCS-7020TRA-48-R	* DCS-7170-32CD	* FAN-7001C
* DCS-7050CX3-32S	* DCS-7170-32CD-F	* FAN-7001C-F
* DCS-7050CX3-32S-F	* DCS-7170-32CD-R	* FAN-7001C-R
* DCS-7050CX3-32S-R	* DCS-7170-64C	* FAN-7001D
* DCS-7050CX3-32S-SSD-F	* DCS-7170-64C-F	* FAN-7001D-F
* DCS-7050CX3-32S-SSD-R	* DCS-7170-64C-M-F	* FAN-7001DH
* DCS-7050CX3M-32S	* DCS-7170-64C-M-R	* FAN-7001DH-F
* DCS-7050CX3M-32S-F	* DCS-7170-64C-R	* FAN-7002H
* DCS-7050CX3M-32S-R	* DCS-7170B-64C	* FAN-7002H-R
* DCS-7050DX4-32S	* DCS-7170B-64C-F	* FAN-7010
* DCS-7050DX4-32S-F	* DCS-7260CX-64	* FAN-7010X
* DCS-7050PX4-32S	* DCS-7260CX-64-F	* FAN-7011H
* DCS-7050PX4-32S-F	* DCS-7260CX-64-R	* FAN-7011H-F
* DCS-7050SX3-48C8	* DCS-7260CX-64-SSD-F	* FAN-7011H-R
* DCS-7050SX3-48C8-F	* DCS-7260CX-64-SSD-R	* FAN-7011M

* DCS-7050SX3-48C8-R	* DCS-7260CX3-64	* FAN-7011M-F
* DCS-7050SX3-48YC12	* DCS-7260CX3-64-F	* FAN-7011M-R
* DCS-7050SX3-48YC12-F	* DCS-7260CX3-64-R	* FAN-7012H
* DCS-7050SX3-48YC8	* DCS-7260CX3-64E	* FAN-7012H-BLUE
* DCS-7050SX3-48YC8-F	* DCS-7260CX3-64E-F	* FAN-7012H-RED
* DCS-7050SX3-48YC8-R	* DCS-7260CX3-64E-R	* FAN-7012M
* DCS-7050SX3-96YC8	* DCS-7260CX3-64LQ	* FAN-7012M-BLUE
* DCS-7050SX3-96YC8-F	* DCS-7260CX3-64LQ-F	* FAN-7012M-RED
* DCS-7050SX3-96YC8-R	* DCS-7260CX3-64LQ-R	* FAN-7130-1U
* DCS-7050TX3-48C8	* DCS-7260QX-64	* FAN-7130-1U-F
* DCS-7050TX3-48C8-F	* DCS-7260QX-64-F	* FAN-7130-1U-R
* DCS-7050TX3-48C8-R	* DCS-7260QX-64-R	* FAN-7130-2U
* DCS-7060CX-32C	* DCS-7260QX-64-SSD-F	* FAN-7130-2U-F
* DCS-7060CX-32C-F	* DCS-7260QX-64-SSD-R	* FAN-7130-2U-R
* DCS-7060CX-32C-R	* DCS-7280CR-48	* PWR-1011-AC
* DCS-7060CX-32S	* DCS-7280CR-48-F	* PWR-1011-AC-BLUE
* DCS-7060CX-32S-F	* DCS-7280CR2-60	* PWR-1011-AC-RED
* DCS-7060CX-32S-F (HwRev:13.20)	* DCS-7280CR2-60-F	* PWR-1011-DC
* DCS-7060CX-32S-R	* DCS-7280CR2A-30	* PWR-1011-DC-BLUE
* DCS-7060CX-32S-R (HwRev:13.20)	* DCS-7280CR2A-30-F	* PWR-1011-DC-RED
* DCS-7060CX2-32S	* DCS-7280CR2A-60	* PWR-1021-AC
* DCS-7060CX2-32S-F	* DCS-7280CR2A-60-F	* PWR-1021-AC-RED
* DCS-7060CX2-32S-R	* DCS-7280CR2K-30	* PWR-1100AC
* DCS-7060DX4-32	* DCS-7280CR2K-30-F	* PWR-1100AC-F
* DCS-7060DX4-32-F	* DCS-7280CR2K-60	* PWR-1100AC-R
* DCS-7060DX4-32S	* DCS-7280CR2K-60-F	* PWR-1511-AC
* DCS-7060DX4-32S-F	* DCS-7280CR2M-30	* PWR-1511-AC-BLUE
* DCS-7060DX4-32S-MON	* DCS-7280CR2M-30-F	* PWR-1511-AC-RED
* DCS-7060DX4-32S-MON2	* DCS-7280CR3-32D4	* PWR-1511-DC
* DCS-7060DX4-32S-PEM	* DCS-7280CR3-32D4-F	* PWR-1511-DC-BLUE
* DCS-7060DX4-32SON	* DCS-7280CR3-32D4-M-F	* PWR-1511-DC-RED
* DCS-7060DX4-32SON2	* DCS-7280CR3-32D4-M-R	* PWR-1600AC
* DCS-7060DX5-64	* DCS-7280CR3-32D4-R	* PWR-1600AC-F
* DCS-7060DX5-64-F	* DCS-7280CR3-32P4	* PWR-1611-AC
* DCS-7060DX5-64S	* DCS-7280CR3-32P4-F	* PWR-1611-AC-RED
* DCS-7060DX5-64S-F	* DCS-7280CR3-32P4-M-F	* PWR-1611-DC
* DCS-7060DX5-64S-F (HwRev: 12.00)	* DCS-7280CR3-32P4-M-R	* PWR-1611-DC-RED
* DCS-7060DX5-64S-R	* DCS-7280CR3-32P4-R	* PWR-1900-DC
* DCS-7060DX5-64S-R (HwRev: 12.00)	* DCS-7280CR3-36S	* PWR-1900-DC-F
* DCS-7060PX4-32	* DCS-7280CR3-36S-F	* PWR-1900-DC-R
* DCS-7060PX4-32-F	* DCS-7280CR3-36S-R	* PWR-1900AC

* DCS-7060PX5-64E	* DCS-7280CR3-96	* PWR-1900AC-F
* DCS-7060PX5-64E-F	* DCS-7280CR3-96-F	* PWR-1900AC-R
* DCS-7060SX2-48YC6	* DCS-7280CR3K-32D4	* PWR-2021-AC
* DCS-7060SX2-48YC6-F	* DCS-7280CR3K-32D4-F	* PWR-2021-AC-RED
* DCS-7060SX2-48YC6-R	* DCS-7280CR3K-32D4-R	* PWR-2411-AC
* DCS-7130-16G3	* DCS-7280CR3K-32D4A	* PWR-2411-AC-BLUE
* DCS-7130-16G3-F (HwRev:C1)	* DCS-7280CR3K-32D4A-F	* PWR-2411-AC-RED
* DCS-7130-16G3-R (HwRev:C1)	* DCS-7280CR3K-32D4A-R	* PWR-2411-DC
* DCS-7130-16G3S	* DCS-7280CR3K-32P4	* PWR-2411-DC-BLUE
* DCS-7130-16G3S-F (HwRev:D1)	* DCS-7280CR3K-32P4-F	* PWR-2411-DC-RED
* DCS-7130-16G3S-R (HwRev:D1)	* DCS-7280CR3K-32P4-R	* PWR-3001-AC
* DCS-7130-48E	* DCS-7280CR3K-32P4A	* PWR-3001-AC-BLUE
* DCS-7130-48E-F (HwRev:A4)	* DCS-7280CR3K-32P4A-F	* PWR-3001-AC-RED
* DCS-7130-48E-R (HwRev:A4)	* DCS-7280CR3K-32P4A-R	* PWR-3001-DC
* DCS-7130-48EH	* DCS-7280CR3K-36A	* PWR-3001-DC-BLUE
* DCS-7130-48EH-F (HwRev:A4)	* DCS-7280CR3K-36A-F	* PWR-3001-DC-RED
* DCS-7130-48EH-R (HwRev:A4)	* DCS-7280CR3K-36A-R	* PWR-400-DC
* DCS-7130-48EHS	* DCS-7280CR3K-36S	* PWR-400-DC-BLUE
* DCS-7130-48EHS-F (HwRev:B1)	* DCS-7280CR3K-36S-F	* PWR-400-DC-RED
* DCS-7130-48EHS-F (HwRev:B2)	* DCS-7280CR3K-36S-R	* PWR-400AC
* DCS-7130-48EHS-R (HwRev:B1)	* DCS-7280CR3K-96	* PWR-400AC-BLUE
* DCS-7130-48EHS-R (HwRev:B2)	* DCS-7280CR3K-96-F	* PWR-400AC-RED
* DCS-7130-48G3	* DCS-7280CR3MK-32D4	* PWR-460AC
* DCS-7130-48G3-F (HwRev:B1)	* DCS-7280CR3MK-32D4-F	* PWR-460AC-F
* DCS-7130-48G3-F (HwRev:B2)	* DCS-7280CR3MK-32D4-R	* PWR-460AC-R
* DCS-7130-48G3-R (HwRev:B1)	* DCS-7280CR3MK-32D4S	* PWR-460DC
* DCS-7130-48G3-R (HwRev:B2)	* DCS-7280CR3MK-32D4S-F	* PWR-460DC-F
* DCS-7130-48G3S	* DCS-7280CR3MK-32D4S-R	* PWR-460DC-R
* DCS-7130-48G3S-F (HwRev:C1)	* DCS-7280CR3MK-32P4	* PWR-500-DC
* DCS-7130-48G3S-R (HwRev:C1)	* DCS-7280CR3MK-32P4-F	* PWR-500-DC-F
* DCS-7130-48L	* DCS-7280CR3MK-32P4-R	* PWR-500-DC-R
* DCS-7130-48L-F (HwRev:A2)	* DCS-7280CR3MK-32P4S	* PWR-500AC
* DCS-7130-48L-F (HwRev:A3)	* DCS-7280CR3MK-32P4S-F	* PWR-500AC-F
* DCS-7130-48L-R (HwRev:A2)	* DCS-7280CR3MK-32P4S-R	* PWR-500AC-R
* DCS-7130-48L-R (HwRev:A3)	* DCS-7280DR3-24	* PWR-501AC
* DCS-7130-48LA	* DCS-7280DR3-24-F	* PWR-501AC-F
* DCS-7130-48LA-F (HwRev:A2)	* DCS-7280DR3-24-M-F	* PWR-501AC-R
* DCS-7130-48LA-F (HwRev:A3)	* DCS-7280DR3A-54	* PWR-511-AC
* DCS-7130-48LA-R (HwRev:A2)	* DCS-7280DR3A-54-F	* PWR-511-AC-BLUE
* DCS-7130-48LA-R (HwRev:A3)	* DCS-7280DR3AK-54	* PWR-511-AC-RED
* DCS-7130-48LAS	* DCS-7280DR3AK-54-F	* PWR-511-DC

* DCS-7130-48LAS-F (HwRev:B2)	* DCS-7280DR3AM-54	* PWR-511-DC-BLUE
* DCS-7130-48LAS-F (HwRev:B3)	* DCS-7280DR3AM-54-F	* PWR-511-DC-RED
* DCS-7130-48LAS-R (HwRev:B2)	* DCS-7280DR3K-24	* PWR-621-AC
* DCS-7130-48LAS-R (HwRev:B3)	* DCS-7280DR3K-24-F	* PWR-621-AC-BLUE
* DCS-7130-48LB	* DCS-7280PR3-24	* PWR-621-AC-RED
* DCS-7130-48LB-F (HwRev:A2)	* DCS-7280PR3-24-F	* PWR-721-DC
* DCS-7130-48LB-F (HwRev:A3)	* DCS-7280PR3-24-M-F	* PWR-721-DC-RED
* DCS-7130-48LB-R (HwRev:A2)	* DCS-7280PR3K-24	* PWR-745AC
* DCS-7130-48LB-R (HwRev:A3)	* DCS-7280PR3K-24-F	* PWR-745AC-F
* DCS-7130-48LBA	* DCS-7280QR-C36	* PWR-745AC-R
* DCS-7130-48LBA-F (HwRev:A2)	* DCS-7280QR-C36-F	* PWR-747AC
* DCS-7130-48LBA-F (HwRev:A3)	* DCS-7280QR-C36-M-F	* PWR-747AC-F
* DCS-7130-48LBA-R (HwRev:A2)	* DCS-7280QR-C36-M-R	* PWR-747AC-R
* DCS-7130-48LBA-R (HwRev:A3)	* DCS-7280QR-C36-R	* PWR-750AC
* DCS-7130-48LBAS	* DCS-7280QR-C72	* PWR-750AC-F
* DCS-7130-48LBAS-F (HwRev:B2)	* DCS-7280QR-C72-F	* PWR-750AC-R
* DCS-7130-48LBAS-F (HwRev:B3)	* DCS-7280QR-C72-M-F	* FAN-7000H
* DCS-7130-48LBAS-R (HwRev:B2)	* DCS-7280QR-C72-M-R	* FAN-7000H-F
* DCS-7130-48LBAS-R (HwRev:B3)	* DCS-7280QR-C72-R	* FAN-7002
* DCS-7130-48LBS	* DCS-7280QRA-C36S	* FAN-7002-F
* DCS-7130-48LBS-F (HwRev:B2)	* DCS-7280QRA-C36S-F	* FAN-7002-R
* DCS-7130-48LBS-F (HwRev:B3)	* DCS-7280QRA-C36S-M-F	* FAN-7002H-F
* DCS-7130-48LBS-R (HwRev:B2)	* DCS-7280QRA-C36S-M-R	* FAN-7012MP
* DCS-7130-48LBS-R (HwRev:B3)	* DCS-7280QRA-C36S-R	* FAN-7012MP-RED
* DCS-7130-48LS		

Modular

* FAN-7000H	* 7500R-8CFPX-LC	* CCS-750X-48ZP-LC
* FAN-7000H-F	* 7500R2-18CQ-LC	* CCS-750X-48ZXP-LC
* FAN-7002	* 7500R2-36CQ-LC	* CCS-755-CH
* FAN-7002-F	* 7500R2A-36CQ-LC	* CCS-755-X3-SC
* FAN-7002-R	* 7500R2AK-36CQ-LC	* CCS-758-CH
* FAN-7002H-F	* 7500R2AK-48YCQ-LC	* CCS-758-X3-SC
* FAN-7012MP	* 7500R2AM-36CQ-LC	* DCS-7304
* FAN-7012MP-RED	* 7500R2M-36CQ-LC	* DCS-7308
* 7289	* 7500R3-24D-LC	* DCS-7500-SUP2
* 7289-SUP	* 7500R3-24P-LC	* DCS-7500-SUP2-D
* 7289-SUP-D	* 7500R3-36CQ-LC	* DCS-7504
* 7289-SUP-S	* 7500R3-48Y4D-LC	* DCS-7504N
* 7289-SUP-S-D	* 7500R3K-36CQ-LC	* DCS-7508
* 7289R3A-SC	* 7500R3K-48Y4D-LC	* DCS-7508N
* 7289R3AK-SC	* 7500RM-36CQ-LC	* DCS-7512N
* 7289R3AM-SC	* 7504R-FM	* DCS-7516-SUP2
* 7300-SUP	* 7504R3-FM	* DCS-7516-SUP2-D
* 7300-SUP-D	* 7508R-FM	* DCS-7516N
* 7300-SUP2	* 7508R3-FM	* DCS-7800-SUP
* 7300-SUP2-D	* 7512R-FM	* DCS-7800-SUP1A
* 7300X3-32C-LC	* 7512R3-FM	* DCS-7800-SUP1S
* 7300X3-48TC4-LC	* 7516R-FM	* DCS-7804-CH
* 7300X3-48YC4-LC	* 7800R3-36D-LC	* DCS-7808-CH
* 7304X3-FM	* 7800R3-36P-LC	* DCS-7812-CH
* 7304X3-FM2	* 7800R3-48CQ-LC	* DCS-7816-CH
* 7308X3-FM	* 7800R3-48CQ2-LC	* DCS-7816-SUP
* 7308X3-FM2	* 7800R3-48CQM-LC	* DCS-7816-SUP1S
* 7320X-32C-LC	* 7800R3-48CQM2-LC	* FAN-752M
* 7324X-FM	* 7800R3-48CQMS-LC	* FAN-752M-RED
* 7328X-FM	* 7800R3A-36D-LC	* PWR-2422-HV
* 7358	* 7800R3A-36D2-LC	* PWR-2422-HV-RED
* 7358-16C	* 7800R3A-36DM-LC	* PWR-2700-DC
* 7358X4-SC	* 7800R3A-36DM2-LC	* PWR-2700-DC-F
* 7368	* 7800R3A-36P-LC	* PWR-2700-DC-R
* 7368-16C	* 7800R3A-36PM-LC	* PWR-2900AC
* 7368-16S	* 7800R3AK-36DM-LC	* PWR-3351-AC
* 7368-4D	* 7800R3AK-36DM2-LC	* PWR-3351-AC-RED
* 7368-4P	* 7800R3AK-36PM-LC	* PWR-3K-AC
* 7368-SUP	* 7800R3K-36DM-LC	* PWR-3K-AC-F

* 7368-SUP-D	* 7800R3K-48CQ-LC	* PWR-3K-AC-R
* 7368X4-SC	* 7800R3K-72Y-LC	* PWR-3K-DC
* 7388	* 7804R3-FM	* PWR-3K-DC-BLUE
* 7388-16CD	* 7808R3-FM	* PWR-3K-DC-RED
* 7388-16CD2	* 7808R3-FM2	* PWR-3KT-AC
* 7388-8D	* 7812R3-FM	* PWR-3KT-AC-BLUE
* 7388-SUP	* 7816R3-FM	* PWR-3KT-AC-RED
* 7388-SUP-D	* CCS-750-SUP100	* PWR-D1-3041-AC
* 7388X5-SC	* CCS-750-SUP25	* PWR-D1-3041-AC-BLUE
* 7500R-36CQ-LC	* CCS-750X-48THP-LC	* PWR-D4-3041-AC
* 7500R-36Q-LC	* CCS-750X-48TP-LC	* PWR-D4-3041-AC-BLUE
* 7500R-48S2CQ-LC		

Transceiver

* 1000BASE-BX10	* 25GBASE-SR	* OSFP-400G-LR8
* 1000BASE-BX10-D	* 25GBASE-XSR	* OSFP-400G-PLR4
* 1000BASE-BX10-U	* 40GBASE-AOC	* OSFP-400G-PLR4-S
* 1000BASE-LX	* 40GBASE-CR4	* OSFP-400G-SR8
* 1000BASE-SX	* 40GBASE-ER4	* OSFP-400G-XDR4
* 1000BASE-T	* 40GBASE-LR4	* OSFP-400G-XDR4-S
* 100BASE-FX	* 40GBASE-LRL4	* OSFP-400G-ZR
* 100GBASE-AOC	* 40GBASE-PLR4	* OSFP-400G-ZRP
* 100GBASE-CR4	* 40GBASE-PLRL4	* OSFP-800G-2FR4
* 100GBASE-CWDM4	* 40GBASE-SR4	* OSFP-800G-2LR4
* 100GBASE-DR	* 40GBASE-SRBD	* OSFP-800G-2PLR4
* 100GBASE-ERL4	* 40GBASE-UNIV	* OSFP-800G-2XDR4
* 100GBASE-FR	* 40GBASE-XSR4	* OSFP-AMP-ZR
* 100GBASE-LR	* 50GBASE-CR	* QDD-200G-2LR4
* 100GBASE-LR4	* ADPT-O-Q-100G	* QDD-400G-DR4
* 100GBASE-LRL4	* AOC-D-D-400G	* QDD-400G-FR4
* 100GBASE-PLR4	* AOC-O-O-400G	* QDD-400G-LR4
* 100GBASE-PLRL4	* CAB-D-2Q-200G	* QDD-400G-LR8
* 100GBASE-PSM4	* CAB-D-2Q-400G	* QDD-400G-PLR4
* 100GBASE-SR4	* CAB-D-4Q-200G	* QDD-400G-PLR4-S
* 100GBASE-SRBD	* CAB-D-4Q-400G	* QDD-400G-SR8
* 100GBASE-SWDM4	* CAB-D-8S-200G	* QDD-400G-SR8-C
* 100GBASE-XCWDM4	* CAB-D-8S-400G	* QDD-400G-XDR4
* 100GBASE-XSR4	* CAB-D-D-400G	* QDD-400G-XDR4-S
* 100GBASE-ZR4	* CAB-O-2Q-200G	* QDD-400G-ZR
* 100GE-DWDM2	* CAB-O-2Q-400G	* QDD-400G-ZRP
* 10GBASE-AOC	* CAB-O-4Q-200G	* QDD-800G-2FR4
* 10GBASE-CR	* CAB-O-4Q-400G	* QDD-800G-2LR4
* 10GBASE-DWDM	* CAB-O-8S-200G	* QDD-800G-2PLR4
* 10GBASE-DWDM-T	* CAB-O-8S-400G	* QDD-800G-2XDR4
* 10GBASE-DWDM-ZR	* CAB-O-O-400G	* QSFP-100G-ER
* 10GBASE-ER	* CAB-Q-2Q-100G	* QSFP-200G-AR4
* 10GBASE-ERBD	* CFP2-100G-DWDM-DCO	* QSFP-200G-FR4
* 10GBASE-ERBD-D	* CFP2-100GBASE-ER4	* QSFP-AMP-ZR
* 10GBASE-ERBD-U	* CFP2-100GBASE-LR4	* SFP-10G-MRA-T
* 10GBASE-ERLBD	* CFP2-100GBASE-XSR10	* SFP-10G-RA-1G-LX
* 10GBASE-ERLBD-D	* CFPX-100G-DWDM	* SFP-10G-RA-1G-SX
* 10GBASE-ERLBD-U	* CFPX-200G-DWDM	* SFP-10G-T
* 10GBASE-LR	* E-D800-D800	* SFP-10GRA-T

* 10GBASE-LRL	* E-O800-O800	* SFP-25G-LR-E
* 10GBASE-SR	* H-D400-4Q100	* SFP-25G-MR-LR
* 10GBASE-SRL	* H-O400-4Q100	* SFP-25G-MR-SR
* 10GBASE-ZR	* OSFP-200G-2LR4	* SFP-25G-MR-XSR
* 200GBASE-CR4	* OSFP-400-2FR4-LC	* SFP-25G-SR-E
* 200GBASE-SR4	* OSFP-400G-2FR4	* SFP-25GR-LR
* 25GBASE-AOC	* OSFP-400G-DR4	* SFP-25GR-SR
* 25GBASE-CR	* OSFP-400G-FR4	* SFP-25GR-XSR
* 25GBASE-LR	* OSFP-400G-LR4	

Reference to MLAG ISSU Upgrade/Downgrade Table

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

Resolved Caveats in 4.29.9.2M

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7260X, DCS-7260X3 and DCS-7300X Series

DCS-7060X5 Series

- Forwarding agent "Strata-FixedSystem" may restart when a large number of ports are shutdown or unshutdown using interface range CLI and if interfaces have one or more TCAM based features applied to them like security ACL, policy QoS, storm-control, etc.
Agent recovers after restart. (821616)

Known Software Caveats in 4.29.9.2M

AristaSAI

- A programming of a dynamic default route with no nexthop will cause a crash of the orchagent. To avoid this crash, a route policy map rule can be installed to prevent such routes from being installed. (969989)
- When all the members of a nextHop group are deleted, the group's state is updated to DROP, which does not allow the route to be in FORWARD state to a new nextHop, unless route action is set to FORWARD explicitly, before trying to program the new nextHop. The route action, with which the route was programmed with, can be stashed and used instead of using the current state of the nextHop to help avoid this issue. (1022352)

Accelerated System Update

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

- SSU can cause a momentary change in ECMP hashing, leading to extra traffic loss (738594)
- In certain EVPN deployments where the VTEPs share the same virtual IPV6 address, it is possible that the devices after undergoing hitless reboot are unable to claim the virtual interface address due to an otherwise avoidable IPV6 Duplicate Address Detection.

As a workaround to this issue, the SVIs with Duplicate Address Detection failure can be flapped to successfully register the virtual IPV6 address for the SVI. (832854)

- A system may reboot during an ASU if at least one CMIS transceiver (i.e. QSFP200, OSFP, QSFP-DD) is unresponsive. (845097)
- Performing SSU when OSPF router ID is not explicitly configured may result in extended traffic outage.

A workaround is to configure router ID before performing SSU (856714)

- Interfaces in a port-channel that are configured with `no switchport` may flap during SSU.
The workaround is to configure `default switchport` on port-channel interfaces (860647)
- QSFP optical transceivers may experience a link flap after configuring "reload fastboot". (866709)
- Connections using tunable SFP transceivers like the 10GBASE-DWDM-T and 10GBASE-DWDM-ZR may flap during 'reload fastboot'. There is no workaround, but after the flap(s), the link should be operational again. (878971)
SKUs Affected: 10GBASE-DWDM-T and 10GBASE-DWDM-ZR
- During SSU, network churn events (such as remote VTEP flaps and route churns) may result in extended traffic outage. (927699)
- Performing SSU upgrade from an unsupported release of CCS-750X-48SX-LC may result in a fatal error reboot causing extended traffic outage.

The workaround is to remove the CCS-750X-48SX-LC Linecard before starting the upgrade. (932515)

SKUs Affected: CCS-755-CH and CCS-758-CH

- During ASU, subinterfaces can see updates to the status of the corresponding kernel device. This can cause a flap in the BGP peer. (948901)
- SSU for the below release combinations is not supported:
 - From the source release 4.28.8.3M and later releases in the 4.28 train to the target releases in the range between 4.29.0F (inclusive) and 4.29.8 (exclusive) <- this is the first release the patch was casted to 4.29 train

- From the source release 4.28.8.3M and later releases in the 4.28 train to the target releases in the range between 4.30.0F (inclusive) and 4.30.6 (inclusive)
 - From the source release 4.28.8.3M and later releases in the 4.28 train to the target releases in the range between 4.31.0F (inclusive) and 4.31.2 (inclusive)
 - From the source release 4.28.8.3M and later releases in the 4.28 train to the target releases in the range between 4.32.0F (inclusive) and 4.32.1F (exclusive)
 - From the source release 4.29.8M and later releases in the 4.29 train to the target releases in the range between 4.30.0F (inclusive) and 4.30.6 (inclusive)
 - From the source release 4.29.8M and later releases in the 4.29 train to the target releases in the range between 4.31.0F (inclusive) and 4.31.2 (inclusive)
 - From the source release 4.29.8M and later releases in the 4.29 train to the target releases in the range between 4.32.0F (inclusive) and 4.32.1F (exclusive)
 - From the source release 4.31.3M and later releases in the 4.31 train to the target releases in the range between 4.32.0F (inclusive) and 4.32.1F (exclusive) (960833)
- If an interface is configured with "phy link training", then a SSU reload may result in a fatal error reboot that will cause extended traffic outage. (983323) (1)
 - Hardware BFD acceleration is not correctly reinitialized during the SSU which may cause a post-upgrade device reload.

Workaround: disable hardware BFD offload prior to SSU until the next full reload. (1001142)

- When performing SSU upgrade from releases older than EOS-4.29.2F to the affected releases, SSU upgrade may fail with fatal error, if at least one of the interfaces is errdisabled with the speed-misconfigured cause. Fatal error will result in a regular reload, causing link flaps and traffic disruption.

The workaround is to remove the invalid speed configurations from the system, prior to performing the SSU upgrade. (1090095)

SKUs Affected: DCS-7050CX3-32S, DCS-7050SX3-48YC12 and DCS-7050SX3-96YC8

BGP EVPN L2/L3 VXLAN/MPLS

- On a BGP router, EVPN creates dynamic VLANs for L3 VNIs on receiving EVPN RT2 or RT5 advertisements even when the router does not have any L3 VRF or import route-target configurations. (294328)
- A policy change (such as changing the color extended communities) which causes an L2 EVPN MAC-IP or IMET route to resolve over a different SR-TE policy or tunnel may cause a brief traffic loss. (521090)
- A remote VTEP reached through unnumbered BGP over IPv6 will be incorrectly displayed as a configuration error with "FAIL" state due to "No route" (525842)
- Some MAC addresses may get blacklisted when the BGP agent is restarted on a multihomed peer. The workaround is to clear the blacklist using the "clear bgp evpn host-flap" command. (538177)
- In an EVPN multi-homing deployment, designated forwarder election for a given Ethernet segment and EVPN instance pair is based on the complete set of VLANs within the EVPN instance. This differs from the standard in which only the instance VLANs that are in the Ethernet segment interface's allowed VLAN list are used in the election algorithm. This deviation may cause interoperability issues with other vendors. It may also cause devices to be elected DF for an EVI on which they can't forward, if there is disjoint configuration of allowed VLANs for the ES interface on different peers in the same ES.

As a workaround, a single VLAN-aware bundle EVPN instance may need to be split into multiple instances, such that for every instance and Ethernet segment either every VLAN is in the Ethernet segment interface's allowed VLAN list, or none are. (583126)

- On Multi-Domain EVPN VXLAN deployment using anycast gateways, administratively enabling and disabling the VXLAN interface may cause and ARP entry to be lost and trigger traffic lost. (632955)
- In an L2 EVPN MPLS configuration, disabling MAC address learning on a VLAN that is part of the EVPN bridging domain does not take effect. (677182)
Series Affected: 7500R3, DCS-7280CR2, DCS-7280CR3, DCS-7280DR3, DCS-7280PR3, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280SR3, DCS-7280TR, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 Series
- When configuring "set_nexthop_self()" using RCF in the EVPN outbound context, the BGP agent can restart. (717957)

- In an EVPN VXLAN deployment using gateway overlay index, flapping configuration for "route import overlay-index gateway" may cause vias for overlay-index routes to become empty, as seen in the "show ip route <prefix>" command. This situation can be avoided by waiting for ip route convergence between the "no route import overlay-index gateway" and "route import overlay-index gateway" commands.

If the situation is encountered, as a workaround the routes can be reset to the functioning state by resetting the BGP session with "clear ip bgp". (723279)

- In an EVPN VXLAN deployment, the BGP process can leak memory when the VLAN to VNI mapping is continuously provisioned and unprovisioned. The amount of memory leaked does not scale with the number of routes but instead with the number of MAC-VRFs. (774066)
- In a EVPN VXLAN deployment with L2 ARP or ND proxy enabled ("arp learning bridged" for IPv4 hosts and/or "nd learning bridged" for IPv6 hosts), when a GARP is generated by a host to update its MAC in an ARP or ND binding, the multi-homed PEs may advertise stale and conflicting MAC to ARP/ND binding. (838924)
- If a DCI gateway changed from connecting VXLAN and VXLAN domains to connecting VXLAN and MPLS domains, EVPN MAC-IP and IMET routes sent by this gateway may be missing on the remote DCI gateways or route reflectors.

Use "clear ip bgp <neighbor> soft in" on the remote gateways or route reflectors to recover the missing EVPN routes. (841750)

- On an EVPN router, the system might run out of dynamic VLANs available, in the event of a VXLAN VRF VNI mapping configurations churn.

To workaround the problem, restart the BGP agent. (846021)

- In a EVPN VXLAN deployment where L2 ARP and ND proxy is enabled, executing the "clear arp|nd bridged" or the "no arp|nd learning bridged" command might not clear all the locally learned L2 ARP entries in the "show arp bridged" or "show ipv6 neighbors bridged" output. (882754)
- Changing certain MLAG configuration parameters (MLAG domain ID, local MLAG interface, MLAG peer address, MLAG peer interface) will cause the MLAG session to temporarily go down while the connection is renegotiated. This can lead to transient L2 loops in the network until the MLAG session is re-established. If EVPN is also being used, this can lead to rapid MAC flaps between devices in the network, causing those MACs to become blacklisted by EVPN. To see if a particular MAC has been added to the EVPN blacklist, use the "show bgp evpn host-flap" command.

As a workaround, configuring EVPN host-flap recovery will allow the

blacklisted MACs to age out after a configured time once the MLAG session has been re-established. Blacklisted MACs may also be cleared manually with the "clear bgp evpn host-flap" command. (891929) (1)

- If we have multihoming PEs PE1 and PE2, upon the Bgp agent restarting on PE1, for a host with an ARP entry locally learned on PE1 and with a MAC locally learned on PE2, Bgp may not advertise the type 2 MAC - IP route or the type 2 MAC only proxy route for this host upon restart even though the ARP entry will still be learned on PE1 and the MAC will still be installed in the MAC table on PE1.

Workaround is to restart the Bgp agent on PE1 until the problem resolves or flap the vlan interface this is occurring on. On Sand platforms, having a host send an ARP reply to PE1 or having the host send local traffic to Pe2 will fix the issue. (904679)

- In an EVPN VXLAN IRB MLAG deployment, out-of-order processing of dynamic VLAN-VNI mappings received on the secondary MLAG peer can cause some EVPN routes to be unresolved.

A workaround is to terminate the Vxlan agent. (957183)

- In a EVPN VXLAN deployment, if an L2 VTEP is configure with "arp|nd learning bridged" receives a VXLAN encapsulated NA in response to a IPv6 DAD request, the response may not be decapsulated and bridged.

A encapsulated NA to a solicited NS will decapsulated and bridged as expected. (976554)

- On MLAG devices running EVPN OISM, `mlag shut` `no shut` may result in traffic loss.
On devices running EVPN OISM in a non MLAG scenario could result in PIM neighborship being formed across the EVPN fabric, this could prevent traffic from being routed and result in traffic loss. (990767) (1)
- When transitioning IPv6 L3 VARP VTEP functionality from one VTEP to another, IPv6 neighbor entries may not be properly updated from being remotely learnt through EVPN, to being locally learnt. This may cause traffic to be encapsulated with stale MAC addresses. (1005904)
- In EVPN deployments using the overlay index feature, rapid configuration changes on a remote device may cause learned IP prefixes to fail to resolve.

As a workaround, clear the BGP session. (1016559)

- In EVPN VXLAN OISM deployment the VxlanSwFwd agent may restart modifying the VXLAN flood group in a configuration session. (1026200)

- When migrating from a VCS deployment with virtual gateway IP to an EVPN VXLAN A-IRB deployment, a per VLAN to MAC VRF migration might result in the failure to resolve ARP or ND of a remote host due to the fact that ARP and ND synchronization is disabled in an EVPN VXLAN environment.

The workaround is to enable VCS arp proxy prior to migration:

```
interface vxlan 1
    vxlan controller-client arp proxy
    vxlan controller-client nd proxy (1044263)
```

- On an EVPN L3 gateway, BGP agent may crash if a RCF outbound policy is configured to change the next-hop of the paths. (1058549)
- If we have 2 multihoming EVPN active-active peers, peer A and peer B, if an ARP or ND entry is learned only on peer A and then removed from peer A and there is continuous TCP traffic between the peer B and the host, if the host does not respond to the first ARP packet sent out by peer B or the ARP packet is dropped, peer B may permanently advertise a proxy MAC-IP route for the host for as long as the TCP traffic is ongoing.

Workaround is to clear the ARP or ND entry from peer B. (1076292)

- In an EVPN MPLS setup, a PE may never age out an ARP entry as long as the corresponding MAC address exists in the MAC table. In some cases, if the entry was originally on a different PE and moved to this PE, the PE may also repeatedly send out ARP requests for an entry.

Note that this does not affect IPv6 neighbors.

Workaround is to shut/no shut the interface the ARP entry is learned on. (1076851)

- In EVPN Multicast transit network with a RP on a stick topology, some multicast receivers connected to the PIM External Gateway that is on the path to the source from the RP may not get traffic. The workaround is to add a receiver on another PIM External Gateway. (1079151) (1)
- In an EVPN MPLS setup, if a host was originally on a different PE and moved to this PE, the PE may repeatedly send out ARP requests for the entry.

Note that this does not affect IPv6 neighbors.

Workaround is to shut/no shut the interface the ARP entry is learned on. (1081385)

- In EVPN MPLS multihome active-active setups, if an IPv6 device changes IPv6 address, all active multi home peers will repeatedly send NS packets for the old IPv6 address and will repeatedly switch between advertising proxy MAC-IP routes and non-proxy MAC-IP routes.

Workaround is to add a static drop MAC on all but one of the multihoming devices until all problematic ND entries are removed if "redistribute static" is not configured. If "redistribute static" is configured, workaround is shut the vlan on all multihoming devices until all problematic ND entries are removed. (1081806)

- When a host's ARP entry is learned on two MH A-A PEs and the host changes MAC address and sends an ARP packet with the same IP and the new MAC, the PE receiving the ARP packet may delete its local ARP entry and replace it with a remote ARP entry with the old MAC instead of updating the ARP entry with the new MAC. This will cause traffic destined to the IP to continue to be sent to the old MAC

Workaround is to add an RCF rule to not advertise EVPN type 2 routes for hosts whose MACs are expected to change or have the host resend the ARP packet upon changing MACs. (1117682)

- In an EVPN VXLAN deployment, configuring a VLAN-VNI mapping without creating corresponding VLANs could prevent some dynamic VLANs from being allocated. (1118105) (1)

(1) This item is in two or more sections on this document

vEOS Router / CloudEOS

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

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

- A reboot due to a kernel panic can happen during first 10 minutes of vEOS bootup on Microsoft Azure. The kernel panic occurs because of a task blocking without being scheduled for more than 300 seconds in "D" state. (249618)
- Interface may be allowed to be configured with un-supported speed of NIC using "speed <speed>", which will cause an unexpected PhyEthtool agent restart. Do not use "speed <speed>" command to configure interface speed that NIC does not support. (264395)

- When 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)
- 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 an IPsec connection is established with a primary IP address of an interface and a secondary IP address is configured on an interface and that secondary IP address is used to establish an IPsec connection, the IPsec connection with the secondary IP address doesn't come up.
The workaround is to restart the IPsec agent. (572808)
- When running CloudEOS on ESXI hypervisor, while using Intel Ethernet Controller X710 NIC in SRIOV mode, the TX traffic out of the interface, may stop working after either the router reload and/or SFE agent (Software Forwarding Engine) restart for any reason.
The workaround is to set "trust-mode" ON and "spoof-check" OFF for the SRIOV interface on the Vmware hypervisor.
Another possible workaround is to log into the ESXI VCenter console and change the SRIOV port's assigned VLAN to any other VLAN and then revert it back to the original VLAN. (647410)
- The /var/log/ filesystem on AWS CloudEOS VMs may get filled up with awslogs.*backup files. The workaround is to delete the /etc/cron.d/awslogs_log_rotate file. This workaround is better applied using an on-boot event handler that deletes this file on every reboot of the VM. (713562)
- CloudHA uses older Boto SDK version than some of the AWS regions allow. CloudHA may not work in such AWS regions.

The workaround is to create a boto config file as indicated below and follow it with shut/no shut action on the CloudHA feature. This file needs to be created in the designated path upon every restart. An event-handler can be used to achieve it.

Boto config file content with the exact path:

```
-----
[ec2-user@localhost ~]$ cat /etc/boto.cfg
[Boto]
use_endpoint_heuristics = True
```

```

-----

Event handler configuration:
-----

event-handler AWS
    trigger on-boot
    action bash sudo echo -e "[Boto]\nuse_endpoint_heuristics = True" > /etc/
boto.cfg
    delay 0
----- (758301)

```

- When auto-discovery of path MTU is enabled in a DPS path-group, and an IPsec profile with tunnel mode is configured in the same path-group, Sfe agent may restart unexpectedly.
Workaround is to avoid such a configuration by configuring a static path MTU value in the path-group or changing the IPsec profile from Tunnel mode to Transport mode. (811452)
- When static tunnel interface and autovpn dynamic path are configured with the same source ip address and destination ip address, then Ipsec agent asserts that these two types of connections with the same connection parameters are not allowed at the same time.

The workaround is to shutdown or delete the static tunnel interface.
(980193)

CVP

Network Provisioning

- ZeroTouch agent may keep restarting continuously during Zero Touch Provisioning (ZTP) if the DHCP option 67 (boot file name) is a hostname that needs to be resolved as opposed to an IPv4 or IPv6 address. This will result in ZTP getting stuck. The workaround is to use an IPv4 or IPv6 address in the boot file URL. (718251)
- During Zero Touch Provisioning(ZTP), repeated failures to execute a bootstrap script that forks other scripts may result in the switch running out of memory.
The workaround is to rewrite the bootstrap script to either avoid forking or properly cleanup forked child processes when the script exits, and then reloading the switch to restart ZTP. (820844)
- Device enrollment and consequently ZTPaaS onboarding could fail for a device, that was enrolled to CVaaS before, with the error: "Failure when receiving data from the peer". The workaround is to delete the TerminAttr certs stored at /persist/secure/ssl/terminattr and attempt ZTPaaS onboarding again. (916124)

- ZeroTouch provisioning over DHCPv6 is not possible if DAD failed for any IPv6 link-local address. (1092540)

CVX

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

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

- When a security policy on a CheckPoint firewall includes unsupported services, the Macro-Segmentation Service (MSS) intercepts traffic for all services for the specified end-points in the policy. (433067)
- When the Policy Control Service is enabled on CVX, under certain circumstances such as re-numbering IP on a host connected to a tagged port, the groups received from NSX controller may not be resolved into IP addresses. This may result in incorrect programming of ACLs on the switches.

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

- When MSS is configured in a VXLAN routing deployment with VXLAN controller service (VCS), an MLAG ISSU upgrade from a pre 4.22.1 release to a post 4.22.1 will be hitful and MLAG roles will not be re-established until both MLAG peers have been upgraded.

The workaround is to disable MSS from the VCS prior to the ISSU upgrade. Note that disabling MSS will cause ARP entries on the switch to be relearned and is disruptive. (694514)

- In a VCS deployment with MLAG, a rapid MAC move may result in a MAC being learned locally on multiple VTEPs. The workaround is to run "clear mac address-table dynamic vlan VLAN" on the MLAG secondary which should not have the MAC learned locally. (702518)
- The MCS agent might restart when sending notifications after processing multiple POST delSenders and delReceivers requests. (767011)
- In a VXLAN deployment using VXLAN Controller Service (VCS) to distribute MAC addresses, the MAC address of an orphan host behind a secondary MLAG switch may not get published to VCS. (770762)
- If MCS client agent restarts, it could affect multicast traffic flows for the secondary cvx. (783230)
- The Hardware Switch Controller (HSC) service available on CVX deployments will no longer be supported.

For existing deployments using this feature, please continue to run and upgrade EOS to the latest release on the release train where this defect remains open (that is a release where the feature has not been removed). (950884)

- The VXLAN Controller Service (VCS) feature available on CVX deployments will no longer be supported.

Please migrate the deployment to using an BGP/EVPN control plane. If the deployment must continue to use VCS, please continue to upgrade EOS to the latest version in the train where this defect remains open (that is a release where the feature has not been removed). (950887)

- The Policy Control Service (PCS) available on CVX deployments will no longer be supported.

The feature is being deprecated and the solution shall no longer be supported by Arista or the partner vendor. In the event that you are impacted by the feature deprecation, please reach out to your account team or TAC. (960209)

DANZ Monitoring Fabric

- The CLI command "reset management dmf controller sync" might time out if run on switches with slow flash storage. While the command might timeout, the actual effect of triggering a state sync with the DMF controller still takes effect and is not aborted

Workaround: None (942183)

Enterprise routing

- When IPsec is configured on DPS path groups, applying the same configuration using "configure replace" CLI can cause DPS paths to flap. (814121)
- A DPS path can become inactive and not come back up when any of the encapsulation parameters change.
The workaround is to manually shut/no shut the interface used by the DPS path. (910694)

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)

- The OpenFlow agent fails to return counters associated with a flow, when queried by the OpenFlow controller using protocol version 1.0 (132812)
- On switches using the tg3 driver for the management interface a kernel panic can happen resulting in a switch reload. This is due to a Single Event Upset (SEU) affecting the management interface ethernet controller, and is a rare occurrence. (136944)

Series Affected: 7368 Series

SKUs Affected: DCS-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-7050SX3-48YC12-F, 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-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)
- 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 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)
- In rare conditions, the kernel may panic with the message containing "BUG at ../mm/slub.c:3334" (476225)
- On system reload, the kernel may crash with the "PANIC: double fault" message. (493516)
- Modular switch may restart in case of Fabric or Linecard failure. (554503)
- Local (admin) user is allowed to login with empty password when RADIUS over TLS is enabled (569426)
- "reload peer" configured with a delay such as "reload peer in 1 force" results in both supervisors being reloaded after the specified delay. (621572)
- The OpenConfig & Octa agents can restart unexpectedly if a subscriber stops consuming data (or slows consuming data) for a long period of time. (647535)
- The Octa or OpenConfig agent may restart if there are a large number of sample subscriptions, or if sample subscriptions are made to a large number of paths. (648223)
- ConnectivityMonitor may restart unexpectedly if more than 250 http probes are configured. (672269)
- If number of probes exceed 100 with a minimum of 5 second interval configured, the ConnectivityMonitor agent can restart unexpectedly, causing all of the probes to go down momentarily. (685392)
- The default python interpreter in EOS is now python3 and the python2 interpreter has been removed beginning with this release. (703281)
- BOM change on PCA-01601-03 and change of riser pol image from SFT-01392-01 to SFT-01392-02 resolves issue. (732299)
SKUs Affected: DCS-7280DR3A-54-F, DCS-7280DR3AK-54 and DCS-7280DR3AM-54-F

- EOS may not send traffic out of management interface after the ethernet cable is removed and inserted back.

Workaround is to issue "shutdown" followed by "no shutdown" of the management interface from the CLI. (736552)

SKUs Affected: DCS-7130LBR-48S6QD, DCS-7280DR3A-54, DCS-7280DR3AK-54 and DCS-7280DR3AM-54

- When an event handler has the following 3 conditions:
 - The "trigger on-config disabled" command is configured
 - A delay greater than 0 seconds
 - The first event is triggered sooner than the delay time since the handler is configured

In this case, the handler's action is triggered after the delay seconds since the time the handler is configured, not since the actual event is detected.

Workaround:

Configure the delay to be 0 seconds and add a "sleep" command to the action bash command (778181)

- On reload, it is possible for the VXLAN datapath learning setting to become out of sync with its state prior to the reload. This could lead to hosts and VTEPs not being learned from dataplane traffic when they're expected to be, or being learned from dataplane traffic when they're not expected to be.

As a workaround, flapping the VTI should re-synchronize the setting and fix the issue (789611)

- Removing an interface and reassigning its IP addresses to another interface is not possible via OpenConfig in a single transaction. A workaround is to split the transaction into two or use a mixed-mode (CLI + YANG tree) OpenConfig transaction (791299)
- In an EVPN VXLAN multihoming setup, administratively shutting down and re-enabling the VXLAN source loopback interface on a remote VTEP can result in MAC-IP routes of some hosts, received from the multihoming VTEPs, not being installed causing the traffic to fail to those hosts, in the absence of other covering routes. (804491)
- An EOS agent may restart unexpectedly with a log message such as:
 %FWK-3-SOCKET_CLOSE_LOCAL: Closing connection to Sysdb (pid:3033) at tbl://sysdb/+n (Received an unexpected Tac::Exception)

Additionally, the agent log should contain a message such as:

Exception thrown is Tac::AttrLogEntityException("not found eid 0xd675, sAdapt: NULL")

There is no workaround but the agent is expected to recover normally after a restart in most cases. (807110)

- Units may be unable to complete exiting from maintenance mode if entering maintenance mode is cancelled before reaching Under Maintenance state. (813375)
- Using an event-handler config with multi-line action will fail in running the action on the newly active supervisor in an SSO setup when the switchover happen.

Workaround: Try to convert the the multi-line action to a single line if possible (821674)

- The KernelNetworkInfo IP address shared memory collection can grow over time if interfaces and IP addresses are repeatedly removed and added back again. If this causes an issue, the KernelNetworkInfo agent can be restarted and the IP address collection will return to its original size. (824906)
- Multiple restarts of platform forwarding agents, e.g. Strata-FixedSystem, Strata-Linecard, SandFap-FixedSystem, SandFapNi-Linecard, etc., can result in an out of memory condition.

To recover from this situation, power cycle the switch. (840871)

- OpenConfig/Octa agent may restart while a large number of routes are added using EOS SDK (843671)
- If a PTP boundary clock receives its own messages on a different port, it will fail to sync. (852891)
- When the OpenConfig, Octa, or TerminAttr agents are configured, repeated creation and deletion of VRFs may cause other agents to unexpectedly restart with "
"Marco::InvalidInitializeException("InotifyMgr::addWatch()
inotify_add_watch()" (884562)
- Switch may reboot if the ethernet cable is inserted and then removed from the management interface port during transmission. (886530)
SKUs Affected: DCS-7130LBR-48S6QD, DCS-7280DR3A-54, DCS-7280DR3AK-54 and DCS-7280DR3AM-54-F
- A DirectFlow flow configured with a set nexthop action may fail to be updated if the resolution for the nexthop changes and therefore misforward traffic. Deleting and readding the flow will clear the problem. (897660)
- Under heavy load, ConfigAgent may restart (910907)

- Configuring an untagged L3 subinterface, I', under some parent interface I, in a non-default VRF will prevent sampling of packets received on any subinterface under the given parent interface, I, for features like sFlow and IPFIX. (912274)
- Event handlers with the on-logging trigger might in rare condition rescan the Syslog message file from the beginning, causing messages that fit the trigger regular expression to trigger the handler a second time. (948963)
- When the next hop of a route is updated, neighbors associated with the updated route are removed from the ARP table. As a result, there may be temporary traffic loss while the neighbors are re-resolved. (973419)
- A malformed update message with the Partial Bit set for a non-optional or non-transitive path attribute will result in a BGP session reset rather than treat-as-withdraw. (975141) (1)
- In case of a GM switchover when the GMs' times are significantly different, downstream S-clock, connected to the M-clock experiencing the said switchover, might enter a brief period (~120 Sync message intervals) of incorrect elevated OFM/skew values. The workaround is to ensure that GMs are reliable sources of time. (983258)
- Multiple "dhcp server ipv4" or "dhcp server ipv6" configuration commands might be shown under an interface config if interface VRF change is performed via CVP.

To avoid this issue, change the VRF and add the server configuration in separate CVP transactions.

To work around when the switch is already in this state:

1) via CVP/capiCli, disable CLI caching by issuing `agent ConfigAgent environment DISABLE\_CLI\_CACHE=1` in configure mode, restart ConfigAgent and then reconfigure the interface to the correct vrf.

2) via EOS CLI, reconfigure interface to the old VRF(s) and then perform "no dhcp server ipv4". (986518)

- Python3.9 now includes improved cookie header parsing to prevent DoS attacks caused by malicious cookie payloads. (994288)
- When a SSO switchover happens, the Sysdb agent can restart unexpectedly with "write to read-only proxy" error, causing the system to reset. (996074)
- When the switch is powered on or restarted, Sysdb may close the connection to another agent with "invalid mount profile" error. Sysdb could also close the connection to an agent when it restarts, with the same error. (996465)

- OpenConfig agent might not successfully boot up in "arista-segmentation" exclusive module group. A workaround could be to restart the agent. (1000288)
- Python3.9 now includes improved tarfile parsing to mitigate issues caused by malformed tarfile headers. (1000431)
- The L3 MTU setting includes a safeguard that applies only to interactive SSH sessions (CLI); however, this safeguard doesn't take effect when configured via the non-interactive SSH sessions or eAPI and is set on a per-interface basis (1015561)
- On affected platforms running Arista EOS with OpenConfig configured, a gNOI request can be run when it should have been rejected. This can result in users retrieving data that should not have been available. Further information can be found at <https://www.arista.com/en/support/advisories-notices/security-advisory/21098-security-advisory-0111> (1015821)
- On affected platforms running Arista EOS with OpenConfig configured, a gNOI request can be run when it should have been rejected. This can result in users retrieving data that should not have been available. Further information can be found at <https://www.arista.com/en/support/advisories-notices/security-advisory/21098-security-advisory-0111> (1015822)
- When first booting up or after a SSO, one or more linecards may remain in the "Initialized" state in the "show ptp hardware-sync" command output, causing the device to be unable to participate in PTP.

To workaround this issue, power cycle the affected module using the "no power enable module <MODULE>" then "power enable module <MODULE>" commands. (1022284)

Series Affected: CCS-750 Series

- Update python3 interpreter to improve thread safety (1022998)
- When the system is initialized or a certain subset of EOS agents restart, routing may be disabled in the default VRF. (1029984)
- OpenConfig/Octa agent memory can grow with each Linecard power reset. A workaround is to restart the agent. (1031311)
- When an interface to a peer requires a non-default speed group configuration to be up, ZeroTouch may sometimes fail to connect to bootstrap servers.

Workaround is to either use another interface on an unaffected product, or use an interface with a speed that is supported by the default speed group. (1031963)

SKUs Affected: 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC,

7800R3A-36DM2-LC, 7800R3A-36P-LC, 7800R3A-36PM-LC, 7800R3AK-36DM-LC, 7800R3AK-36DM2-LC and 7800R3AK-36PM-LC

- ACL Agent can terminate on network namespace change on Layer3 interfaces (1038673)
- EosSdkRpc agent may restart when the sum of replicate_count values of next hop group next hops is greater than equal to 65536 (1039484)
- MLAG configuration in conjunction VXLAN flood group configuration and IPv6 underlay may result in DAD seen on hosts.

Workaround: There is no workaround without removing VXLAN flood group configuration, IPv6 underlay or MLAG. (1039765)

- If an EOS event-handler has these 3 conditions:
 - The action is triggered on config
 - The handler has a delay
 - The first event is detected within the delay

The action trigger of the first event will happen at the config time + the delay, not after the event detection + the delay

Workaround: If possible, use the "trigger on-config disable" config command (1040000)

- OpenConfig/Octa agent may restart when a large set of prefix lists, sourced from a file, are configured on the switch using CLI. (1049101)
- When the switch is configured as a PTP Boundary Clock, receiving a PTP message with a negative correctionField value will cause the time sent downstream to be incorrectly adjusted. (1057320)
- Performing a config replace may cause routing to be disabled in some VRFs. A workaround is to perform another config replace operation. If that does not resolve the issue, then the switch needs to be rebooted. (1117760)
- Supervisor peers may have syslog messages not replicated in peer's logs if configured both to not forward messages in RFC5424 format and to use the "traditional" syslog timestamp format with timezone included.

A workaround: Enable forwarding syslog messages in RFC5424 format with "logging format rfc5424" configuration from CLI. (1122129)

- Enabling secure boot on certain versions of About will prevent EOS from booting after May 21, 2025. See <https://www.arista.com/en/support/advisories-notices/field-notice/21317-field-notice-0099> for details. (1133280)

IPSEC

- Under some circumstances IPsec rekey time may show up as 'N/A'. (561575)
- In Jericho2 IPsec tunnel mode, both peers should either enable or disable replay protection. Asymmetric configuration will lead to complete packet loss over the tunnel. (605866)
- On a topology configured with IPsec AutoVPN, doing a shut/no shut several times on interfaces can cause the IPsec agent to restart unexpectedly. (634820)
- Few packet drops are observed intermittently during IPsec Rekey (847831)
- After performing a configuration replace, TunnelIntf agent may restart with a large scale of tunnels if the hostname is used as the tunnel destination. To resolve the agent restart issue for a large scale of tunnels, the IP address can be used as the tunnel destination. (868618)
 SKUs Affected: DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R
- When IPsec packets are received with flow entropy enabled(UDP source port changed based on inner flow), some packets may get dropped because of anti replay errors. The workaround is to configure the replay window to 65k or disable the replay window.

```
To disable anti replay window
ip security
    sa policy <sa-policy-name>
        no anti-replay detection
```

```
To increase the replay window to 65k
ip security
    sa policy <sa-policy-name>
        anti-replay window 65535 (1101823)
Series Affected: CloudEOS VM Series
```

Layer 1

- The forwarding agent may unexpectedly restart after a sequence of speed changes involving speeds of 10G, 25G, 40G, 50G and 100G. (358317)
 Series Affected: DCS-7060X, DCS-7060X2, DCS-7260CX, DCS-7260QX, DCS-7260X and DCS-7260X3 Series

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

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

- Sequences of speed changes involving speeds of 10G, 25G, 40G, 50G and 100G may result in a Strata forwarding agent restart. (456835)

Series Affected: DCS-7260CX3 Series

- If EthernetX/1 and EthernetX/3 are configured for two lane auto-negotiation on a QSFP-DD/OSFP port EthernetX, EthernetX/1 linking down can cause EthernetX/3 to link down. The above is true for EthernetX/5 and EthernetX/7 as well.

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

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

SKUs Affected: 7500R3-24D-LC, 7500R3-24P-LC, 7500R3-48Y4D-LC, 7500R3K-48Y4D-LC, 7800R3-36P-LC, 7800R3-48CQ2-LC, 7800R3-48CQMS-LC, 7800R3K-72Y-LC, DCS-7280CR3K-32P4-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F, DCS-7280PR3K-24-F, DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R, DCS-7280SR3K-48YC8-F and DCS-7280SR3K-48YC8-R

- 10GBASE-T SFP adapters will always advertise 100M/1G/10G link modes during BASE-T auto-negotiation, regardless of the CLI configuration. With the default speed configuration on an interface, they might fail to link up or might report linkdown while the peer has a linkup.

Workaround is to configure both peers with the desired linkup speed explicitly. For Arista switches, `speed auto SPEED` CLI command can be used. (675238)

Series Affected: CCS-710P, CCS-720DP and CCS-720DT Series

SKUs Affected: CCS-720DF-48Y and CCS-720DP-48S

- BASE-T ports might link up but blackhole traffic when they are configured to not use BASE-T auto-negotiation and force speed to 100M using `speed 100full` CLI command.

Workaround is to run `shutdown`, wait for a minute, and then run `no shutdown` on the affected interfaces. (681902)

Series Affected: CCS-710P and DCS-7010TX Series

SKUs Affected: CCS-720DP-24S, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-48S, CCS-720DP-48S-2F, CCS-720DP-48S-F and CCS-722XPM-48Y4

- The /1 interface of a QSFP28 port may get stuck in multi-lane mode when changing speed from `speed 100g` to `speed 10g` after inserting a 40G transceiver. Some of the other interfaces for the port may remain inactive rather than coming up as expected.

Workaround is to configure the /1 interface to 40g and then 10g again, or configure 10g before inserting the transceiver. (694587)

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

- Interfaces with SFP transceivers and link-debounce configuration may experience a link down debounce period double than the configured period. The workaround is to configure "phy link detection aggressive" on affected interfaces (this may lead to an interface flap at the time of applying this configuration). (702920)
- Inserting a 1Gbps SFP transceiver or 10Gbps BASE-T SFP transceiver configured at 1G into a QSFP port using a QSFP-SFP Adapter (QSA) causes 1 logical port to be allocated. If the interface is configured with "speed forced 1000full" or "speed auto 1000full", transceiver removal causes 3 additional logical ports to be allocated. These logical ports may be reallocated to other active interfaces and cause them to become inactive.

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

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

- Inserting a 40Gbps transceiver into a port with a default speed configuration may cause the speed to be auto-adjusted to 4x10G. This may lead to reallocation of 3 logical ports and inactivation of other interfaces.

The workaround is to explicitly configure "speed forced 40gfull" on affected interfaces prior to inserting the transceiver, or configure "transceiver qsfp default-mode 4x10G" globally. (706067)

Series Affected: DCS-7050SX3, DCS-7050TX3 and DCS-7260X3 Series

- When inserting a linecard into a running system there is a chance that the forwarding agent on the switchcard may unexpectedly restart. The same issue may occur on the linecards' forwarding agent if and when a second switchcard is inserted. After restarting all traffic should forward normally.

During cold boot the forwarding agent on one or more of the cards may also restart, but will stabilize before starting to pass traffic. (747464)

SKUs Affected: CCS-755-CH and CCS-758-CH

- L1 Profiles configuration is not supported with OpenConfig. (789141)

- Executing "shutdown" interface configuration command after SSU may result in an unexpected forwarding agent restart and a brief traffic outage on all ports. (792791)

SKUs Affected: CCS-720DP-48S, CCS-722XPM-48Y4 and DCS-7010TX-48

- ZTP may not complete if all of the ports within a speed-group on the system are populated with transceivers that do not support the default speed-group configuration. If these interfaces are the interfaces needed to connect to the DHCP/configuration server this will result in ZTP failing to complete.

Possible workarounds are using a different interface for DHCP/server connectivity, temporarily removing one of the unneeded transceivers from the group until ZTP succeeds, or changing the wiring on the system so that the speed-group has at least one transceiver that supports the default configuration. (797607)

- In some very rare cases the XcvrAgent will fail to program BASE-T transceivers, leading to a failure to bring the link up.

The workaround is to shut/no shut the interface. (798262)

SKUs Affected: DCS-7130-16G3-F (HwRev:C1), DCS-7130-16G3-R (HwRev:C1), DCS-7130-16G3S-F (HwRev:D1), DCS-7130-16G3S-R (HwRev:D1), DCS-7130-48E-F (HwRev:A4), DCS-7130-48E-R (HwRev:A4), DCS-7130-48EH-F (HwRev:A4), DCS-7130-48EH-R (HwRev:A4), DCS-7130-48EHS-F (HwRev:B1), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48G3-F (HwRev:B1), DCS-7130-48G3-F (HwRev:B2), DCS-7130-48G3-R (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48G3S-F (HwRev:C1), DCS-7130-48G3S-R (HwRev:C1), DCS-7130-48L-F (HwRev:A2), DCS-7130-48L-F (HwRev:A3), DCS-7130-48L-R (HwRev:A2), DCS-7130-48L-R (HwRev:A3), DCS-7130-48LA-F (HwRev:A2), DCS-7130-48LA-F (HwRev:A3), DCS-7130-48LA-R (HwRev:A2), DCS-7130-48LA-R (HwRev:A3), DCS-7130-48LAS-F (HwRev:B2), DCS-7130-48LAS-F (HwRev:B3), DCS-7130-48LAS-R (HwRev:B2), DCS-7130-48LAS-R (HwRev:B3), DCS-7130-48LB-F (HwRev:A2), DCS-7130-48LB-F (HwRev:A3), DCS-7130-48LB-R (HwRev:A2), DCS-7130-48LB-R (HwRev:A3), DCS-7130-48LBA-F (HwRev:A2), DCS-7130-48LBA-F (HwRev:A3), DCS-7130-48LBA-R (HwRev:A2), DCS-7130-48LBA-R (HwRev:A3), DCS-7130-48LBAS-F (HwRev:B2), DCS-7130-48LBAS-F (HwRev:B3), DCS-7130-48LBAS-R (HwRev:B2), DCS-7130-48LBAS-R (HwRev:B3), DCS-7130-48LBS-F (HwRev:B2), DCS-7130-48LBS-F (HwRev:B3), DCS-7130-48LBS-R (HwRev:B2), DCS-7130-48LBS-R (HwRev:B3), DCS-7130-48LS-F (HwRev:B2), DCS-7130-48LS-F (HwRev:B3), DCS-7130-48LS-R (HwRev:B2), DCS-7130-48LS-R (HwRev:B3), DCS-7130-96-F (HwRev:A7), DCS-7130-96-R (HwRev:A7), DCS-7130-96-R (HwRev:A8), DCS-7130-96L-F (HwRev:A2), DCS-7130-96L-F (HwRev:A3), DCS-7130-96L-R (HwRev:A2), DCS-7130-96L-R (HwRev:A3), DCS-7130-96LA-F (HwRev:A2), DCS-7130-96LA-F (HwRev:A3), DCS-7130-96LA-R (HwRev:A2), DCS-7130-96LA-R (HwRev:A3), DCS-7130-96LAS-F (HwRev:B1), DCS-7130-96LAS-F (HwRev:B2), DCS-7130-96LAS-R (HwRev:B1), DCS-7130-96LAS-R (HwRev:B2), DCS-7130-96LB-F (HwRev:A2), DCS-7130-96LB-F (HwRev:A3), DCS-7130-96LB-R (HwRev:A2), DCS-7130-96LB-R (HwRev:A3), DCS-7130-96LBA-F (HwRev:A2),

DCS-7130-96LBA-F (HwRev:A3), DCS-7130-96LBA-R (HwRev:A2), DCS-7130-96LBA-R (HwRev:A3), DCS-7130-96LBAS-F (HwRev:B1), DCS-7130-96LBAS-R (HwRev:B1), DCS-7130-96LBS-F (HwRev:B1), DCS-7130-96LBS-F (HwRev:B2), DCS-7130-96LBS-R (HwRev:B1), DCS-7130-96LBS-R (HwRev:B2), DCS-7130-96LS-F (HwRev:B1), DCS-7130-96LS-F (HwRev:B2), DCS-7130-96LS-R (HwRev:B1), DCS-7130-96LS-R (HwRev:B2), DCS-7130-96S-F (HwRev:B1) and DCS-7130-96S-R (HwRev:B1)

- When changing the speed of an interface managed by a CRT55321 PHY (et1-2,31-34,63-64), the Babbage agent may restart once causing all of these interfaces to flap. (837387)
SKUs Affected: DCS-7060DX5-64S
- Using the rc.eos script to add port numbering to the boot config will result in the port numbering not being applied.

Workaround: Please add port numbering through the CLI using 'boot port numbering octal interfaces INTERFACE_COUNT' commands and then reboot the switch. (843173)

SKUs Affected: 7800R3-36D-LC, 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3A-36P-LC, 7800R3A-36PM-LC, 7800R3AK-36DM-LC, 7800R3AK-36DM2-LC, 7800R3AK-36PM-LC and 7800R3K-36DM-LC

- Interfaces with 1GBASE-T SFPs installed may not link up. Workaround is to shut/no shut the port to recover. (847035)
Series Affected: CCS-720DF Series
SKUs Affected: CCS-720DF-48Y, CCS-720DF-48Y-2, CCS-720DF-48Y-2F and CCS-720DF-48Y-F
- Configuration commands which push traffic-loopback configs on inactive interfaces will fail to apply. Applying loopback and speed config in the same config push may result in the loopback config relying on the previously inactive interfaces rather than inactive interfaces from the new speed changes. (898297)
- A link flap on BASE-T ports configured with the default speed ("no speed" or "default speed") may cause an unexpected restart of the forwarding agent. This leads to a flap of all interfaces on the line card and a brief traffic outage. (899785)
SKUs Affected: CCS-750X-48TP-LC
- MAC faults may fail to propagate to the internal CMS42550 PHY chip. This may lead to a one way link up. (917641)
SKUs Affected: DCS-7170B-64C-F
- CRT55321 interfaces that are not /1 might not be able to get configured at all. (917897)

- Interfaces with optical transceivers may experience link issues including failure to link up, link flaps or FCS/CRC errors. (1017076)

SKUs Affected: DCS-7050SX3-48YC12

- A 400GBASE-ZRP transceiver with a custom Media Interface ID in the first application advertisement may cause XcvrAgent to restart repeatedly. Issuing the interface configuration command "transceiver diag simulate removed" on the affected ports will allow other transceivers to operate correctly. There is no workaround to allow the affected 400GBASE-ZRP transceiver to operate. (1021845)

- MACsec enabled interfaces on the same speed-group may receive only FCS-Errors or not transmit frames after Evora-FixedSystem agent restart or system reload. This condition is resolved by restarting the agent while MACsec enabled interfaces are shutdown.

Since this issue occurs only while traffic is present, the workaround is to do the system reload or the Evora-FixedSystem agent restart in the absence of traffic. (1053494)

SKUs Affected: 7500RM-36CQ-LC, DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R

- Removing a transceiver from a port without a speed configuration that belongs to a port channel may cause all port channel members to temporarily leave the port channel before rejoining.

In order to avoid this issue, configure a speed on the port to the desired speed. (1056277)

- There may be a brief traffic disruption during the first stateful switchover (SSO) after a stateful switchcard switchover (SSS) event. (1058599)

SKUs Affected: CCS-755-CH and CCS-758-CH

- Removing secure VXLAN profile from a remote VTEP can result in traffic loss on all front panel ports, workaround is to shut the affected interfaces and unshut them. (1064076) (1)

SKUs Affected: DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

- Link flaps during SSU reload may cause one way link up.

Workaround is to configure "shutdown" and "no shutdown" on the affected interface that was undergoing SSU reload. (1066200)

SKUs Affected: DCS-7060DX5-64S

- On QSFP-DD/OSFP port EthernetX, if EthernetX/1 and EthernetX/3 are configured for two lane auto-negotiation, EthernetX/1 linking down can cause EthernetX/3 to link down. The above is true for EthernetX/5 and EthernetX/7 as well.

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

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

Another workaround is to save the configuration to startup-config and restart. (1087237)

Series Affected: DCS-7060X4 Series

- SSO may fail to complete successfully on fully loaded systems, resulting in a traffic outage while the system reboots automatically. (1098072)
SKUs Affected: CCS-758-CH
- When the switch is booted into EOS, certain ports may exhibit PCS errors or link flaps. These ports will always stand a chance exhibiting symptoms. (1149629)
SKUs Affected: DCS-7130LBR-48S6QD, DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R

(1) This item is in two or more sections on this document

Layer 2

- When running rapid-PVST at high scales of interfaces and vlans, the device may occasionally not transmit a BPDU, which can lead to a BPDU timeout on the neighbor. (442663)
- Performing an SSO switchover on an MLAG primary peer running Multiple Spanning Tree Protocol may result in STP topology reconverging. (470962)
- Restarting STP agent with a high number of interfaces and vlans in PVST mode may result in STP topology reconverging. (472675)
- On EVPN setup using VXLAN and Multihomed active hosts, when the ownership of the host MAC address changes to a new peer, the MAC/IP route may go missing since the new peer is unable to install the MAC address as a dynamic learned entry. (598730)
- Dot1x agent may restart unexpectedly when there are a large number of supplicants present on a port with "dot1x aaa unresponsive action traffic allow vlan <>" and "dot1x mac based authentication always" configured. (731691)

- Pseudowire agent may restart unexpectedly when a large number of flexible cross-connect pseudowires are configured/unconfigured. (732883)
- The following CLI is used to set the source IP address of all outgoing RADIUS packets to that of a particular interface.

```
"[no] ip radius [ VRF ] source-interface <interface-name>"
```

Currently proactive probing of RADIUS servers does not honour this command and will not use the source IP according to this command, the probes which are sent out will still continue to have the source IP of an interface that the switch chooses based on an internal algorithm. (798136)

- When the VXLAN configuration is removed and re-added or VXLAN tunnel interface is shut/no shut in quick succession while VXLAN traffic is being received and MACs are being learned, then VXLAN traffic could be disrupted, and new VXLAN configurations might not take effect.

Workaround is to restart L2 forwarding agent. (834567)

- After SSU on MLAG secondary, spanning-tree portfast ports might stay in learning state. Spanning tree learning state does not impact the STP restartability.

To recover, shut and no shut down the affected ports. (850584)

- 802.1x authenticator in a switch is sending the RADIUS accounting stop messages with the same identifier value in the RADIUS header message. Due to this a RADIUS proxy server may drop those messages, resulting in the 802.1x authenticator not receiving response for those RADIUS requests. (873700)
- Under some conditions, the Arp agent can restart repeatedly when removing attached routes. (912697) (1)
- In an EVPN VXLAN deployment, if a sub interface with dot1q-tunnel configuration is used for VXLAN forwarding, then the c-tag of the overlay packet will not be preserved. (968892)
- In case RADIUS servers are placed in a VRF which has not been configured on the switch, then we will see the ServerProbe agent restart unexpectedly. The workaround to avoid the unexpected restarts would be to remove the misconfigured RADIUS configuration and configure the VRF correctly and add the RADIUS server configuration back again. (975131)
- In case there are 802.1X MBA supplicants connected to the switch and an SSU is performed, then there is a possibility that these MBA device might face traffic loss due to the AAA server name not being resolved in time. A simple workaround for this is to configure the server IP directly instead of using the hostname. (976313)

- In a VXLAN deployment with a virtual IPV6 VTEP address, a device reload may result in the lose of the virtual VTEP configuration.

A workaround is to restart the VXLAN agent or reapply the "vxlan virtual-vtep local-interface" configuration (981821)

- With MBA always (dot1x mac based authentication always) configured, Dot1x agent may restart unexpectedly when a supplicant authenticated via MBA sends EAPOL-logoff in the middle of its EAPOL authentication. (998808)
- With "dot1x mac based authentication always" and "monitor server radius" feature for 802.1X configured, when AAA is unreachable, the switch may continue to send unicast identity request to an EAPOL supplicant even after if fails authentication due to unreachable AAA servers. This may happen if deadtime is configured for RADIUS servers instead of "monitor server radius" feature.
The switch recovers from this state when the AAA servers are reachable. Another potential workaround is to disable "dot1x mac based authentication always" configuration. (1020318)
- The Mvrp agent may restart when a large number of VLANs are registered on several interfaces at a frequent interval. This can lead to traffic loss. (1034789)
- The Mvrp agent may experience continuous restarts when a large number of VLAN registrations occur. This issue persists even after the VLAN Join PDUs are stopped and can cause traffic loss. (1036730)
- If spanning tree protocol puts the root port into a discarding state while the MLAG peer link flapped, it may not recover that interface on its own.

Shut and no shut the affected port to resolve the issue. (1040001)

- The Mvrp agent may experience continuous restarts when a large number of VLAN registrations are sent on several interfaces. (1045352)
- If AAA unresponsive ACL is configured along with cached authentication feature, the unresponsive ACL may be removed on device restart, leading to clean up of supplicants in that ACL. Moreover, the show commands "show running-config" and "show active" do not display the ACL part of the command. (1058621)
- In MLAG devices in active-partial mode with EVPN peers configured over routes exchanged in IPv4/IPv6 AF, if SSU is initiated on the active MLAG peer, there can be a temporary loss in traffic during SSU because some EVPN peers fail to establish on time

As a workaround, the slow peer time configured with `bgp convergence slow-

peer time` can be adjusted to a high value of 180 seconds to help prevent running into this issue. (1068076)

- In an EVPN VXLAN deployment, configuring a VLAN-VNI mapping without creating corresponding VLANs could prevent some dynamic VLANs from being allocated. (1118105) (1)
- In a VXLAN deploy MACs prevent from being learned with "vxlan learn-restrict" are not relearned when "vxlan learn-restrict" is unconfigured. This can lead to traffic drops and flooding for traffic destined the MACs.

A workaround is to flap the VXLAN interface or the restart the appropriate L2 forwarding agent. (1131264)

(1) This item is in two or more sections on this document

Layer 3

- Removing addressless forwarding configuration disables IPv4 routing in the hardware causing packet drops. This issue can be worked around by disabling and enabling IPv4 routing. (254548)
- 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)
- Auto discovery of multiple IPv6 link-local BGP neighbors over a single interface does not work and will result in only one of the neighbors transitioning to established state. (289875)
- BFD configurations with large numbers of sessions may see BFD session flaps during initial session bringup shortly after boot. (375773)
- When graceful restart is enabled under "router isis" and other protocol routes are redistributed into IS-IS then on ASU or SSO, some traffic that is forwarded using redistributed routes might be lost. (413706)
- Removing SVI configuration at scale via a config-replace operation may cause the BFD agent to unexpectedly restart. (414702)

- 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)
- 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)
- When large number of distribute lists are configured in RIP then it may result in route flaps and the protocol agent CPU utilization can increase significantly. (536500)
- KernelFib agent may restart on interface flap, if there are a large number of route resolving through that interface (539757)
- KernelFib agent may restart, if EVPN config is toggled in quick succession. (557358)
- RSVP tunnels might not get established or may use non-best paths when IS-IS system ID or OSPF router ID of any two neighbouring nodes in the IGP network is changed back-to-back and the best path goes via that node. (572080)
- KernelFib agent may restart when interface flaps lead to a large number of routes (entire routing table) being withdrawn (597557)
- Moving the managementactive(ma0) interface to a VRF and restarting Sysdb would lead to managementactive interface staying down in the default VRF after Sysdb(and other agents) come back up. (613014)
- During Zero Touch Provisioning(ZTP) in an IPv6 environment, even if the Router Advertisement(RA) packets do not have the M and the O bits set, the Arista switch still tries to contact a DHCPv6 server for information other than the IPv6 address. (691387)
- Cspf agent may sometimes restart when there are two or more routers with duplicate TE router ID in the network (700468)
- Nodes marked with the IS-IS overload bit will not be avoided if they are specified in an include-hop constraint (754609)
- Upon the Arp agent starting up in setups with a large number of routes, the Arp agent may restart. (764196)
- After reload, static ARP or ND entries may not get correctly restored.

Workaround is to restart the Arp agent. (790370)

- VRF names with "." result in repeated restart of the Launcher agent while starting a process with a name like "Rib-vrf-foo.bar".

To restore the system, first remove the VRF with "." in its name with "no vrf instance <vrf.name>". Then, delete any file in /etc/ProcMgr.d/run/ ownedByLauncher/ of the form "<agent>-vrf.<vrf.name>". For an example VRF named "foo.bar", the file to delete might be "Rib-vrf.foo.bar" or, when using the multi-agent routing protocol model, a file for a protocol-specific agent such as "Ospf-vrf.foo.bar" or "Isis-vrf.foo.bar". (792987)

- Router Solicitation (RS) packets with the source address that is not the unspecified address and no source link-layer address option, causes a Router Advertisement (RA) to be sent with the destination address as unicast address and the destination link-layer address as all-nodes multicast address. (802490)
- RSVP tunnels can fail to come up if an include-hop is specified on a path that causes the shortest path to the destination to contain a loop (827083)
- If local adjacency segment protection is enabled and IS-IS agent restarts at scale, MPLS agent may restart sometimes. (893203)
- After the Cspf agent restarts, tunnels using IS-IS flexible algorithm can experience momentary traffic loss. (894389)
- When a covering route is unconfigured and reconfigured or the interface the route was configured on is flapped, the ARP or ND entries that are covered by the route may not be installed.

Workaround is to remove and readd the route, flap the interface, or restart the Arp agent (895053)

- Under some conditions, the Arp agent can restart repeatedly when removing attached routes. (912697) (1)
- When EOS receives two Router Capability TLV in IS-IS LSP, it only processes the first one. This can result in EOS not to function correctly viz if the Flex Algo TLV is received in the first Router Capability TLV and the Segment Routing Capability TLV is received in the second Router Capability TLV, EOS will ignore the second Router Capability TLV thereby assuming that the node that originated the LSP does not support Segment Routing. (918334)
- When a static MPLS route's next-hop becomes unviable, EOS fails to program the MPLS route with a next-hop-group adjacency in hardware, even when the MPLS static route with next-hop-group is programmed with a lower metric. (923930)
- The OSPF agent may restart when there are duplicate addresses on the same broadcast network. (930800)

- IS-IS agent will restart if it receives prefix segments with Algo-ID 1 (941111)
- An RSVP authentication sequence number rollover in EOS occurred for RSVP agents which restarted on or after June 10, 2024. This can lead to session timeouts if in a pair of upstream/downstream neighbors, both neighbors established the RSVP session before June 10 and one of them restarted RSVP after June 10.

A further symptom of this issue occurring is an incrementing "Invalid sequence number" RSVP rx error counter.

A workaround is to turn RSVP authentication off and back on again. A change of secrets afterwards is recommended as well to prevent replay attacks. These steps can also be applied proactively. (964630)

- If there is a high churn in IPv4/IPv6 routes in the FIB while the L3 forwarding agent is down, the L3 forwarding agent might restart unexpectedly when the agent comes back up. (985291)
- In configurations which use hierarchical FECs (eg. VXLAN, MPLS tunnels), port or BFD session failures may result in longer than expected traffic loss. The issue affects routes resolved via multiple paths (eg. multiple tunnels) or have backup paths, where the underlay paths for these are non-ECMP. (1008054)
- There may be traffic loss when performing a SSO with OSPFv2 configured. (1009676)
- There may be traffic loss due to OSPFv3 routes being modified during SSO (1015073)
- OSPF agent can unexpectedly restart on config replace with an interface configured with OSPF area but without an IP address.
A workaround is to configure an IP address on such an interface. (1016903)
- In a scenario where multiple RIP networks are configured such that those IP addresses are configured on the same interface, if one RIP network is unconfigured then RIP gets disabled on the interface and all the adjacencies on that interface will go down.

Workaround is to do shut/no-shut on the interface. (1039428)

- IS-IS authentication will stop working after reload of the router if "isis enable" command is part of "interface profile" configuration and "isis authentication" command is part of interface configuration

Workaround is to configure "isis enable" command directly under the interface instead of "interface profile" (1039795)

- BGP agent may consume excessive memory when there is a significantly high volume of peer flaps. This may result in restarting of BGP agent. (1042063)
- Deletion and addition of nexthop groups in quick sequence might cause the L3 forwarding agent to restart unexpectedly. (1052200)
- If a static link local neighbor is removed from the kernel, the neighbor may not correctly be added back to the kernel.

Workaround is to remove and add back the static link local neighbor configuration or restart the Arp agent (1060365)

- With TI-LFA FRR enabled, LFIB routes on a PLR router may not have TI-LFA backup paths programmed for local or remote adjacency segments which are reachable via ECMP and when these ECMP paths share the same failure risk, like same SRLG or same nexthop router. (1086677)
- In cases where DHCP smart relay is enabled, when the DHCP Relay agent receives an IPv4 relay renew request that needs to cross VRFs, it will always insert the primary IP address in the option-82, even if the requested IP is in the secondary subnet. (1099157)
- If a route is deleted or changes next hops, ARP or ND entries matching that route that are removed immediately after may not be correctly cleaned up. If this is caused by an SVI changing VRFs, new ARP or ND entries in the new VRF on the same SVI matching the IP of a previously removed entry that was not correctly cleaned up may not be installed in hardware in the new VRF

Workaround is to restart the Arp agent. (1101903)

- BGP routes which are imported from BGP VPN or BGP EVPN (Type 5) address family and further leaked from the imported VRF to another source VRF using route-leak routing policy may not be correctly programmed in the kernel software forwarding table resulting in either misforwarding or blackholing of SW forwarded packets in the new source VRF. (1104410)

Multi-agent

- Performing a stateful switchover in a switch with a large number of VRFs configured, with BGP config in all VRFs and IGP configured in a large number of VRFs, may cause the system to restart unexpectedly, leading to traffic loss. (376773)
- The BGP agent may restart when EVPN imported routes churn or "clear ip bgp *" is issued when the BGP agent is under load. (456044)
- SSO on ASBR with MPLS L3 VPN Inter-AS Option B configured may result in temporary traffic loss due to some VPN labels being reallocated (498848)

- SSO with MPLS L3 VPN configured may result in temporary loss of traffic flowing from MPLS core towards the tenant sites (507353)
- Under rare circumstances, Bgp Agent may restart when VRF configuration is removed (845945)
- Setting community no-advertise, no-export, and local-as inline in route-map results in integer values in the YANG path /routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/set-community/inline instead of NO_ADVERTISE, NO_EXPORT, and NO_EXPORT_SUBCONFED (858821)
- The OSPF agent may restart when an extended community list with many entries is configured. (917762)
- Running "show mpls bgp labeled-unicast bindings detail" may cause ConfigAgent to restart unexpectedly if there are too many bindings. (918712)
- The BGP inheritance model requires BGP to be configured in default VRF before it can be configured in non-default VRFs.
The workaround is to enable BGP in the default network-instance (AS number at minimum) before enabling BGP in non-default network-instance. (927063)
- If using BGP MPLS VPN and the neighbor received attributes discard command is enabled to discard attribute 16 (BGP Extended Community), then all previously received and imported VPN routes should be removed from any VRF where they were imported, since the Route Target Extended Community is now removed. However, on scaled setups, it is possible that some routes may remain imported in some VRFs. (937337)
- Next hops may become unreachable when quickly deleting and re-creating VRFs with the same name.

In order to workaround this issue, delete the VRF, and wait some time before creating with the same name. (960502)

- In an EVPN type 2 deployment, if there is high MAC VRF and peer scale, and verbose tracing is enabled for the EVPN plugin (for example with 'trace Bgp setting BgpEthernetSegmentSm/*') then it is possible for the Bgp agent to crash on startup or config replace.

Workaround is to reduce the trace verbosity. (962403)

- If BGP "update wait-install" is configured along with route map "match aggregate-role contributor", BGP will not advertise aggregate contributors out to neighbors.

As a workaround, replace "match aggregate-role contributor" with a prefix list which matches on just the contributors (when possible) (971101)

- A malformed update message with the Partial Bit set for a non-optional or non-transitive path attribute will result in a BGP session reset rather than treat-as-withdraw. (975141) (1)
- A malformed AS4_PATH path attribute will result in a session reset rather than attribute discard. (975145)
- If an ARP or ND entry is added while the covering route for this entry is being deleted and added back, the ARP or ND entry may never be correctly added to the system.

Workaround is to restart the Arp agent. (975220)

- An update message with malformed ATTR_SET attribute may result in a BGP session reset rather than treat-as-withdraw. (985850)
- In IS-IS, configuring proxy node segments with a large range might result in erroneous prefix to proxy-node mappings. (991006)
- ARP entry might get prematurely removed for a host if a static interface route is configured for it and the interface has "ip attached-host route export" configured, and the ARP is cleared.
Workaround: de-configure, then re-configure the static route. (1002200)

- Performance issue when BGP Link-State is in use. <https://www.arista.com/en/support/advisories-notices/security-advisory/21092-security-advisory-0110> (1006114)
- While a BGP peer is being moved into an update group following an outbound policy change, if a BGP route is deleted the switch may miss sending withdrawals for the route to other peers in the same update group. (1013516)
- Arp agent may restart unexpectedly if the router has BGP/static ipv4 next hops recursive resolving through other ipv4 routes with ipv6 egress next hops. (1019807)
- If a community list is deleted in a config session, or if the system is under load when the list is deleted, it is possible for the community list deletion to end up in an incomplete state. The knock on effect of this is that any route map which referenes this incomplete community list will not react to configuration changes. The route map configuration changes will not be reflected in running policy.

Workaround is to remove any references to the deleted community list from any running route map. (1019978)

- All IPv6 link-local BGP peering sessions, configured using 'neighbor interface ...' command, flap during SSO or SSU or when 'Ipv6RouterAdvt' agent restarts. This may result in traffic drop. (1021849)

- When BGP receives routes with nexthop points to its own peering address, those routes are discarded, because such nexthop is invalid per RFC4271 Section 5.1.3. However a new Via Id is consumed for the invalid nexthop. Via Id may wrap around if BGP receives such UPDATE messages at a very high rate for a prolonged time. Existing Via may be overwritten after Via Id wraps around. (1029184)
- BGP may not advertise paths with a new local next hop to a peer if the peer is moving between update groups. (1034626)
- The default route may not be advertised to a peer with "default-originate always" configured if that peer shares a "Policy group ID" with another peer that is in a different update group. (1035861)
- BGP may transiently advertise unexpected paths to a neighbor if that neighbor joins an update group that is reevaluating paths.

For example when renaming multiple RCF functions applied to neighbors the new RCF functions names will be inactive temporarily. These neighbors will now share an update group as they all have inactive RCF functions. The neighbors new to the update group will start advertising paths based on the previous update group state until the update group finishes reevaluating paths based on the inactive RCF function and withdraws them.

A workaround is to separate policy configuration changes and BGP configuration changes. For example add duplicate RCF functions with the new names and wait until they are active, then change BGP to reference the new RCF functions, and finally remove the RCF functions with the old names. (1060542)

- A switch configured with BGP RT membership may fail to advertise membership routes for some of the configured import route-targets. As a workaround, modifying any route-target import command in the configuration should unblock the missing route-targets. (1060979)
- In rare situations, when addresses in the 127.0.0.0/8 reserved block are transmitted by protocols or used in certain configurations, the Ira agent may unexpectedly restart until the offending address is removed. Consider using addresses from private-use address blocks like 10.0.0.0/8 (see <https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml>) instead. (1064362)
- When BGP peer flap damping is enabled, and one of the peer restarts after the peering was established, the remote peer may be unable to re-establish peering. If a peer is stuck in such state, configuring either the Idle-restart timer or manually clearing the peering with 'clear bgp <peer_ip>' on the remote peer, will recover from this condition. (1076795)

- During SSU, link-local BGP peering sessions - configured using 'neighbor interface ...' command - may not be reestablished before BGP convergence timeout which may result in traffic loss until the router neighbors are learned via router advertisement messages. (1084729)
- In EVPN or MPLS VPN deployments, when reachability between PEs flaps repeatedly, updating the reachability of dependent VPN routes may take significant time to complete and may delay bestpath updates. During this period the Bgp agent will have high CPU consumption. (1097913)
- Modifying a peer filter used by BGP interface peers (configured using 'neighbor interface ...' command) will flap all IPv6 link local peers associated with those interfaces. (1101198)
- If MPLS LDP sync for OSPF is enabled and multiple OSPF areas are configured, OSPF interfaces might not update the cost after the LDP session changes state. (1104259)
- If there is an error with the new BGP neighbor password, the existing password might also be removed. (1104776)
- When BGP aggregates are configured, the system may stop adding or updating aggregate routes in the BGP RIB once BGP convergence has been declared. The problem is more likely to occur with high route scale combined with hierarchical aggregates or aggregates which have a large number of contributing routes. (1143364)

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)

- In an EVPN VXLAN deployment where EVPN Type-2 or Type-5 routes have been received, configuring a static EVPN route that reuses a nexthop VTEP received from EVPN can cause received routes to become unresolved.

A workaround is to unconfigure the static EVPN route. (443964)

- The Rib agent may restart when large number of interfaces go down at the same time and there is a large scale of BGP routes reachable over these interfaces. (549816)
- Next hops for static and BGP routes undergo proactive ARP resolution, and routes for which the next hop is not resolved behave as drop routes. When a subnet (associated with a local IP address) is moved from one interface to another, next hops reachable via that network change to reflect the new interface, but the proactive ARP resolution may not change accordingly, resulting in traffic loss despite the affected routes appearing correctly in "show ip(v6) route" output.

This will principally affect static and iBGP routes because by default eBGP routes' next hop is the address used for peering, so ARP resolution is completed before the route is learned. The straightforward workaround is to restart the Rib agent for the affected VRF. If all next hops affected by an address move can be identified, actively connecting to them from the switch, such as with ping, will also result in ARP resolution and restore forwarding. (804746)

- When OSPF3 is using message authentication, the Ospf3 agent may restart unexpectedly. (816808)
- When OSPF3 is using message authentication, the Ospf3 agent may restart unexpectedly. (882020)

(1) This item is in two or more sections on this document

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)
- IPv6 overlay virtual interface may get disabled with Duplicate Address Detected (DAD) state with VXLAN over MLAG configuration. Flapping the overlay IPv6 L3 interface ('shutdown' followed by a 'no shutdown') will resolve the issue (798299)
- When 'switchport mode trunk phone' is enabled on one peer of an MLAG switch, it must also be manually configured on the other MLAG peer to allow for the switchport configuration to take effect. (877198)
- Changing certain MLAG configuration parameters (MLAG domain ID, local MLAG interface, MLAG peer address, MLAG peer interface) will cause the MLAG session to temporarily go down while the connection is renegotiated. This can lead to transient L2 loops in the network until the MLAG session is re-established. If EVPN is also being used, this can lead to rapid MAC flaps between devices in the network, causing those MACs to become blacklisted by EVPN. To see if a particular MAC has been added to the EVPN blacklist, use the "show bgp evpn host-flap" command.

As a workaround, configuring EVPN host-flap recovery will allow the blacklisted MACs to age out after a configured time once the MLAG session has been re-established. Blacklisted MACs may also be cleared manually with the "clear bgp evpn host-flap" command. (891929) (1)

- An interface which is an MLAG peer-link in trunk mode belonging to the underlay (core VLAN interface) and overlay VLANs may lose ingress membership on the overlay VLAN when removed as a peer-link. (961096)
- On MLAG devices running EVPN OISM, `mlog shut` `no shut` may result in traffic loss.
On devices running EVPN OISM in a non MLAG scenario could result in PIM neighborship being formed across the EVPN fabric, this could prevent traffic from being routed and result in traffic loss. (990767) (1)
- When the MLAG agent restarts and the MLAG domain ID is changed at roughly the same time (within 250ms or so), MLAG may report an Active and Connected state in "show mlag", despite the differing domain IDs between peers. This will break the exchange of all MLAG state between peers.

The workaround is to ensure that the domain IDs match between peers. If the domain IDs match, the exchange of MLAG state will resume. (1126109)

MPLS

- FlexAlgo LFIB routes and tunnel entries may not be computed for FlexAlgo segments sometimes if the corresponding FlexAlgo definitions are configured in router isis mode after the FlexAlgo segments are already advertised by other routers. (820797)
- MLDP mappings may be temporarily withdrawn and then sent again when a FEC transitions from bud role to leaf role. (979220)
- Lfib routes are not created for IS-IS SR prefix segments for prefixes redistributed from other routing protocols with no-php option. This will cause traffic loss as the no-php flag is advertised to peers and the top label is not popped by the penultimate hop router and there is no lfib route programmed at tunnel end point to pop the top label and perform routing (995953)
- MPLS traffic of IS-IS SR prefix-segment can loop if the originator of the prefix-segment advertises the segment with 0 metric and no-php option in one instance and receives the same segment from another node in another instances as a result of the segment getting leaked across instances by some other node in the topology (995977)
- When a physical interface is moved from one IS-IS instance to another all segment routing MPLS traffic via the interface could be dropped (1000327)
- In rare cases, Pseudowire agent keeps continuously restarting after config replace. (1034329)

Multicast

- An S,G route may not properly cleanup previously sent S,G joins. This would prevent S,Gs from being deleted in the upstream router. A work around is to flap the interface in which the join is sent on. (291756)
- If large number of IGMP reports are received in a very short interval, some of them might get dropped resulting in the multicast route not being created. The routes will be re-learned on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIcmp", consider increasing the bandwidth for this Copp class. (356675)
- show ip igmp vrf <name> groups command does not filter the IGMP interfaces as expected. (368071)
- High CPU utilization might be observed for the forwarding agent when PIM Bidir routes are present. (479119)

- Removing router pim bsr also remove router pim sparse mode configuration. Workaround is to just remove specific configuration under router pim bsr mode. (516188)
- In a PIM-SM network when an RP is also the first hop router, it may take upto a minute to advertise source to MSDP if the source goes from inactive state to active state. (584263)
- Configuring "ip igmp host-proxy <ACL-NAME>" where the access list has a large multicast group prefix length may cause the IgmpHostProxy to be repeatedly restart resulting in traffic loss. Use smaller prefix lengths to work around the problem. (586865)
- In certain topologies with EVPN Multicast deployments and Multicast Transit feature enabled, with the RP in external network, learns about the source through PIM SGRPT prunes rather than PIM Register messages, RP may never join the source tree later upon receiving IGMP joins.

Workaround: Always ensure there are local receivers on RP for the same group. (616245)

- Igmp host proxy configuration using overlapping prefixes where a permit is followed by a deny rule, the groups in the permit range are not proxied. As a workaround, remove `ip igmp host proxy access-list <acl>` and use `ip igmp host proxy <group>` to explicitly configure the permitted groups. (619688)
- On devices running EVPN OISM with multihoming, when there are multiple hosts interested in the same group traffic, when one of the receivers sends a leave, then remaining receivers may experience traffic loss. The receivers will recover on the next general query.

Workaround for *,G receivers use v2 version of querier `ip igmp version 2` for the SVI or `ip igmp snooping querier version` L2 querier (636045)

- Running MFA feature on x86_64 can lead to memory contention. (662221)
Series Affected: CCS-710P, CCS-720XP and DCS-7010TX Series
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- MVPN related multicast route might be programmed incorrectly after BGP restart on a scaled MVPN PE (677280)
- On devices running EVPN VXLAN multicast OISM, Overlay multicast traffic with TTL1 arriving on VXLAN enabled VLAN in a VRF with `evpn multicast` enabled is not VXLAN encapsulated when there is a receiver in the same VLAN on remote VTEP.

Workaround

Use `TTL>1` or enable ``redistribute igmp`` under router BGP for the VLAN on the receiver VTEP. (703609)

Series Affected: CCS-720XP, DCS-7050CX3, DCS-7050SX3, DCS-7050TX3 and DCS-7300X3 Series

- If KernelMfib agent is/gets terminated the device will be unable to discover new multicast sources and/or detect traffic duplication. This can result persistent multicast traffic loss the first case and persistent traffic duplication in the latter.

Workaround is to restart Pimsm and PimReg agent by issuing "agent pimsm terminate", "agent pimreg terminate" command. (707691)

- When the VXLAN interface is shutdown and immediately unshut using the "no shutdown" CLI, for some of the multicast tenant VRFs, the SBD IMET routes may be withdrawn and never re-advertised. This might impact any EVPN multicast traffic flowing through that device. A workaround would be to shutdown the VXLAN interface again, wait for all SBD IMET routes to be withdrawn and then unshut the interface. (722075)
- On MLAGed devices running multicast, ``shut`` followed by ``no shut`` under ``mlag`` submode can result in stale multicast routes (ones existing before the shut was excuted) being programmed in the hardware resulting in traffic loss.

Workaround: Toggle ``routing`` under ``router multicast`` for the affected VRF (727482)

- In an MVPN setup, traffic at ingress PE might get dropped if multiple flaps occur on the source interface. (777298)
- On devices running multicast with MRIB lookup enabled(`rpf resolution ribs multicast-rib`), when the MRIB lookup yields a connected route it is not used as the RPF and results in traffic not being forwarded. (783549)
- On devices running multicast with "software forwarding sfe," BessMgr agent may restart unexpectedly when performing configure replace. (828727)
- On devices running multicast with 'software forwarding-sfe', running 'configure replace' may result in persistent multicast traffic loss in the default VRF.

Use 'agent BessMgr terminate' to workaround the issue. (829341)

- On devices running EVPN L2 multicast with a PIM -ASM underlay, persistent multicast traffic duplication may be seen.
To avoid this issue, consider having PIM-SSM in underlay by configuring ``vxlan multicast protocol pim ssm``. (845134)

- On devices running EVPN multicast OISM, when a PIM External Gateway (PEG) device is configured as RP, receivers in the external network that join the EVPN network via the PEG-nonDR RP may experience traffic loss.

``clear pim ipvX vrf <VRF> sparse-mode route <S> <G> `` on the PEG-non-DR will recover the flow. (848542)

- When a router is acting as DR + RP + MSDP peer, the receivers connected to the peer may experience delay in receiving the traffic. (883260)
- On devices running Multicast with Kernel based software forwarding, toggling multicast routing configuration may result in complete and permanent loss of traffic.
Workaround is to use ``software-forwarding sfe``. Note that changing this config will require a reload. (914121)
- On devices running PIM sparse mode, the device that is in the process of switching from RP tree to Source tree may experience a traffic gap if it happens to receive an RPT prune message from downstream.
The traffic will recover once the source tree is completely built. (922806)
- (*,*) SMET and (*,*) SPMSI routes with AddressFamily of IpUnknown is not supported yet. Upon receiving such a SMET or SPMSI route, Bgp agent will restart. The only way to workaround this issue is to send (*,*) for IPV4 and IPV6 as two different Evpn routes. (951938)
- When a corrupted SPMSI route is received, BGP agent will restart. Only way to workaround and recover is to stop sending the corrupted SPMSI. (962570)
- During restart or at sufficient mroute scale, with MVPN tunnel configuration, SandMcast agent might restart multiple times with "FecId range exhausted" message in traces. This could lead to traffic loss/interruption for certain MVPN tunnels. Removing and re-adding affected tunnel configuration can be used as a workaround. (969935)
- On MLAG devices running EVPN OISM, ``mlag shut`` ``no shut`` may result in traffic loss.
On devices running EVPN OISM in a non MLAG scenario could result in PIM neighborship being formed across the EVPN fabric, this could prevent traffic from being routed and result in traffic loss. (990767) (1)
- If a source starts sending traffic and then stops for a while, and then restarts, continual registers sent may result in issues such as burden on CPU/network and DDOS protection getting triggered. (996739)
- On the device running multicast IPv6, when reloading the device, multicast traffic might be lost and won't recover automatically.

Workaround: Restart the SandMcast agent. (1031956) (1)

- In EVPN Multicast network, if a PEG which is non-DR but is the RP could fail to deliver to external receivers in the PIM domain.
The workaround is to add an overlay receiver for any group in the VRF.
(1037883)
- On devices running EVPN OISM on a PEG, change in the assignment of the dynamic VLAN (SBD VLAN) for the VRF might result in traffic loss.
Workaround is to restart Pimsm agent to recover. (1056971)
- IGMP protocol agent may crash during configuration replacement when the number of IGMP interfaces added is large. The agent recovers after the initial crash. (1066748)
- In EVPN Multicast transit network with a RP on a stick topology, some multicast receivers connected to the PIM External Gateway that is on the path to the source from the RP may not get traffic. The workaround is to add a receiver on another PIM External Gateway. (1079151) (1)
- On devices running PIM/IGMP on an interface with "pim ipv4 local-interface" IGMP may not have multicast traffic delivered to this interface.
Workaround is to toggle "pim ipv4 sparse-mode" or "ip igmp" to recover.
(1081685)
- On MLAG devices running EVPN OISM, receivers connected to orphan VLANs may not receive multicast traffic.
To work around issue, configure the VLAN and SVI and enable multicast on the MLAG peer as well. (1089742)
- On devices running EVPN OISM when the dynamic VLAN(SBD VLAN) for the VRF changes, traffic loss might be observed.
Restart BessMgr to recover. (1095059)
- On MLAG devices running EVPN OISM, receivers connected to orphan VLANs may not receive multicast traffic.
To work around issue, configure the VLAN and SVI and enable multicast on the MLAG peer as well. (1104327)
- Doing config replace with EVPN Multicast and PIM BFD configurations present may cause duplicate multicast traffic.
Restart PIM agent to fix the problem. (1105152)
- If the connection between the IgmpSnooping agents running on MLAG peer switches hits an unexpected problem, it may not reconnect successfully. The connection remains in "connecting" state with the IgmpSnooping state out of sync. The workaround is to restart the IgmpSnooping agent. (1172545)

NAT

- On MLAG set-up using NAT synchronization between the two peers, TCP connection might reset while the connection is being established.

A workaround is to make sure that all the control packets for a specific connection arrive on the same peer. Careful network design can be used to achieve this goal. (681428)

Series Affected: DCS-7170 Series

- When "ip nat tcp flow creation trigger syn" is set, if dynamic NAT configuration is modified or if the admin shutdown/no shutdown performed on an interface configured with dynamic NAT while NAT traffic is flowing, NAT connections may be left untranslated.

Workaround is to manually flush all the entries using the command "conntrack -F" in the corresponding network namespace bash mode. If configured on a non-default VRF (vrfName, for example), run the command "ip netns exec ns-vrfName conntrack -F".0 (748487)

- NAT translations may not happen if either the ingress or egress interface is repeatedly flapped. (777877)
Series Affected: DCS-7130 Series
- Hardware translations may not properly programmed if Dynamic NAT profile is changed while high amount of NAT entries are created. (792730)
- NAT agent might restart, if there are large number of address-only NAT flows from the same host. This is caused due to an inefficient loop that updates the last activity time for each of these flows causing NAT agent to fill the attrlog buffer. (798605)
- ACL agent may restart unexpectedly once after doing config-replace with non-NAT to NAT config and vice versa (897496)
- In Active-Active NAT synchronization mode between the two peers, TCP and ICMP connection establishment might fail if the control packets are received in the other peer.
A workaround is to make sure that all the control packets for a specific connection arrive on the same peer. (954033) (1)
- If the device is restarted when overload mode dynamic network address translation is configured and the NAT traffic is running then the NAT translations may brake and some or all of the flows subject to NAT are being forwarded without translation after the device completes the restart.

Workaround is to manually flush all the NAT entries using the command

"conntrack -F" in the corresponding network namespace bash mode. If configured on a non-default VRF (vrfName, for example), run the command "ip netns exec ns-vrfName conntrack -F". (1122152) (1)

(1) This item is in two or more sections on this document

SwitchApp

- Heavy CPU bound traffic may cause BGP flaps despite Copp being configured with ingressCpuDrop counts being observed under "show fpga switch counters". (748835)
- MLAG peer link does not come up post reboot until the non-mlag reload delay timer expires (758022)
- A Security ACL can not be applied to the same VLAN as a NAT ACL (772289)
- If the MakoCentral agent is restarted, the MakoSharedResources can restart unexpectedly, and some configured NAT sessions may not be correctly programmed in hardware. (773579)
- Switch interfaces may discard transmitted traffic for up to 30s after the corresponding front panel port is brought up despite having an UP status. If being UP causes the interface to attract traffic that would otherwise take a different path (e.g. if the port is a routed port it will install its connected prefix and this may attract traffic from neighboring routers) then this may cause traffic loss in the network. (778047)
Series Affected: DCS-7130 Series
- When a member is added to a port-channel, reverse NAT translation may not take effect for the new member.

The workaround is to remove and re-add the relevant NAT configuration. (778996)

Series Affected: DCS-7130 Series

- After MakoL3Unicast and/or MakoSharedResources agent restarts routes for connected hosts may remain programmed in the FPGA (and be visible via the 'show fpga switch ip route') even though the route and ARP entry have been removed (no longer shown in 'show ip route' and 'show arp').

Restart SwitchApp using 'switch default' followed by 'disabled' followed by 'no disabled' to clear this state. (800353)

- ACLs configured with a lt match on ports will not match on packets with port equal to 0. The workaround is to configure an explicit rule to match on port 0 (809597)

- If a SwitchApp device has configuration that matches following condition:
 1. There are trunk interfaces configured on the switch
 2. ACL applied on an SVI that is included in the trunk VLAN range
 3. The ACL involved has counters enabled
 4. There is at least one rule in ACL involved uses L4 port range
 5. There are multiple SVIs with ACL configured and all in trunk VLAN range
 6. The involved SVIs must have members sharing the same pipeline
 7. When ACL applied on the SVI, members of SVI involved needs to be active

When links of the SVI involved are down, MakoSharedResource agent might (not always) crash with the following failure: ``assert(!hasCounter || conterRef)``

There are many workarounds to avoid the issues by breaking any of the conditions above, for example:

1. disable counters in ACL involved
 2. Do not use L4 port range in ACL rules, split it into individual rules with single L4 port filter
 3. Split members of SVI involved into different pipelines
 4. Narrow trunk interface VLAN range using CLI 'switchport trunk allowed vlan xxx'
- ... (869876)
- When switchApp is used and ACL is configured, on MakoSharedResources agent restarting alone (either caused by bug or manually), the internal ACL states is lost, which causes inconsistency between hardware and software.

The workaround is to reboot the device. (869902)

- L3 multicast routes in non-default VRFs may incorrectly age out even when there is still traffic activity.

The workaround is to restart the MakoMcast agent once. Restarting the agent may cause a brief disruption of both L2 and L3 multicast forwarded traffic across the device. (959974)

SKUs Affected: DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R

- Upon many IGMP snooping membership or multicast route churn in a small duration, the VLAN Rewrite table may reach exhaustion resulting in multicast routes not being programmed. MROUTE-3-HW_RESOURCE_FULL logging messages will also be displayed. The membership and route churn causes VLAN Rewrite table resources (VlanGroupId) to be allocated inefficiently, and exhaustion is reached even when there is a small number of access ports that make up outgoing interface lists (OIL) for multicast replication. (993432)

- Two or more security ACLs with different priorities and overlapping ip ranges may not be applied.
A workaround is to not overlap IP ranges. (1000156)
- Multicast FIB routes that are not programmed due to hardware resource overflow are not automatically reprogrammed when resources become available.

One workaround is to manually clear the unprogrammed multicast FIB entries, and have them be re-established for programming. (1015170)

- Flaps of interfaces with multicast NAT configured may result in a brief period of resource exhaustion. If the translations configured is at a supported scale, the system should recover automatically without user intervention. (1017562)
- Upon a large number of multicast route churn in a short duration, the MakoMcast agent may unexpectedly restart. The system is expected to stabilize and recover automatically. (1044933)
- The number of multicast (S, G) routes we are able to install in the hardware may vary between reboots - a set of routes that previous fitted in the table may not fit after a reboot and this may occur at a utilisation that is significantly lower than the expected 90%.

One workaround is to try reloading the device, this may cause the routes to fit. (1070380)

- Upon many IGMP snooping membership or multicast route churn in a small duration, the VLAN Rewrite table may reach exhaustion resulting in multicast routes not being programmed. MROUTE-3-HW_RESOURCE_FULL logging messages will also be displayed. The membership and route churn causes VLAN Rewrite table resources (VlanGroupId) to be allocated inefficiently, and exhaustion is reached even when there is a small number of trunk ports that make up outgoing interface lists (OIL) for multicast replication. (1079108)

vEOS Router / CloudEOS

- When IPsec phase 2 negotiation (mis-match in SA policy parameters such as ESP encryption algorithm or ESP integrity algorithm) fails or IPsec negotiation fails with a wrong shared secret, the IPsec security parameter Index(SPI) allocated for that negotiation is not freed. When this happens continuously, all SPI space can be allocated but not freed. This can result in SPI allocation failures and the IPsec connections can't come up.

The workaround is to restart the IPsec agent. If an IPsec connection continuously fails, then that connection can be either shut down or deleted for prevent further SPI leaks. (1055216)

DCS-7170 Series

- If a LAG member has "no switchport" configured incorrectly, it may sometimes not be programmed correctly in hardware for VLAN membership checks causing traffic loss.
Workaround is to not configure "no switchport" for LAG members. (855423)

DCS-7170 Series

- PTP transparent mode is currently not supported. Attempting to configure PTP transparent mode with the `ptp mode e2etransparent` or `ptp mode p2ptransparent` command will cause the BfnSlice agent to unexpectedly restart. (609143)
Series Affected: DCS-7170 Series
- Poor LAG member hashing can occur on a 7170 when flows are hashed along a path that contains both ECMP and LAG, where the LAG has a number of members that is a multiple of the number of ECMP paths.

As a workaround, the number of members in the LAG can be increased/decreased to no longer be a multiple of the number of ECMP paths. (933531)

- Dynamic (LACP) LAGs may flap during ASU reload, which can result in delayed convergence. (1034873)
- The BfnLoadBalancer agent can restart unexpectedly when the CLI command "show platform barefoot lb nexthop-group <nh-group-name> hash <hash>" is invoked. (1050455)
Series Affected: DCS-7170 Series

CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM and CCS-750 Series

- If link speed is changed on a 2.5G interface, Strata forwarding plane agent may unexpectedly restart with "Process died with signal 3 \ (SIGQUIT\)" signature. Subsequent restart of the Strata agent will resolve the issue. (368424)
- The forwarding slice agent may restart unexpectedly when a SEU is encountered in the BSC_AG_AGE_TABLE during initialization. This may result in brief traffic outage. Switch is expected to start normal operation after the agent restart. (417588)
- Interfaces configured for MAC loopback are limited to 10 Mbps instead of the configured speed. (585395)
SKUs Affected: CCS-720XP-24Y6, CCS-720XP-48Y6 and DCS-7010TX-48

- PhyIsland agent might restart unexpectedly multiple times, while configuring 2.5GBASE-T interfaces or during a reload. This will cause link flaps and traffic disruption on all front-panel BASE-T ports. (586920)

SKUs Affected: CCS-720XP-96ZC2, CCS-722XPM-48ZY8 and CCS-750X-48ZP-LC

- During hitless restart of forwarding plane agent(s), packets which were denied by security ACLs may leak through the chip(s) being managed by the agent(s) for a second. The deny security ACL rules start acting on the traffic once the forwarding plane slice agent(s) is warm.

There is no workaround for this bug. (589311)

Series Affected: CCS-750 Series

- 2.5GBASE-T interfaces may fail to pass traffic in the ingress direction after a reload, linecard insertion (for modular systems) or reconfiguration, if they are configured with `speed 100full` command.

Workaround is to run `shutdown` followed by `no shutdown` on the affected interfaces. (598487)

SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-48ZC2, CCS-720XP-96ZC2, CCS-722XPM-48ZY8 and CCS-750X-48ZP-LC

- When MACsec is disabled on an interface which is used for VXLAN underlay, non-unicast traffic may not seamlessly transition to clear text. The work around is to use "shutdown" followed by "no shutdown" after disabling MACsec on an interface. (637199)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- EVPN VXLAN A-A multihoming is not supported when MACSec is enabled on the core port. Configuring it may result in undesirable double delivery of BUM traffic. (642181)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- When MACsec is disabled on an interface, multicast traffic may not seamlessly transition to clear text. The work around is to use "shutdown" followed by "no shutdown" after disabling MACsec on an interface. (645584)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- On switchcard switchover, there may be extended traffic outage. (682181)

Series Affected: CCS-750 Series

- 2.5GBASE-T ports configured with `speed 100full` (100M forced speed) configuration may fail to pass traffic even if they report linkup after a failing cable test run with the `test 11 cable base-t interfaces INTFS` CLI command.

Workaround is to run `shutdown` followed by a `no shutdown` on the affected interfaces. (800901)

SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-48ZC2, CCS-720XP-96ZC2, CCS-722XPM-48ZY8 and CCS-750X-48ZP-LC

- 5GBASE-T/10GBASE-T ports may stop passing traffic or start causing packet corruption with errors. (e.g. FCS/Symbol/Rx/Runt packet errors).

Workaround is to run `shutdown` followed by a `no shutdown` on the affected interfaces.

The issue may persist after Leaf Smart Software Upgrade (Leaf SSU) reloads from the affected EOS versions. In order to apply the fix on those cases, please run the `reset phy interface Ethernet1 chip <CHIP>` CLI command during a maintenance window. The <CHIP> argument should be "bcm54998e" on CCS-720XP-96ZC2 and "bcm54998es" on CCS-720XP-48ZC2 and CCS-720XP-24ZY4. Please note that this command will cause link flaps and brief traffic disruption on all front panel BASE-T interfaces. (801926)

SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-48ZC2 and CCS-720XP-96ZC2

- 5GBASE-T/10GBASE-T ports configured with `speed 100full` configuration may fail to link up after a cable test is run with the `test 11 cable base-t interfaces INTFS` CLI command, even if the link partners report a linkup.

Workaround is to run `shutdown` followed by `no shutdown` on the affected interfaces. (804553)

SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-48ZC2, CCS-720XP-96ZC2 and CCS-750X-48ZXP-LC

- Under rare circumstances, after the system reloads, certain switch chips may fail to initialize correctly, resulting in Ethernet links being established but network traffic being dropped.

The workaround is to reset the affected switch chip using `platform trident <x> reset`, replacing <x> with the name of the switch chip as shown by `show platform trident system port`. Alternatively, `platform trident \* reset` can be used to reset all switch chips in the system. (808236)

Series Affected: CCS-750 Series

- Excessive link flaps can cause high CPU usage, potentially causing agents to restart. Agent restarts may cause brief traffic disruptions.

Workaround is to configure the desired linkup speed using `speed forced <SPEED>` on the affected interfaces. (815477)

SKUs Affected: CCS-750X-48ZXP-LC and DCS-7050TX3-48C8

- BASE-T ports may stop passing traffic or start causing packet corruption with errors. (e.g. FCS/Symbol/Rx/Runt packet errors).

Workaround is to run `shutdown` followed by a `no shutdown` on the affected interfaces.

The issue may persist after Leaf Smart Software Upgrade (Leaf SSU) reloads from the affected EOS versions. In order to apply the fix on those cases, please run the `reset phy interface <INTF> chip bcm84898` CLI command during a maintenance window. The <INTF> argument should be Ethernet1 on fixed systems and EthernetX/1 on modular systems where X is the number of the affected linecard. Please note that this command will cause link flaps and brief traffic disruption on all front panel BASE-T interfaces of the fixed system switches or the linecards. (816890)

SKUs Affected: 7300X3-48TC4-LC, CCS-750X-48ZXP-LC and DCS-7050TX3-48C8

- 1000BASE-T ports may fail to link up after a reload or forwarding agent restart.

Workaround is to restart the forwarding agent with `platform trident reset` CLI command. This will cause link flaps and brief traffic disruption on all front panel ports.

Affected ports:

DCS-7010TX-48: 25-48

DCS-7010TX-48-DC: 25-48

DCS-7010TX-48C: 25-48

CCS-710P-12: 1-12

CCS-710P-16P: 1-12

CCS-720DP-24S: 1-24

CCS-720DP-48S: 25-48

CCS-720DT-24S: 1-24

CCS-720DT-48S: 25-48

CCS-722XPM-48Y4: 25-48 (818224)

SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DP-24S, CCS-720DP-48S, CCS-720DT-24S, CCS-722XPM-48Y4, DCS-7010TX-48 and DCS-7010TX-48-DC

- 2.5GBASE-T ports configured with `speed auto 100full`/`no speed` configurations and linked up at 100M speed may fail to pass traffic after an SSU reload (i.e. reload fast-boot) if they link down while the SSU reload is in progress, and come back up after the SSU reload is complete.

Workaround is to run `shutdown` followed by a `no shutdown` on the affected interfaces. The issue can also be avoided by configuring the susceptible interfaces with `speed forced 100full` configuration prior to the SSU reload. Please note that the configuration change will flap the interfaces. (840024)

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

- PhyIsland agent restart after an SSO (stateful switchover) may be hitful, causing link flaps and brief traffic disruption on all front panel BASE-T

ports of the affected linecard(s). (841719)

SKUs Affected: 7300X3-48TC4-LC, CCS-750X-48THP-LC, CCS-750X-48TP-LC, CCS-750X-48ZP-LC and CCS-750X-48ZXP-LC

- BASE-T interfaces may stay down indefinitely after forwarding plane agents restart if the interfaces are configured with the default speed or the `speed auto` config.

Workaround is to configure the BASE-T interfaces with `speed auto <SPEED>` command where the <SPEED> is the desired linkup speed in order to avoid the issue, or to run `shutdown` followed by `no shutdown` on the affected interfaces in order to recover. (852418)

SKUs Affected: 7300X3-48TC4-LC, CCS-720XP-24Y6, CCS-720XP-24ZY4, CCS-720XP-48Y6, CCS-720XP-48ZC2, CCS-720XP-96ZC2, CCS-722XPM-48ZY8, CCS-750X-48THP-LC, CCS-750X-48TP-LC, CCS-750X-48ZP-LC, CCS-750X-48ZXP-LC and DCS-7050TX3-48C8

- BASE-T interfaces in the range Et1-Et24 may stay down indefinitely after forwarding plane agents restart if the interfaces are configured with the default speed or the `speed auto` config.

Workaround is to configure the BASE-T interfaces with `speed auto <SPEED>` command where the <SPEED> is the desired linkup speed in order to avoid the issue, or to run `shutdown` followed by `no shutdown` on the affected interfaces in order to recover. (852419)

SKUs Affected: CCS-720DP-48S, CCS-722XPM-48Y4, DCS-7010TX-48 and DCS-7010TX-48-DC

- 5GBASE-T interfaces may stay down indefinitely after forwarding plane agents restart if the interfaces are configured with the default speed or the `speed auto` config.

Workaround is to configure the BASE-T interfaces with `speed auto <SPEED>` command where the <SPEED> is the desired linkup speed in order to avoid the issue, or to run `shutdown` followed by `no shutdown` on the affected interfaces in order to recover. (852421)

SKUs Affected: CCS-710P-16P

- The QUIC crash issue could be seen in case of truncated or missing Quic CHLO packet. It may also happen during the reconstruction of the deciphered Quic CHLO when there are several fragments. (940087)
- BASE-T ports may fail to link up, even if the link partners report a linkup.

On EOS versions prior to 4.29.0F, on fixed systems, the workaround is to run `platform trident reset`; on modular systems, the workaround is to run `platform trident <linecard> reset` on the affected linecard. On EOS versions 4.29.0F and onwards, the workaround is to run `reset platform

trident phy interface INTF pcs` on one of the affected interfaces.

Affected ports:

- CCS-720DP-48S 1-24
- CCS-720DT-48S 1-24
- CCS-720XP-24Y6 1-24
- CCS-720XP-48Y6: 1-48
- CCS-722XPM-48Y4: 1-24
- CCS-750X-48THP-LC: 1-48
- CCS-750X-48TP-LC: 1-48
- DCS-7010TX: 1-24 (1031804)

Series Affected: DCS-7010TX Series

SKUs Affected: CCS-720DP-48S, CCS-720XP-24Y6, CCS-720XP-48Y6, CCS-722XPM-48Y4, CCS-750X-48THP-LC and CCS-750X-48TP-LC

- When fan 1 is removed, EOS may lose control of the 2nd fan leading to overheating. (1086825)

CCS-750 and DCS-7060X4 Series

- In a chassis with redundant supervisors configured with the SSO redundancy protocol, after a switchover has been performed, a forwarding agent restart may cause the active supervisor to experience a fatal CPU error and restart, which will cause another switchover. (586096)

Series Affected: CCS-750, DCS-7300X and DCS-7300X3 Series

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

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

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

DCS-7130 Series

- When configured to an illegal speed, interfaces will go to inactive rather than errdisabled. This will cause the interface to seemingly disappear from 'show interface status' unless 'service interface inactive expose' setting is used to display it.

Workaround: change the speed configuration of the interface to a valid speed. To check which speeds are valid, please use 'show interface hardware'. (457130)

SKUs Affected: DCS-7130LBR-48S6QD, DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R

- Some transceivers sometimes report an invalid EEPROM checksum.

Workaround is to unplug/replug the module. (670128)

SKUs Affected: DCS-7130-16G3-F (HwRev:C1), DCS-7130-16G3-R (HwRev:C1), DCS-7130-16G3S-F (HwRev:D1), DCS-7130-16G3S-R (HwRev:D1), DCS-7130-48E-F (HwRev:A4), DCS-7130-48E-R (HwRev:A4), DCS-7130-48EH-F (HwRev:A4), DCS-7130-48EH-R (HwRev:A4), DCS-7130-48EHS-F (HwRev:B1), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48EHS-R (HwRev:B1), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48G3-F (HwRev:B1), DCS-7130-48G3-F (HwRev:B2), DCS-7130-48G3-R (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48G3S-F (HwRev:C1), DCS-7130-48G3S-R (HwRev:C1), DCS-7130-48L-F (HwRev:A2), DCS-7130-48L-F (HwRev:A3), DCS-7130-48L-R (HwRev:A2), DCS-7130-48L-R (HwRev:A3), DCS-7130-48LA-F (HwRev:A2), DCS-7130-48LA-F (HwRev:A3), DCS-7130-48LA-R (HwRev:A2), DCS-7130-48LA-R (HwRev:A3), DCS-7130-48LAS-F (HwRev:B2), DCS-7130-48LAS-F (HwRev:B3), DCS-7130-48LAS-R (HwRev:B2), DCS-7130-48LAS-R (HwRev:B3), DCS-7130-48LB-F (HwRev:A2), DCS-7130-48LB-F (HwRev:A3), DCS-7130-48LB-R (HwRev:A2), DCS-7130-48LB-R (HwRev:A3), DCS-7130-48LBA-F (HwRev:A2), DCS-7130-48LBA-F (HwRev:A3), DCS-7130-48LBA-R (HwRev:A2), DCS-7130-48LBA-R (HwRev:A3), DCS-7130-48LBAS-F (HwRev:B2), DCS-7130-48LBAS-F (HwRev:B3), DCS-7130-48LBAS-R (HwRev:B2), DCS-7130-48LBAS-R (HwRev:B3), DCS-7130-48LBS-F (HwRev:B2), DCS-7130-48LBS-F (HwRev:B3), DCS-7130-48LBS-R (HwRev:B2), DCS-7130-48LBS-R (HwRev:B3), DCS-7130-48LS-F (HwRev:B2), DCS-7130-48LS-F (HwRev:B3), DCS-7130-48LS-R (HwRev:B2), DCS-7130-48LS-R (HwRev:B3), DCS-7130-96-F (HwRev:A7), DCS-7130-96-R (HwRev:A7), DCS-7130-96-R (HwRev:A8), DCS-7130-96L-F (HwRev:A2), DCS-7130-96L-F (HwRev:A3), DCS-7130-96L-R (HwRev:A2), DCS-7130-96L-R (HwRev:A3), DCS-7130-96LA-F (HwRev:A2), DCS-7130-96LA-F (HwRev:A3), DCS-7130-96LA-R (HwRev:A2), DCS-7130-96LA-R (HwRev:A3), DCS-7130-96LAS-F (HwRev:B1), DCS-7130-96LAS-F (HwRev:B2), DCS-7130-96LAS-R (HwRev:B1), DCS-7130-96LAS-R (HwRev:B2), DCS-7130-96LB-F (HwRev:A2), DCS-7130-96LB-F (HwRev:A3), DCS-7130-96LB-R (HwRev:A2), DCS-7130-96LB-R (HwRev:A3), DCS-7130-96LBA-F (HwRev:A2), DCS-7130-96LBA-F (HwRev:A3), DCS-7130-96LBA-R (HwRev:A2), DCS-7130-96LBA-R (HwRev:A3), DCS-7130-96LBAS-F (HwRev:B1), DCS-7130-96LBAS-R (HwRev:B1), DCS-7130-96LBS-F (HwRev:B1), DCS-7130-96LBS-F (HwRev:B2), DCS-7130-96LBS-R (HwRev:B1), DCS-7130-96LBS-R (HwRev:B2), DCS-7130-96LS-F (HwRev:B1),

DCS-7130-96LS-F (HwRev:B2), DCS-7130-96LS-R (HwRev:B1), DCS-7130-96LS-R (HwRev:B2), DCS-7130-96S-F (HwRev:B1) and DCS-7130-96S-R (HwRev:B1)

- Removing the cable from a BASE-T transceiver or shutting down the peer interface may leave the FPGA interface behind the transceiver in a connected state. (797593)

Series Affected: DCS-7130 Series

- In some cases when there is simultaneously a routed and a dynamically source-NAT-ted traffic, the rate of the NAT-ted one might be limited. Workaround: don't configure a behind-the-NAT subnet and a routed subnet within the same port group (ports are grouped in fours). (884065)

7358, 7368, 7388, CCS-750, DCS-7300X, DCS-7500R, DCS-7800 and DCS-7800R3 Series

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

- A Linecard with link issues on the control plane may cause the Supervisor to reboot (670003)
 - Correctable PCIe errors may be seen on a forwarding chip control plane link. (772195)
- Series Affected: DCS-7800 Series

- If ManagementActive interface is configured with redundancy enabled, then ManagementActive agent can restart unexpectedly during supervisor switchover. (805315)
- The SandFap agent may restart unexpectedly if a PTP packet for one-step boundary clock mode is received when PTP is not enabled.

To prevent such restarts, configure "ptp mode boundary" temporarily. PTP can be turned off safely after approximately one minute. (810008) (1)

Series Affected: DCS-7500R2 Series

- Supervisor may become unresponsive on reboot and require re-seating or chassis power cycle to recover (921207)

SKUs Affected: DCS-7800-SUP, DCS-7800-SUP1A, DCS-7800-SUP1S, DCS-7816-SUP and DCS-7816-SUP1S

- Scaled ACL configuration may cause traffic drop of few seconds during SSU upgrade. Workaround is to use CLI "hardware access-list update default-result permit" to disable port blocking while applying ACL (1037945) (1)
- Continuous PCIe corrected communication issues can lead to high CPU usage (1049285)

(1) This item is in two or more sections on this document

7500R3, DCS-7020, DCS-7130, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280R3A, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800, DCS-7800R3 and DCS-7800R3A 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)
- PBR policies applied on interfaces do not work if the interface has some subinterfaces that also have PBR policies applied. (243286)
Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series
- VTEP cannot learn a remote host's MAC address from VXLAN encapsulated IPv6 Neighbor Advertisement packets received from that host. This might be an issue in a pure IPv6 deployment where the VTEP does not receive any bridged traffic from the host. In such a scenario, the traffic toward the host will be flooded in the VLAN.

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

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

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

- VXLAN bridging traffic sent out of LAG interfaces may be temporarily dropped after an SSO event. (251398)
Series Affected: DCS-7500R Series
- When a DirectFlow flow is created with an action to drop the outer tag, the packet is still leaving the interface tagged. (264336)
Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

- When a physical port (e.g. Et1) is removed from a Port-Channel (e.g. Po1) that has subinterfaces configured (e.g. Po1.1), and there are subinterfaces configured under that physical port (e.g. Et1.1), the subinterfaces under the physical port can be incorrectly programmed. This results in traffic being dropped for the incorrectly programmed subinterfaces.

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

- If one or more linecards in a modular system is running in tap aggregation mode, a small percentage of multicast traffic on linecards running in normal mode may be dropped, following a change to membership of some tap aggregation destination groups. (275286)
- 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)
- 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)

- 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)
- While a line card is being power cycled, corrupt VXLAN encapsulated packets may be sent. (441141)
- Performing many updates to a PBR policy in a short amount of time may cause the SandAcl agent to restart unexpectedly. (441912)

- 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)
- "match nexthop-group" filter is not supported in PBR rules when the corresponding nexthop-group is referenced by ECMP routes. (458189)
- 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)
- 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
- 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)
- With MPLS VPN, during stateful switchover, some VPN traffic after MPLS decapsulation may be transiently forwarded in the default VRF instead of the appropriate tenant VRF. (541206)
- An over-voltage condition on the POS1V2_HBM_IO_JE0 power rail can cause the system to power-cycle. (544659)
SKUs Affected: DCS-7280CR3-96 and DCS-7280CR3K-96
- 'Persistent Internal Drop DeqDeletePktCnt' error messages are seen due to incorrect/missing system ACEs. Hitful restart of SandAcl agent is needed to resolve the issue. The trigger for this issue is unknown. (548524)
- When 'ip-length' key field is missing from the "ipsec" TCAM feature in the operational TCAM profile, all traffic on the IPsec tunnels may be trapped to the CPU, leading to low tunnel throughput.

The workaround is to enable the 'ip-length' key field in the "ipsec" TCAM feature. (588127)

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

- On a shaped subinterface tx-queue, configuring the bandwidth doesn't take effect unless no-priority is also configured. (588133)
- If 'cpu traffic-policy <policyName> vrf all' is configured, where the traffic-policy contains rules with the policer action, and the operating TCAM profile does not support the CPU traffic-policy feature, then configuring any

routed port in the switch may lead to an unexpected SandAcl agent restart. The workaround is to remove the 'cpu traffic-policy <policyName> vrf all' configuration. (619998)

- Postcard sampling doesn't work for traffic that is forwarded out of a sub-interface that has shaping applied. (621932)
- Postcard Telemetry sampling doesn't work for IPv6 packets if the egress interface has an IPv6 egress ACL applied on it. (624010)
- Continuous shut/no shut on a Port-Channel with 1000s of subinterfaces can cause the SandTunnel agent to restart unexpectedly. (718924)
- When using a large scale of PBR rules, removal and addition of member links used as nexthop groups destinations of PBR policies may cause the SandAcl agent to restart. (759846)
- On a config-replace, a change in the PBR config that qualifies for in-place update may still cause traffic loss. (760884)
- Performing many updates to a traffic policy in a short amount of time may cause the SandAcl and/or SandAegis agent to restart unexpectedly. (770478)
- When a LAG subinterface is configured and a LAG member is subsequently removed, the outer VLAN tag for packets egressing that subinterface may be improperly removed.

To recover proper behavior, issue a shut/no shut sequence on the malfunctioning subinterface. (783098)

- Added support for QoS policy-map names with more than 64 characters. (791679)
- Flapping the interface that is the destination for a large number of redirect actions in a traffic-policy may result in a long programming delay to updating redirect actions in hardware and to any future changes in the traffic-policy config. (814453)
- In rare instances, entries in a TCAM bank may fail to be deleted before the TCAM bank is released by one feature and then allocated to another feature, resulting in non-deterministic behavior in the new feature. (826437)
- In rare cases, it is possible when making egress interface traffic-policy changes for SandTcam to continuously crash upon policy application or removal. The workaround for this is to removal policy application from all interfaces and then re-apply the policies. (849779)
- During normal operation, the message 'HARDWARE-3-TIMESYNC_ERROR ... (slave frequency check error)' may be logged if the internal time synchronization

in a switch is disrupted. The internal time synchronization and PTP will recover shortly without any user intervention. (901828)

- When inband or postcard telemetry is enabled, an interface disappearing (e.g. due to a forwarding plane agent restart) can cause the system to reload unexpectedly due to a kernel panic.

Disable telemetry features to workaround this issue. (933485)

- CPU traffic-policy installs incorrect iptables filter entries for 'initial' and 'not-*' TCP flag matches (e.g. 'not-fin', 'not-sync', 'not-rst', 'not-psh', 'not-ack', 'not-urg'). (949376)
- Multiple updates to access-lists matched by PBR policy-maps in a short period of time may cause transient misforwarding of traffic. The issue is expected to be resolved once the feature agent processes all the updates. (956556)
- CPU traffic-policy may fail to program if 'cpu traffic-policy POLICY vrf all' is configured on a switch with lots of VRFs and the applied policy contains a large number of matches with police actions. (957530)

- In an EVPN VXLAN network deploying active/active multi-homing using shared Ethernet segments, when the MAC tables on the VTEPs become momentarily inconsistent, there might a brief period of duplicate delivery or packet loss.

Typically, the MAC tables become inconsistent when a new MAC is learnt on one of the VTEPs and is yet to be advertised to other VTEPs or when a MAC is removed on one of the VTEPs and it is yet to be withdrawn from other VTEPs. (957981)

- Security Advisory for 0108 for CVE-2024-6437. For more information, see <https://www.arista.com/en/support/advisories-notices/security-advisory/20689-security-advisory-0108> (962149) (1)
- With VXLAN routing configurations and the maximum hardware FEC hierarchy level set to 2, packets that are routed over a VXLAN tunnel might be corrupted or missing the inner Ethernet header, if the underlay next-hop MAC address is not learned.

Workaround is to force the underlay next-hop MAC address to be learned (e.g. by pinging the next-hop IP) or using static MAC address. (967128)

- In a VXLAN routing deployment with "ipv6 address virtual" configured on VXLAN VLAN that has VLAN translation:
 - the switch may incorrectly advertise the virtual anycast with switch physical MAC instead of the virtual-router MAC
 - the switch may not respond to neighbor solicitations for the shared link local address corresponding to the virtual-router MAC

- duplicate address detection may fail

The workaround is not use VLAN translation on any ethernet or port-channel interface that is on a VXLAN VLAN. (968877)

Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280R3A, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800R3 and DCS-7800R3A Series

- When a forwarding ASIC gets reset, the device may unexpectedly reload. (990910) (1)
- Applying a traffic-policy with 'redirect next-hop group NAME' on a non-default VRF interface causes SandAegis to restart continuously. (1001193)
- Sometimes Macsecproxy patch interface incorrectly does not terminate the internal VLAN tag resulting in an inoperational Macsecproxy subinterface.

Workaround is to delete and readd the MacsecProxy patch sub interface. Then unconfigure and configure MACsec proxy on MACsec proxy subinterface. (1005850)

- In an EVPN VXLAN active-active multi-homing setup, if the designated forwarder (DF) election result changes rapidly or the interfaces in local ethernet segments flap rapidly, then multi-homed hosts connected to the impacted ethernet segments could start receiving duplicate BUM packets. (1008046) (1)
- When running Packet Link Qualification via gNOI, the hardware programming setup duration can be longer.

The minimum setup duration reported in Packet Link Qualification gNOI Capabilities API should be used during the Packet Link Qualification gNOI Create API request. (1015120)

- In an EVPN VXLAN Campus Wifi gateway switch where VTEP-to-VTEP bridging is configured, SandL3Unicast agent restart may result in some traffic loss towards the AP VTEPs. (1019630)
- When the active configuration on a device contains a URL field-set, performing a config-replace where the new configuration removes the field-set definition will still leave the field-set defined. To workaround this, delete the field-set via interactive CLI before perform the config-replace. If the issue is encountered after config-replace then re-defining and removing the field-set via interactive CLI gets the device out of the issue state. (1021152)
- Suboptimal voltage configuration may lead to power faults requiring a reload to recover (1037944)
SKUs Affected: 7289

- If an ACL applied on a monitor session contains a rule that expands into multiple TCAM entries in hardware, only the matches against the final TCAM entry will result in the packet being mirrored. Rules that may expand into multiple TCAM entries include range matches (lt or gt) on the following match criteria: TTL, hop-length, IP-length, l4-src-port, and l4-dst-port. (1049418)
- SandL3Unicast agent might restart unexpectedly if a route resolves over an ECMP of next hops that includes both management and front-panel ports. (1051007)
- Fans may erroneously be marked as failed after hot swap (1057286)
Series Affected: 7368 Series
- In OISM setups with the ingress replication mode enabled, bridged IPv4 multicast packets can be head end replicated to remote VTEP devices when underlay multicast tunnels do not have remote receivers. A workaround is to change the global replication mode to the fabric-egress replication mode. (1086873)
- When "hardware next-hop fast-failover" is configured, the hardware programming acknowledgement might be premature for certain adjacencies. (1135705)
- When a traffic-policy has rules with redirect next-hop group actions and next-hop group is updated from NHG1 => NHG2, where both next-hop groups are already resolved, the hardware will not be updated and still remain with a redirect to NHG1.

To work around this issue, remove the affected match and action from the policy and then add it back. Alternatively, unapply the affected policy from all interfaces and then re-apply the policy. (1145425)

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

- 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)

- When MACsec is enabled and a SSO failover occurs, the Evora agent may restart due to "LMI Access timeout", which may cause links to flap resulting in extended traffic loss. (494072)
SKUs Affected: 7500RM-36CQ-LC
- Private VLAN may not work correctly after a stateful switchover. The workaround is to restart the SandTunnel agent. (505258)
- IQM_RST_USCNT_ERR, IQM_FR_FIMC_DB_ROLL_OVER, IQM_FR_MNMC_DB_ROLL_OVER interrupts of Jericho chip indicate traffic to and from ports of the affected chip to be dropped. A full reset of the chip could workaround this issue. (540881)
- Performing ASU on a Hopland while hardware accelerated BFD is running, may cause SandBfd to continuously restart.

Workaround: Disable BFD hardware acceleration before ASU is performed. (840813)

DCS-7280R2 and DCS-7500R2 Series

- The default forwarding behavior for IS-IS packets received on subinterface or vlan-tag-based pseudowire connectors is now "trap" rather than "forward".

To achieve forwarding of IS-IS packets received on these connector types, apply an L2 Forwarding Profile, as described in the TOI "L2 Protocol Forwarding". This capability for IS-IS is supported in newer releases only. (611787)

- Multicast traffic on non-default VRF will be lost for as long as 15 minutes when default VRF is multicast enabled without any L3 interfaces. Workaround is to disable mulitcast on default VRF or configure a dummy L3 interface. (619693)

7500R3, 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)
- ECN bits may not be correctly copied from outer header when packets are routed after VXLAN decapsulation. (489998)
- 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 continuously restart. (492549)
- If the SandFapNi agent is restarted there is a chance it will enter a restart loop due to the pcieSwitchDynamicConfigDir object no longer being available.

Please reload the device to resolve this issue. (560980) (1)

- When an interface traffic-policy is changed such that hardware resources are exhausted by the change, SandAegis agent may restart. After the agent restart, the switch will recover after traffic-policy rollback kicks in. (574639)
- On applying/removing a Egress TC-to-CoS map on L2 LAG Subinterfaces, the packets going out of the Subinterfaces can have wrong dot1q tag. Workaround is to shut the subinterface, apply the map and then unshut it (582063)
- ECN bits are not preserved when packets are bridged over a VXLAN tunnel. (587486)
- When speed of a port is changed while it is carrying traffic, speed change may fail, causing interface to become inactive and potentially leading to a forwarding agent restart.

One workaround is to change the speed-group configuration of problematic port and revert it back to the previously applied speed-group configuration. Another workaround is to change the speed of primary port to 400g-8 and revert it back to the previously applied speed. (597294)

- Following a Stateful Switchover, disabling power to either supervisor using "no power enable module Supervisor" may cause some control plane agents to restart unexpectedly and repeatedly.

To resolve the agent restarts, restart the Linecard agents one at a time. (606781)

- Priority-based flow control fail to work after SSO on some instances. Workaround is to restart SandFapNi agent for the linecard after scheduling downtime. (606784)
- The system could end up in an unrecoverable power cycle loop. (773897)
- Configuring 256 Ipsec or Secure Vxlan tunnels while also configuring MACsec can result in continous B52 agent restarts. Workaround is to reduce tunnel scale or to remove MACsec configuration. (871256)
- When hardware accelerated sFlow is configured at the same time as "sflow sample input subinterface" sampling packets with sizes 64-196 may cause the acceleration hardware to become stuck and stop sampling packets. The workaround is to disable sFlow , restart the SandFapNi agent for the corresponding linecard, and re-enable sFlow. Note that if the "sflow sample input subinterface" is not removed, the bug may reoccur. (947343) (1)
- Applying speed change configuration on underlay interface of an IPsec tunnel or flapping underlay interface of an IPsec tunnel configured with anti-replay

detection may lead to complete traffic loss on the tunnel.

A workaround is to flap the IPsec tunnel. (980557)

- Unexpected behavior in edge cases with the Tunnelsec agent and VXLAN. This is tracked under <https://www.arista.com/en/support/advisories-notices/security-advisory/21289-security-advisory-0113>.

The workaround is to remove and re-apply security profiles for each secure VTEP. (997526)

SKUs Affected: DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

- Multicast packets egressing from an L3 subinterface may have incorrect VLAN tags when the following conditions are all true:
 1. A patch panel is programmed using a subinterface connector on the same port as the L3 subinterface.
 2. The patch subinterface connector has an outer VID which matches the L3 subinterface internal VLAN ID.
 3. The patch panel or subinterface connector is shut (or the patch goes down due to remote endpoint), thereby causing the subinterface connector to become unprogrammed.

Workaround is to flap the L3 subinterface after patch is down. (1016040)

- On the device running multicast IPv6, when reloading the device, multicast traffic might be lost and won't recover automatically.

Workaround: Restart the SandMcast agent. (1031956) (1)

- Transient traffic loss may be seen during ECMP expansion for a route that pushes MPLS labels and resolves over an RSVP LER Sub-tunnel. This includes SR or LDP or BGP tunnel pushing a VPN or unicast label and resolving over the RSVP LER Sub-tunnel. (1044328)
- Removing secure VXLAN profile from a remote VTEP can result in traffic loss on all front panel ports, workaround is to shut the affected interfaces and unshut them. (1064076) (1)

SKUs Affected: DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

7500R3, DCS-7130, DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series

- If the SandFapNi agent is restarted there is a chance it will enter a restart loop due to the pcieSwitchDynamicConfigDir object no longer being available.

Please reload the device to resolve this issue. (560980) (1)

- Systems shipped early in the life cycle of this product have been programmed with an incorrect ISL chip image for the voltage regulator monitor (VRM), which may cause HBM CRC errors and lead to packet drops.

A swix is available with an updated ISL image to fix this issue. (778115)

SKUs Affected: DCS-7280SR3-48YC8 and DCS-7280SR3K-48YC8

- When hardware accelerated sFlow is configured at the same time as "sflow sample input subinterface" sampling packets with sizes 64-196 may cause the acceleration hardware to become stuck and stop sampling packets. The workaround is to disable sFlow, restart the SandFapNi agent for the corresponding linecard, and re-enable sFlow. Note that if the "sflow sample input subinterface" is not removed, the bug may reoccur. (947343) (1)
- Packets dropped due to split horizon can cause malformed packets being transmitted causing FCS errors on connected device. (1018301) (1)

DCS-7280R3A and DCS-7800R3A Series

- If the SandFapNi agent is restarted there is a chance it will enter a restart loop due to the pcieSwitchDynamicConfigDir object no longer being available.

Please reload the device to resolve this issue. (560980) (1)

- Link may not come up on certain 100GBASE-PSM4 modules. (929400) (1)
SKUs Affected: DCS-7280DR3A-54, DCS-7280DR3AK-54 and DCS-7280DR3AM-54
- Packets dropped due to split horizon can cause malformed packets being transmitted causing FCS errors on connected device. (1018301) (1)
- After configuring Macsec on breakout ports, some MKA controlled ports may not come up. (1024465)

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)
- Hardware corruption can lead to switch dropping all packets. (853697)
- When an IPsec tunnel is established with a non-default MTU for the tunnel interface, MTU violations may not be detected. This could result in packets greater than the MTU size being forwarded without fragmentation, or dropped.

In order to get proper MTU handling for that tunnel, the tunnel interface should be shut/no shut. (1036933)

DCS-7280R3 Series

- With Jumbo MTU packets we could have some FCS errors on the peer. (747305)

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

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

(1) This item is in two or more sections on this document

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

- A linecard addition or hotswap operation on a switch configured with "port-channel load-balance sand replication egress" may lead to continuous RQP discards and STP flaps.
A workaround is to unconfigure "port-channel load-balance sand replication egress". (353249)
Series Affected: DCS-7500R2 Series
- PBR, QOS, and BGP Flowspec policies applied on routed interfaces will not take effect if sub-interfaces are configured on that routed interface. PBR, QOS, and BGP Flowspec policies applied on SVIs will also not take effect on traffic ingressing SVI member interfaces if VLAN mapping configuration is present. (372390)
- 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)
- 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)

- Traffic on a subinterface will be subject to BGP Flowspec matches applied to the parent interface, unless the TCAM profile has the flowspec feature configured with 'port qualifier size N', where N > 1. (541950)
- Traffic policy matches on VLAN tags may incorrectly match untagged packets by interpreting part of the untagged packet as a VLAN tag. (671474)
- If a large number of subinterfaces is configured on a parent interface which is a breakout port, and a speed change of another interface causes that breakout port to become inactive, the SandFap agent may unexpectedly restart once, flapping all links. After the SandFap agent restarts, the system will recover, the speed change will be completed, and the affected subinterfaces will no longer be configured in the hardware.

As a workaround, remove all subinterfaces from the associated breakout ports before changing the speed of a front-panel port. (691732)

- For a traffic-policy containing matches with police actions applied to a linecard's interfaces, the police action may fail to reprogram on that linecard after a linecard hot swap.

To recover, re-install the traffic policy on the interfaces where the police action failed to program. (714936)

- In an EVPN network, SandMact agent might restart due to a SIGQUIT while processing the MAC addresses polled from the hardware MAC table if the number of MAC addresses is close to the hardware MAC table limit. (718204)
- When a static MAC is configured on an MLAG port-channel, traffic may be dropped while the local port-channel is down.

The workaround is to enable the fast MLAG unicast convergence feature. (723709)

- VXLAN source VTEP pruning with VTEP-to-VTEP bridging does not work for VXLAN ARP packets. (797624)
- VXLAN source VTEP pruning with VTEP-to-VTEP bridging might be broken after some underlay route changes, such as ECMP expansion.

The workaround is to run "agent SandL3Unicast terminate" command, which could be traffic impacting depending on which features are configured. (797824)

- Ingress IPv4 and IPv6 router ACL features do not work on GRE, IP-in-IP, and GUE decapsulated IPv4 and IPv6 traffic, respectively. (799085)

- SandL3Unicast agent may unexpectedly restart under VXLAN tunnel ECMP FEC table overflow condition. (801962)
- DirectFlow cannot enforce a new destination with its actions 'output nexthop' and 'output interface' for traffic destined to the CPU. The flow's counters will increase but the traffic will not be redirected. (809100)
- The SandFap agent may restart unexpectedly if a PTP packet for one-step boundary clock mode is received when PTP is not enabled.

To prevent such restarts, configure "ptp mode boundary" temporarily. PTP can be turned off safely after approximately one minute. (810008) (1)

Series Affected: DCS-7500R2 Series

- On a reboot chassis might never come up with agent SandFap restarting continuously due to attrlog overflow. (827602)
- An interface with a 1000BASE-SX or 1000BASE-LX transceiver inserted configured with "speed auto" connected to a peer that has auto-negotiation disabled may incorrectly link up, causing packet loss.

Workaround is to disable auto-negotiation. (829740)

- Packets ingressing in a VLAN without SVI configuration may be unexpectedly routed after the following sequence of events:
 - the VLAN is made secondary to a primary VLAN without SVI configuration
 - SVI is configured for the primary VLAN
 - secondary VLAN configuration is removed for the VLAN

A workaround is to restart the SandFap agents. (857330)

- When VLAN interface egress counters are enabled, the SandTunnel agent may restart following a switchover or transceiver pull/insert, but recovers normally. As a workaround to avoid the SandTunnel restart (and resulting packet loss) during a planned maintenance event, the customer may choose to follow this procedure:

- 1.) Disable counters using "no hardware counter feature vlan-interface out".
- 2.) Using the "show platform fap eedb outac" command, wait until no entries appear in the "poolSviEgress" section.
- 3.) Perform maintenance action, such as transceiver swap.
- 3.) Re-enable counters using "hardware counter feature vlan-interface out". (879576)

- SandTunnel agent may continuously restart on modular switches after a supervisor switchover occurs. This occurs if SandTunnel detects corrupted data in GLEM shadow entries which are stored in DMA memory and part of a Picasso chunk.

Two options to recover from this state:

- 1) Reload the switch, or
- 2) Clear the dma memory used for SandTunnel Picasso chunk, then restart SandTunnel agent:

Step 2a: from bash, "sudo dmamem -f Picasso-SandTunnel"

Step 2b: from CLI config mode, "agent SandTunnel shutdown", then "no agent SandTunnel shutdown". (903750)

- When PFC is configured on some traffic classes, non-PFC traffic classes can drop packets not based on congestion for the destination port but because of the number of buffers used by the ingress port and traffic class. (934999)
- A port in PFC watchdog state can drop all control packets being received. (947091)
- Applying BGP Flowspec rules to a system that has an operating TCAM profile in which "feature flow-spec port ipv4" and/or "feature flow-spec port ipv6" are using 80b key sizes will cause the SandAcl agent to restart continuously. To work around this issue, ensure "feature flow-spec port ipv4" and "feature flow-spec port ipv6" are both using either 160b or 320b key sizes. (972376)
- Packets that match egress security ACL rules which contain a mirror action will be dropped, if the corresponding monitor session is using a GRE destination with an unresolved destination tunnel IP. (980966)
- When forwarding agent restarts or an SSO is triggered after a LC is removed, it can result in traffic loss for flows that are egressing via IP tunnels (1008287) (1)

(1) This item is in two or more sections on this document

7500R3, DCS-7130, DCS-7280R3, DCS-7280R3A, DCS-7500R3, DCS-7800R3 and DCS-7800R3A Series

- When a remote MAC is learnt on an Ingress PE with L2 EVPN MPLS configuration, the traffic destined to the remote MAC may get dropped for a brief amount of time. (551995)
- During normal operation, the message 'HARDWARE-3-TIMESYNC_ERROR ... (master frequency check error)' may be logged if the internal time synchronization in a switch is disrupted. (557400)
SKUs Affected: DCS-7800-SUP
- In the presence of virtual MAC with mask configuration, packets that undergo recirculation may not be properly forwarded. (559242)

- Latency based drop feature is not compatible with features using dynamic rate classes and is applicable only to green packets in traffic stream. Following set of features use dynamic rate classes:
 - 1) Configurable static tail-drops/FADT per speed, per speed+traffic-class tail drop thresholds
 - 2) WRED
 - 3) ECN
 - 4) Interface based per TXQ and per drop precedence tail drop thresholds (618130)
- When the mapping from DSCP, CoS, or MPLS EXP to traffic-class is changed from its default by a qos map, then if the device ingress experiences congestion due to a large number of small packets at line-rate, packets with higher traffic-class with the new mapping may be dropped before packets with lower traffic-class. (626382)
- The "Last LOC Detected" and "Last LOC Cleared" timestamps shown in the output of `show cfm continuity-check end-point` command get cleared if the SandOam agent is restarted. (629830)
- "mac address virtual-router" command no longer works on 7280R3, 7500R3, 7800R3 devices. (633307)
- Egress IPv6 ACL logging do not work on packets subject to layer 3 MPLS termination. (640491)
- If the 'update interface hitless strict' configuration is active, when the FEC associated with the routes from redirect actions used by interface traffic-policies change and traffic-policy hardware usage is near or above 50% of available resources, the FEC programmed in TCAM may not be updated, causing traffic matching a redirect rule to either misroute or drop. The most common reasons for FEC changes are SandL3Unicast restart and/or SandFapNi restart.

To prevent this issue, it is recommended that traffic-policy hardware usage does not exceed 50% of the available resources. (642948)
- SandAcl agent will restart continuously if a TCAM based feature is being configured/modified on a FAP that has linerate traffic ingressing through it. For instance, bootup of linecard. Workaround is to reduce the ingress traffic on that FAP. (645398)
- When config-replace of configuration having VPLS commands is performed back-to-back in quick succession, the SandMact agent can restart unexpectedly. (690059)
- IPv6 linklocal packet received on front panel ports or sourced by CPU do not get headend replicated to remote VTEPs when IPv4 or IPv6 multicast routing is enabled on the VXLAN VLAN. (690105)

- Synthetic Loss Measurement does not work if sFlow, sampled flow tracking or mirroring is active.
Disable or remove sFlow, flow tracking or mirroring configuration. (716811)
- Lots of of route churn can trigger BFD flaps. (723599)
- When segment-security is configured at a high scale, Firewall agent may restart unexpectedly when "shut" or "no shutdown" is issued under "router segment-security", or when a VRF goes down and then comes up. (760155)
- If the SandFapNi agent restarts before it finished initializing, software might not configure the chip properly. Power cycling the affected linecard should restore the correct behavior. (761323)
- Altering MEP attributes, like changing direction, changing profile etc, while CCMs are already running may result in unexpected RDI being reported in the network. A workaround is to restart the SandOam agent on the affected node, i.e. the node sending the RDI after such a churn has been made to the MEP attributes. (766982)
- With EVPN OISM, IPv6 ND link local packets received from the VXLAN tunnel are not flooded to all ports in the VLAN. (772926)
- Running 64 byte packet size RFC2544 generator throughput or frame-loss test causes subsequent RFC2544 Initiator throughput or frame-loss test's of higher packet sizes to fail.
Workaround is to run 128 byte packet size test immediately after running 64 byte packet size test or configure both 64 byte and 128 byte when using multiple packet sizes in RFC2544 generator profile. (773588)
- When removing a port-channel member from a routed port-channel with segment-security configured, traffic may leak through the other port-channel members. (773786)
- When an SVI is configured on the VXLAN VLAN, unknown unicast and multicast packets are unnecessarily sent to the CPU after VXLAN decapsulation. (788956)
- CFM Loss Measurement sessions running on MEPs configured on the same interface at different Maintenance Domain Levels may report incorrect losses. (793764)
- In the presence of a large number of VXLAN MAC addresses, using configure replace to change VXLAN dual stack underlay to IPv4 underlay and also changing the TCAM profile at the same time can lead to permanent traffic loss toward some remote hosts. Workaround is to shut and un-shut the SandMact agent to recover from this state. (802920)

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

- SSO can cause an egress queue to be stuck (809410)
- SandFapNi-Linecard agents may restart continuously with "AttrLog buffer is full" in syslogs (811948)
SKUs Affected: DCS-7812-CH and DCS-7816-CH
- Configurations to disable MAC learning applied to trunk Port-Channels may stop working on VLANs with ACL or policy-maps applied after LAG membership changes. An MLAG peer-link should never learn MACs. If that peer-link is a Port-Channel, then it is susceptible to this issue, and may start to incorrectly learn MACs.

One workaround is to ensure the MAC learning settings for all LAG members match that of the LAG. This can be done by explicit configuration on the LAG members, even though it's dormant configuration when the member is part of a LAG. (813959)

- Hardware accelerated sFlow of Vxlan decapped packet miss first 12B from header. (829460)
- Traffic flowing through the repair path of a VRF selection policy that is forwarding via an IP-in-IP decap and encap tunnel may be transiently misforwarded after either a stateful switchover of supervisor or a Layer3 forwarding agent restart. (834783)
- On a switch having an interface configured with traffic-policy unshared policer, unconfiguring the traffic-policy while the SandAegis agent is restarting may cause the SandAegis agent to unexpectedly restart one more time. (836424)
- Executing a config replace between two scaled traffic-policy configurations with unshared policer actions may fail. Additionally, the subsequent config rollback may also fail causing the switch to unapply traffic policies configured on an interface. (838906)
- When modifying interface traffic-policies where we have enough LPM hardware resources for a hitless install, but we do not have enough TCAM resources for a hitless install, we can get into a state where we are incorrectly matching on prefixes from older policy configuration versions. (847950)
- Sub-optimal convergence may be seen in certain scenarios, with "hardware next-hop fast-failover" and BGP PIC configuration. This issue is observed only on non-BK R3-series systems using non-default MDB profiles (l3-xxl, l3-xxxl, or balanced-xl). (854811)
- Configuring more than the maximum number of ports and doing a stateful switchover causes counters to not get updated. (855905)

SKUs Affected: DCS-7816-CH

- In EVPN VXLAN multicast setup, IPv6 ND packets received on an interface that is attached to a multicast routing enabled SVI might be incorrectly forwarded back to the ingress interface.

One workaround to this problem is to disable ND snooping by apply the following configuration:

```
ipv6 prefix-list ND_PROXY_DISABLED
  deny ::/0
router l2-vpn
  nd proxy prefix-list ND_PROXY_DISABLED (862203)
```

- Upon VLAN membership changes, policy-map, ACL, or similar features may not be applied to access ports in the VLAN. For a full list of impacted features please contact support.

Possible workarounds include: reapplying the feature to the impacted VLANs or SVIs, shut/no-shut the impacted access ports. (864384)

- Under rare circumstances traffic toward virtual router MAC may be punted to the CPU rather than routed out of the switch. The issue might impact the virtual MACs used by VRRP groups and MLAG Peer gateway.

The workaround is to flap the virtual router MAC features on the impacted interfaces. (866110)

- In a VXLAN routing setup using IPv6 underlay, if the overlay SVI is configured with "ipv6 address virtual" then the virtual VTEP functionality will not work correctly. Specifically, neighbor advertisement responses destined to this SVI will fail to sync across to other L3 VTEPs and hence neighbor resolutions may fail. Also, the ND proxy functionality on AG L3 VTEPs in campus deployments will not work. (868642)
- Traffic-policies with matches on inner fields of GTPv1 packets may match incorrectly on inner L4 protocols (like TCP, UDP). (875226)
- IS-IS sessions spanning across an LDP pseudowire not using the FlexEncap CLI for connectors will not be established.

One workaround is to switch to the FlexEncap CLI configuration model for LDP pseudowires. Another workaround is to upgrade to an EOS release where this limitation has been resolved. (882952)

- During normal operation, the internal time synchronization in a switch may be disrupted when the messages 'TIMESYNC_ERROR ... (slave signal CRC error)' and 'TIMESTAMP_COUNTER_RESET' are logged.

A workaround is to power off the supervisor card in the standby slot. (886222)

SKUs Affected: 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC,

7800R3A-36DM2-LC, 7800R3A-36P-LC, 7800R3A-36PM-LC, 7800R3AK-36DM-LC, 7800R3AK-36DM2-LC and 7800R3AK-36PM-LC

- In the EVPN VXLAN OISM setup, when the number of VTEPs or that of the configured multicast decapsulated groups goes beyond the limit, restarting SandTunnel or doing a removal and insertion of a linecard modular would lead to losses of some traffic flows.

A workaround to this issue is to restart SandTunnel after the number of VTEPs or the configured multicast decapsulated groups falls within the support range. (906083)

- SSO can cause an egress queue to be stuck (908064)
- If only egress traffic-policy is configured in the TCAM profile, but ingress traffic-policy is not, running 'show traffic-policy interface counters' for an egress traffic-policy will cause the ConfigAgent to restart. (908122)
- While in Tap Aggregation mode, linerate traffic on tool ports that have "switchport tool allowed vlan" configured may have corrupted FCS. (910965)
- In ZeroTouch provisioning mode, all interfaces using QSFP modules or QSA adaptors whose link partner expects 40G or 10G speeds will not link up and cannot be used for DHCP bootstrap requests.

Workaround this by using an interface at a different configuration or the Management interface for ZeroTouch provisioning instead. (920757)

SKUs Affected: DCS-7280DR3A-54-F, DCS-7280DR3AK-54-F and DCS-7280DR3AM-54-F

- With VXLAN underlay multicast, it is possible that VXLAN encapsulated unknown multicast traffic is looped among VTEPs when a PEG DR exists in the network with external source and receivers. (921213)
- When 3 level hierarchical FECs is enabled in hardware, we might see slower convergence when an IS-IS SR or SR-TE tunnel is configured as a backup and this tunnel in turn also has a backup, and link state change events trigger migration from primary to backup via. Note that 3 level hierarchical FEC is enabled by default in hardware on systems running the following MDB profiles: l3-xxl, l3-xxxl, and balanced-xl.

Workaround is to configure "ip hardware fib hierarchical next-hop max-level 2" to force the use of just 2 hierarchical FEC levels in hardware. (921513)

- If a linecard is removed while traffic policy programming is in progress, the SandAegis forwarding agent may unexpectedly restart. (928128)
- Application or removal of a service-policy of type qos, on an interface may not be successful, if there is a security ACL being applied/removed on that

interface at the same time (scenarios such as configure session, configure replace, reload of the device). (934794)

- When PFC is configured on a device non-PFC traffic drops can happen for packets going to un-congested queues. (935006)
- When performing linecard OIR and the linecard being removed have an egress policy applied to an interface on the linecard and at least 1 other interface on another linecard has an egress policy applied, the SandAegis agent may restart. (943175)
- When packets are dropped at a high rate in the egress pipeline, then a small rate of packet corruption might be seen, which show up as FCS errors in the neighboring devices. Egress packets drops could happen during MTU enforcement or during split horizon related packet filtering with features such as VXLAN and L2 DCI. There is a higher chance of the issue happening at larger packet sizes. (944373)
- In an EVPN VXLAN DCI gateway switch, if initially there was both L2 and L3 DCI configuration, and then the L2 DCI configuration was removed, then subsequently all the L3 traffic destined to the remote domain VTEPs might be lost.
Workaround is to restart the SandL3Unicast agent using "agent SandL3Unicast terminate" CLI command. (960703)
- Configuring sampled flow tracking and IPv4 hardware offload with very aggressive sample rate and inactive timeout for short lived flows may cause SandFlowTracker agent or the system to restart.
The workaround is to disable IPv4 hardware offload for sampled flow tracking. (962803)
- When 3 level hierarchical FECs is enabled in hardware, slower traffic convergence may be observed when an IS-IS SR tunnel adjacency has a backup path that spans multiple FEC levels in hardware, and the primary path does down. Note that 3 level hierarchical FEC is enabled by default in hardware on systems running the following MDB profiles: l3-xxl, l3-xxxl, and balanced-xl.

Workaround is to configure "ip hardware fib hierarchical next-hop max-level 2" to force the use of just 2 hierarchical FEC levels in hardware. (964401)

- When interfaces of a breakout configuration are part of LAG with an interface traffic-policy applied have the breakout configuration removed it is possible for the SandAegis agent to get stuck in a waiting state. The workaround for this is a SandAegis agent restart. (969436)
- When a VXLAN packet is decapsulated and routed over the MLAG peer link interface, it may egress with incorrect TOS value. (978560)

- CFM PDUs that are destined to CPU may result in unexpected kernel restart. (980195)
- On a switch where egress replication mode is being used, in the presence of micro-burst or high rate of BUM or L2/L3 multicast traffic, one or more of the chips on the switch might enter a state where all subsequent replicated traffic is completely dropped. In this state, MAC learning could also be impacted. Note that even if the system is operating in ingress-only global replication mode, egress replication could be used for features such as EVPN OISM over VXLAN or VPLS.

Once a chip gets into this state, it can be recovered by either performing a soft reset using "reset platform fap full" CLI command or performing a reload of the switch.

Workaround: In the absence of above-mentioned features that force the use of egress replication, the switch can be configured to operate in global ingress-only replication mode to avoid this issue. (989455)

- Interfaces may experience intermittent link flaps with certain optics due to noise on the optics power rail. See Field Notice 94 for more information. (997423)
 SKUs Affected: DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R, DCS-7280SR3K-48YC8-F, DCS-7280SR3K-48YC8-R, DCS-7280SR3K-48YC8A-F, DCS-7280SR3K-48YC8A-R, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- Upon configuring trunk ports on ports associated with incompatible speed groups, SandFapNi agent restart, or linecard removal/insertion, the SandTunnel agent may unexpectedly restart. The agent is expect to stablize and recover. (998216)
- When forwarding agent restarts or an SSO is triggered after a LC is removed, it can result in traffic loss for flows that are egressing via IP tunnels (1008287) (1)
- Routed interfaces, untagged sub-interfaces, and interfaces with PVLAN configuration are not supported as member ports of a Port-Channel with "port-channel lacp fallback individual" configured. (1021913)
- We may see drops on lossless queues if the amount of congested traffic per ingress chip exceeds it's DRAM bandwidth. (1023171)
- IPv6 HBH packets where the destination is resolved via a Type-5 EVPN route are mistakenly dropped. (1032862)
- Some memories may not be corrected after a bit flip. The %SAND-3-INTERRUPT_OCCURRED error log will be printed continuously. (1039096)

- The SandCounters agent may unexpectedly restart repeatedly on very large systems or if there is heavy load on the supervisor's CPU. (1041848)
- A traffic-policy applied to an interface that never comes up after a linecard hot-swap may cause traffic-policy to complete stop processing configuration updates. (1044368)
- In a modular switch with both Jericho2 and Jericho2C+ line cards, if the Jericho2 line card comes up after the Jericho2C+ linecard and the route scale is high, then the SandL3Unicast agent might unexpectedly restart.

As a workaround, make sure the Jericho2C+ line card gets inserted after the Jericho2 line card. (1051065)

- SandL3Unicast agent might unexpectedly restart during LAG member expansion in the presence of EEDB HW resource overflow. (1060402)
- If the SandTunnel agent is down and the SandL3Unicast agent restarts before the SandTunnel agent comes back up, traffic destined towards VPLS pseudowires may be misforwarded or lost.
A workaround is to restart the SandL3Unicast agent again while the SandTunnel agent is up. (1071241)

- SandOam will restart continuously when CFM config is disabled and CFM packets will be wrongly forwarded to CPU.
Workaround is to disable domain config before disabling CFM config. (1072728)
- If 'show forwarding destination' is used when there isn't at least one active FAP, or all faps become inactive shortly after the command is run, the command may indicate an internal error occurred. Subsequent commands will indicate an internal error, even after there is at least one active FAP.

As a workaround, restart the SandPacketTracer agent. (1087328)

- When a VLAN translation rule is configured on a trunk interface and that interface has `no switchport mac address learning` configured. A packet may have its MAC address learned if it was match by the VLAN translation rule. (1095629)
- When L2PF profile and MAC ACL are applied on the same interface, L2PF might not work post SandAcl hitless restart
Recovery can be done by removing and reapplying the L2PF profile (1096039)
- When recovering from a Single Event Upset, one or more agents may unexpectedly restart. (1098442)
- All interfaces on a linecard remain down when the linecard is affected by a riser card power problem.

Workaround is to RMA the linecard. (1101153)

SKUs Affected: 7800R3-48CQM2-LC, 7800R3-48CQMS-LC and 7800R3K-36DM-LC

(1) This item is in two or more sections on this document

DCS-7280R and DCS-7280R2 Series

- 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)
- RACLs are not supported for VXLAN traffic ingressing on a MLAG peer-link when AlgoMatch is enabled. (485677)
- For an ACL applied to a LAG subinterface when all members of the LAG subinterface on a given chip are down, subsequent edits to any ACL on that chip may cause the default configured action to be applied to all interfaces on the chip where an ACL is present.

The default configured action is to either drop or permit if the "hardware access-list update default-result permit" is present.

To recover when this bug:

- Remove the "ip access-group" configuration from LAG subinterfaces that are down
 - Remove the LAG members that are down from the LAG that has subinterfaces with an ACL configuration (598255)
- When AlgoMatch is enabled and MPLS is a configured feature in the TCAM profile, mpls-in-ip packets (ip-protocol=137) may bypass ACLs configured on interfaces. (679923)
 - When using AlgoMatch ACLs, certain sequences of configuring a routed port with ACL applied that is later configured to be part of a LAG with an ACL applied may cause traffic to be dropped on all ports with ACLs applied:
 1. configure routed port with an ACL
 2. configure routed LAG with no ACL applied
 3. add routed port as member of routed LAG
 4. configure routed LAG with an ACL

The only way to recover from this is to restart the SandHalo forwarding agent. (800795)

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7260X, DCS-7260X3 and DCS-7300X Series

- When performing reload hitless, link flaps on BASE-T ports might be observed. (152459)
Series Affected: DCS-7050TX Series
- MAC ACLs and PACLs cannot co-exist on the same interface if QoS policy ACLs exist in the switch (186949)
- The StrataVlanTopo agent may unexpectedly restart after many linecard online insertion and removal operations. (193787)
Series Affected: DCS-7300X Series
- Ethernet65, Ethernet66 may experience extra link flap when performing fast boot. (221059)
SKUs Affected: DCS-7260QX-64-F, DCS-7260QX-64-R, DCS-7260QX-64-SSD-F and DCS-7260QX-64-SSD-R
- The Strata fabric or linecard agent may restart due to fabric link flaps leading to traffic outage. (225302)
Series Affected: 7320X and DCS-7300X Series
- DirectFlow flow in table type vfp will not match malformed Arp requests. (233922)
- When traffic is redirected using PBR and when the FIB lookup based on destination IP address indicates ARP miss, the PBR redirected traffic may be rate-policed.

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

- When the utilization of the MAC address table grows close to the full hardware capacity, the StrataL2 agent may encounter an out of memory (OOM) condition in the system, causing the agent to get restarted. (238115)
Series Affected: 7320X, DCS-7050TX, 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-7060X Series
- During an SSO failover, Strata Fabric agent may not detect hardware, resulting in a fatal error reboot with increased traffic outage. (250129)
Series Affected: DCS-7300X Series

- In an L3 EVPN setup, traffic that is sourced by a VTEP that is a pure VXLAN router (i.e., it performs no VXLAN bridging and has no hosts attached to it) might incur some loss at the receiving VTEP, where the packets are decapsulated and forwarded.

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

- BUM traffic may temporarily loop into the MLAG pair when a member interface in the peer-link port channel is brought up. (259213)
- When IPv4 uRPF exceptions are already programmed, configuring IPv6 uRPF exceptions imposes a heavy processing load on the Strata Slice agent to do a transition from IPv4 to IPv4-v6 exception mode in the hardware. Therefore, under a scaled config scenario, the Strata Slice agent may restart repeatedly.

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

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

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

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

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

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

- Egress ACL on SVI for BUM traffic will not match ACL rules for packets being sent across fabric interfaces. (284543)
SKUs Affected: DCS-7304 and DCS-7308
- Applying a large configuration containing many interface changes may result in Strata forwarding agent restart, which could cause traffic loss during the reconfiguration. (330840)
- 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)

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

- Irrespective of IGMP snooping being enabled or disabled, IGMP packets will not be flooded to remote VTEPs. (368106)

Series Affected: 7300X3 and DCS-7050X3 Series

- Link partners might see high number of uncorrected codewords when Reed-Solomon FEC is in use.

A workaround is to run 'shutdown' followed by 'no shutdown' on the affected interfaces. (370583)

Series Affected: DCS-7060X2 and DCS-7260X3 Series

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

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

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

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

- DirectFlow flows do not take precedence over conflicting router ACLs. (408734)

Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX2, DCS-7050SX3 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

- Hosts may incorrectly learn the ARP of virtual IP behind the switch MAC when the switch is reloaded or configuration is flapped. Workaround is to clear the ARP on the host to trigger a new ARP reply with the correct VARP MAC. (418626)
- Configuring more than 16K MPLS pop/swap labels will cause continuous restart of forwarding agent. (419243)
- If the number of next-hop groups configured exceeds maximum supported hardware NHGs, then STRATA-3-MPLS_MAC_DA_PROFILE_TABLE_FULL syslog message is logged and L3 forwarding agent may continuously restart.

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

- Ethernet49-54 may not link up with peers over long copper cables. Workaround is to enable autoneg on both ends of the link with 'speed auto <speed>' command. (421121)
SKUs Affected: DCS-7060SX2-48YC6-F and DCS-7060SX2-48YC6-R
- After a speed change, flows with mirroring actions on affected interfaces may not be mirrored as expected. Strata Slice agent can be restarted to restore expected behavior. (422221)
Series Affected: CCS-720XP, DCS-7050SX3 and DCS-7260CX3 Series

- The forwarding agent may continuously restart if the speed is changed from 100G to 50G when a 50G QSFP transceiver is inserted. This only affects 2-lane 50G transceivers, and does not affect 100G transceivers operating in 50G mode. A workaround is to configure the speed before inserting the transceiver. The forwarding agent restarts can be stopped by temporarily configuring the affected interfaces in 25G mode. (424990)
- On speed change or port going down, QoS rewrite stale rules may be persistent. (425445)
- With scale policy-map configuration, Strata agent may restart continuously with SIGQUIT. (430654)
- During Leaf-SSU or SSO, interfaces with 100GBASE-SRBD transceivers may experience a single link flap. (440451)
- 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)

- Aggregate storm control rate limit is not honored for broadcast and multicast traffic. (490511)
Series Affected: DCS-7060X, DCS-7060X2, DCS-7260X3 and DCS-7300X Series
- If an MBA supplicant is not assigned a dynamic VLAN when it is authenticated on a port, and if the supplicant later moves to a new port which is in unauthorized state, MBA authentication on the new port will not succeed and the port will continue to stay in unauthorized state. (491230)
Series Affected: 7300X3, AS7326, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, DCS-7050SX3, DCS-7060CX, DCS-7060SX2, DCS-7260CX and DCS-7260CX3 Series
SKUs Affected: 7304X3-FM2, 7308X3-FM2, 7320X-32C-LC, 7324X-FM, 7328X-FM, DCS-7010TX-48, DCS-7050CX3-32S, DCS-7304 and DCS-7308
- MAC loopback interfaces may not come up after hitless restart of the Strata forwarding agent with "phy link detection aggressive" configured.

Workaround is to disable "phy link detection aggressive" on the loopback interface, or "shutdown" followed by "no shutdown". (497797)

- MLAG connectivity between the peers might go down after linecard OIR. Flapping MLAG peer-link interfaces will bring the MLAG connectivity back up (500162)
- With "switchport default phone vlan <vlan-id>" configuration, DHCP offer packets tagged with the default phone vlan tag might be sent out untagged on trunk ports (501749)
- 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)
- Enabling per VLAN routing with an L3 interface that only has an IPV6 address will result in no routing of packets received on this L3 interface. Workaround is to flap the interface. (528490)
- A different set of routed interfaces may operate after reload or StrataL3 process restart if the supported scale of routed interfaces is exceeded. (533468)
SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DP-24S, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-48S, CCS-720DP-48S-2F, CCS-720DP-48S-F, CCS-720DT-24S, CCS-720DT-24S-2F, CCS-720DT-24S-F, CCS-720DT-48S-2F, CCS-720DT-48S-2R, CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8, CCS-722XPM-48ZY8-F and DCS-7010TX-48

- Changing the speed of a QSFP or SFP interface from 10G to 1G may cause a Strata forwarding agent restart under some circumstances. This will cause an interruption in traffic on all ports. Workaround is to change from 10G to another speed before changing to 1G. (550269)
Series Affected: 7300X3, CCS-720XP, DCS-7050SX3 and DCS-7260CX3 Series
- TCAM resource exhaustion can lead to VXLAN forwarding to fail.
Workaround: Reduce TCAM resource utilization to allow VXLAN entries in TCAM by unconfiguring other features that occupy TCAM. (561231)
- When "show platform trident vxlan multihoming groups non-peering" is executed, StrataL3 agent may restart unexpectedly. (561420)
- Fabric interface on a linecard can sometimes end up in a state where it can discard all packets at egress.
When this condition happens, the "show platform trident fabric counters queue detail" command will show that the drops increment while there is no increase in packets or bytes transmitted on the problematic fabric interface. To recover from the problem, it is necessary to restart the linecard forwarding plane agent by running "platform trident <linecard> reset". The recovery process can cause traffic disruption across several ports in the switch. (562860)
SKUs Affected: DCS-7304 and DCS-7308
- Committing security access lists through CVP may restart forwarding plane slice agent. Post restart forwarding plane functionality will continue to work as is.
Workaround is to use "hardware access-list update default-result permit" CLI. (571289)
- StrataTcam process may restart when we the primary Switchcard is hotswapped. System should be unaffected from the restart of StrataTcam process.

There is no workaround. (584258)
SKUs Affected: CCS-755-CH and CCS-758-CH
- When many VLANs and/or LAGs are configured, port channels may flap during SSU. (592390)
- PIM over GRE interface does not support over 1K multicast routes (605368)
- Performing SSO with increased route scale while Segment Security is configured may result in a fatal error reboot and extended traffic outage. (623639)
- After an interface flap, multicast traffic may not forward correctly in data plane.
Workaround is to remove and re-apply multicast configuration. (627318)

Series Affected: CCS-710P and DCS-7010TX Series

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- StrataL3 forwarding agent may restart unexpectedly when many L3 interfaces are removed and then added again. (636201)
- Migrating from static VXLAN to EVPN VXLAN configuration may result in stale remote MAC entries in hardware leading to black holing of traffic destined to those MAC addresses.

Workaround is to clear the MAC address table entries for those VLANs. (644589)

- If StrataL3 is restarted when there are insufficient hardware resources to add a virtual port for a VXLAN remote VTEP as indicated by the STRATA-3-VIRTUALPORT_RESOURCE_FULL syslog, StrataL3 agent may not become warm resulting continuous traffic drops.

Work around is to reduce the number of VXLAN remote VTEPs.

Workaround: Reduce VTEP scale so that all VTEPs have an allocated virtual port before restarting StrataL3 to recover. (646694)

- Adding and removing "no switchport mac address learning" command while MAC address flaps between local and remote can trigger stale MAC entry left in MAC address table even after clearing all MAC addresses. (664960)
- When ECN is configured globally, non-ECT packets could be dropped if the total buffer occupancy exceeds the global threshold, even if there is sufficient memory available in reserved space. (666378)
- When "hardware counter feature vtep encap" CLI is configured, the deletion of VTEPs can trigger an unexpected restart of the StrataL3 agent. (678643)
- Use of "no mac address learning" command with VXLAN configuration may result in mis forwarding of the VXLAN encapsulated traffic received.

Workaround is to remove "no mac address learning" configuration on VXLAN VLANs. (681651)

Series Affected: DCS-7050QX, DCS-7050TX, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7300X Series

- An EVPN MLAG router configured with "multi VTEP MLAG" feature may not install hosts attached to a remote EVPN Active-Active Multihoming Ethernet Segment. As a result, the bridging and routed traffic towards those remote EVPN A-A multi-homed hosts could experience unwanted flooding or traffic loss. (684557)

- After a Leaf Smart Software Upgrade (Leaf SSU), Strata agent might restart unexpectedly, resulting in increased traffic outage. In some cases, this might also cause an SSU stage timeout without Strata agent restart, resulting in a fatal error reboot. (689713)

SKUs Affected: DCS-7060DX4-32 and DCS-7060PX4-32

- Removing an SFP BASE-T transceiver from Ethernet33 or Ethernet34 may cause the forwarding agent to restart and a brief traffic outage to occur. (715516)

Series Affected: DCS-7050X4 and DCS-7060X4 Series

- During a link toggle or speed change, Strata forwarding plane agent may unexpectedly restart due to "OVERFLOW IRQ from ISEC_PDF_STATUS or ISEC_IDF_STATUS".

Switch is expected to get back to working state after the restart. (717855)

Series Affected: CCS-750 Series

- If MLAG configuration is applied when there is congestion, MLAG stays in inactive state and StrataVlanTopo might restart. When the congestion is stopped MLAG becomes active (727116)

- StrataLag agent will continuously restart unexpectedly if there are more LAGs configured (including LAGs with no members) than the hardware can support.

The LAG hardware limits can be seen by executing CLI 'show port-channel limits', field 'Max port-channels per group'.

Workaround:

- remove configured LAGs until the number of LAGs is not exceeding the hardware capacity. (732733)

- An interface configured to be a member of a LAG may fail to be programmed in the hardware.

The CLI 'show port-channel detailed' will show the status of the interface as "Waiting for hardware to program port to RX...".

Workaround:

Remove and add the affected member interface from/to the LAG (738580)

- Transitioning to data plane learning from EVPN control plane may result in remote MAC addresses not being learnt.

Workaround is to flap the VXLAN interface. (748461)

- 5GBASE-T ports may always report linkup regardless of the actual link state, when configured with forced 100M speed using CLI command `speed 100full`.

Workaround is to use BASE-T auto-negotiation with capable link partners

instead of forcing the speed to 100M. (759421)

SKUs Affected: CCS-720XP-96ZC2

- MAC initialization failure may cause Strata forwarding plane agent to unexpectedly restart during stateful switchover, and this may cause stateful switchover to fail, and reboot the switch. (776821)
Series Affected: 7300X3 Series
- IP ACL configured on interface receiving VXLAN traffic encapsulated to VARP VTEP IP may result in drops.
Workaround is to configure port range rule in the ACL or apply the following command: `platform trident tcam port-acl vfp disabled`. (784384)
- Looping or double delivery of L2 protocol packets received on the peer-link may be observed in an active-active MLAG setup when L2 Protocol Forwarding is configured on the peer-link. (784857)
- When an interface having "switchport trunk private-vlan secondary" configuration is removed from a port-channel we may observe incorrect tagging on packets egressing the interface.
A workaround is to remove and add the "switchport trunk private-vlan secondary" configuration. (786324)
- Groups of four adjacent BASE-T ports in the range Et1-Et24 share a common hardware resource; for example, Et1-Et4 and Et5-Et8 and so on. Due to an issue with the shared hardware resource, all 4 ports may fail to link up, even if the link partners report a linkup.

On EOS versions prior to 4.29.0F, on fixed systems, the workaround is to run `platform trident reset`; on modular systems, the workaround is to run `platform trident <linecard> reset` on the affected linecard. On EOS versions 4.29.0F and onwards, the workaround is to run `reset platform trident phy interface INTF pcs` on one of the affected interfaces. (792506)

SKUs Affected: CCS-720DP-48S, CCS-722XPM-48Y4, DCS-7010TX-48 and DCS-7010TX-48-DC

- IPv6 Underlay is not supported. Receiving EVPN routes from IPV6 VTEPs may result in continuous restart of L3 forwarding agent. (794373)
Series Affected: 7320X, 7388, DCS-7050QX, DCS-7050QX2, DCS-7050SX, DCS-7050SX2, DCS-7050TX, DCS-7050TX2, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series
- When Reducing number of members in ECMP groups while running traffic with DLB enabled can cause some packets drop.
Workaround is to restart the traffic or disable/enable DLB. (797519)
SKUs Affected: 7358X4-SC, DCS-7050DX4-32S, DCS-7050DX4-32S-F, DCS-7050PX4-32S and DCS-7050PX4-32S-F

- Active supervisor removal may cause the PlxPcie-system to unexpectedly restart and, in some circumstances, may prevent the uplink ports on the supervisor card from being detected.

Workarounds: reinsert the card again or reload the device. (807283)

SKUs Affected: CCS-755-CH and CCS-758-CH

- VLAN translation may stop working on a port when:
 - the port was removed from a LAG
 - that LAG has the same VLAN translation configured as the port

Workaround:

- restart StrataL2 agent (809637)

- If there are interfaces with transceiver inserted but repeatedly flapping or unable to link up due to poor signal, high CPU usage of the Strata slice agent may be experienced.

To workaround this, shut down the interfaces or remove the transceivers. (811800)

- If a MAC flaps between a port-security protect mode interface and another interface rapidly, it may be displayed on the wrong interface in 'show mac address-table' and might not be counted in number of allowed addresses on the port-security interface (814589)
- When the switch is configured with lot of VXLAN VLANs and trunk ports, changing the VXLAN loopback interface's IP or flapping the loopback interface may lead to VXLAN DIP termination entry not programmed in hardware resulting in dropping all the VXLAN encapsulated traffic.

Workaround: Reload of the device or shutdown edge interfaces before flapping VXLAN loopback interface. (819174)

Series Affected: DCS-7050QX, DCS-7050SX, DCS-7050TX, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series

- Pause Flow control might not work after Autoneg resolution.

Workaround :

Reconfigure Pause Flow control.

1. "no flowcontrol send"
2. "flowcontrol send on" (820102)

Series Affected: 7388 Series

SKUs Affected: 7358X4-SC, 7388X5-SC and DCS-7060DX5-64

- If VLAN translation is configured on a port-security protect mode port and the StrataL2 agent is restarted, some MACs learned on that port may be able to move and age.

A workaround is to remove and re-add the port-security configuration for that port. (821390)

- If many continuous MAC moves occur in a short period of time, the device may not be accessible via SSH while the MAC moves continue. (845283)
- When performing configure-replace with traffic-policy which exhausts TCAM, the Strata agent may restart. Workaround would be to lower the scale of the traffic-policy such that TCAM is not exhausted. (848910)
Series Affected: 7300X3, CCS-720DF, CCS-720DP, CCS-720XP, CCS-722XPM, DCS-7010TX, DCS-7050CX3, DCS-7050SX and DCS-7050TX Series
SKUs Affected: 7358X4-SC, DCS-7304 and DCS-7308
- A chip reset may occur as a result of a fatal parity error reported in the MMU egress buffer block if the system is configured to exceed the hardware replication limit of 1024 replications per packet. (850455)
Series Affected: 7358, 7368, 7388, DCS-7050X4, DCS-7060X4 and DCS-7060X5 Series
- If multiple devices behind a phone/hub/switch connected to an Arista switch are authenticated via Dot1x MAC Based Authentication and assigned different VLANs, some of these devices may not be able to move to other interfaces that have Dot1x MAC Based Authentication enabled.

A workaround is to enable per-interface MAC-based VLAN Assignment using the "mac based vlan assignment scope interface" command. (858856)

- Support for the 64-bit image binary of EOS on this platform will be removed in EOS-4.32.0F. Customers should plan to migrate to a 32-bit image binary of EOS if presently deployed with the 64-bit image binary of EOS. Please refer to Field Notice 78 for additional details. (880630)
SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DF-48Y-2F, CCS-720DF-48Y-F, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-48S-2F, CCS-720DP-48S-F, CCS-720DT-24S-2F, CCS-720DT-24S-F, CCS-720XP-24Y6-F, CCS-720XP-24Y6-R, CCS-720XP-24ZY4-F, CCS-720XP-48Y6-F, CCS-720XP-48Y6-R, CCS-720XP-48ZC2-F, CCS-720XP-96ZC2-F, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8-F, DCS-7010TX-48-DC-F, DCS-7010TX-48-DC-R, DCS-7010TX-48-F and DCS-7010TX-48-R
- Configuring and removing static MAC address entries for MAC addresses that have already been learnt may result in high CPU usage and unexpected restart of StrataL2 forwarding agent. (884326)
- When a packet is received with IP Destination address 0.0.0.0 and multicast MAC address 01:00:5e:00:00:00, the packet is either flooded to the VLAN or sent to the multicast router ports in the VLAN.
Workaround is to configure an ACL to drop such packets. (885556)

- Configuring and removing static MAC address entries for MAC addresses that have already been learnt may result in high CPU usage by the StrataL2 agent. (886655)
- ARP packets sent to VRRP MAC don't get copied to control plane if MACsec is enabled on the interface.
Workaround is to disable MACsec on the interface by running "no mac security profile <profile>". (921378)
SKUs Affected: CCS-722XPM-48Y4-F and CCS-722XPM-48ZY8-F
- On speed change, WRED entries may not be removed for some of the ports that become inactive. This may reduce the total amount of available hardware resources. (923751)
Series Affected: 7368 Series
SKUs Affected: 7300X3-48TC4-LC, 7368X4-SC, 7388-SUP-D, CCS-720DP-24S-2F, CCS-722XPM-48ZY8-F, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050CX3-32S-SSD-F and DCS-7050CX3-32S-SSD-R
- With IP routing disabled, self IP packets are not processed when Mirror On Drop feature is enabled.
Workaround: Enable IP routing (943114)
- If a single TCAM based feature which uses statistics tries to configure platform maximum number of entries in egress TCAM then Strata forwarding agent will restart continuously.

Workaround is to remove the offending configuration and reduce the scale by 1. (946410)
SKUs Affected: 7388X5-SC, DCS-7060DX5-64-F, DCS-7060DX5-64S-F, DCS-7060DX5-64S-R and DCS-7060PX5-64E-F
- During a Leaf Smart Software Upgrade (Leaf SSU), Strata agent may restart unexpectedly, resulting in a fatal error reboot. (959348)
- Security Advisory for 0108 for CVE-2024-6437. For more information, see <https://www.arista.com/en/support/advisories-notices/security-advisory/20689-security-advisory-0108> (962149) (1)
- A netdevice refcount that is held on an unresponsive Linecard may prevent Strata from shutting down, resulting in a message similar to:
'unregister_netdevice: waiting for et3_49_1 to become free. Usage count = <n>'. The affected supervisor should be rebooted at the operator's earliest convenience. (967727)
- StrataMirror agent could enter an unstable state if it reloads with a session that has an ingress VLAN and an egress port as sources.
Workaround: remove ingress VLAN source, egress port source or the session completely, then reconfigure the session. (969502)

- DHCP traffic egressing on a trunk port from a MAC-Based VLAN Assignment phone have its VLAN tag removed. (976419)
- When uRPF exceptions are configured, forwarding agent may restart unexpectedly during switch card switchover or port events. (980046)
- While ASU upgrade if link flap happens on DLB ECMP member interface then DLB may not consider that member resulting in not sending traffic on that ECMP member.
Workaround is to do shut/no shut on that interface. (980295)
- "IP address virtual" may not take effect for SVI on Strata devices for IPv4 if the SVI is already up. The workaround is to flap the SVI interface. (987212)
- When a forwarding ASIC gets reset, the device may unexpectedly reload. (990910) (1)
- When the Strata agent, the AsicResourceMgr agent, and the StrataLag agent are restarted, control frame may not be sent to a LAG.

Remove and add the LAG members back when the agents are warm will restore the traffic. (995814)

- Forwarding agent may restart multiple times due to an invalid interface configuration (e.g. incompatible speed configuration).

The workaround is to restart the forwarding agent by using the `platform trident reset` command after the invalid interface configuration is reverted. Please note that this may disrupt traffic on all front panel interfaces. (996281)

- Smart System Upgrade(SSU) may time out and this will result in a cold boot. (1002602)
Series Affected: CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM and DCS-7010TX Series

- If Strata agent is restarted right after applying `shutdown` configuration on interfaces, it may begin restarting repeatedly.

The workaround is to configure "no shutdown" on all the shutdown interfaces. (1004640)

- When configuring "shutdown" on multiple L3 sub-interfaces, followed by configuring "no shutdown" and then applying an access-list (ACL) to one of the L3 sub-interfaces, an ACL timeout error may occur. Further attempts to re-apply any ACL on any interface may continue to time out. To workaround the issue, restart the Strata slice agent by running "[no] agent Strata-FixedSystem shut" (1015426)

- The following logs may be emitted occasionally.

```
PtpTimeSync: %HARDWARE-3-TIMESYNC_ERROR: Linecard<...> has experienced a hardware error (slave frequency check error) with internal time synchronization
```

```
PtpTimeSync: %HARDWARE-3-TIMESYNC_ERROR: Linecard<...> has experienced a hardware error (slave synchronization convergence timeout) with internal time synchronization
```

If PTP is not enabled these have no functional impact and can be ignored. If PTP is enabled, PTP may briefly lose synchronization when this event occurs. (1024866)

SKUs Affected: CCS-755-CH and CCS-758-CH

- Scaled ACL configuration may cause traffic drop of few seconds during SSU upgrade. Workaround is to use CLI "hardware access-list update default-result permit" to disable port blocking while applying ACL (1037945) (1)
- Persistent port-security MAC addresses will remain on an interface after it is moved into a channel group. The MAC addresses may not move or age out. If the interface has port-security protect configured then additional new MACs may not learn on the port-channel. A workaround to clear these addresses is to use `clear port-security interface <intf>` for the interface with port-security configured that was moved into the channel-group. (1051152)
- On a system having DLB configured, if state of interfaces changes during software upgrade via SSU, those interface may not be enabled to carry DLB traffic.

Workaround is to do shut/no shut on that interface. (1061194)

- BASE-T ports may blackhole, drop or receive/transmit traffic with errors (e.g. CRC errors) if the link flaps and a different linkup speed is negotiated with the peer, while the system is in the process of a Leaf Smart Software Upgrade (Leaf SSU). The issue will resolve itself after the forwarding agent warms up.

Workaround is to advertise a single speed during auto-negotiation using `speed auto SPEED` CLI command. (1070944)

- IPV6 multicast packets will be forwarded in the data plane and not reach PTP agent if "ptp forward unicast" is configured with "ptp mode boundary" or "ptp mode boundary one-step". (1086869)
- Secure MAC addresses will not be cleared if port-security configuration is removed while the StrataL2 agent is down. 'clear port-security interface <intf>' and 'clear mac address-table dynamic' may also fail to clear these addresses. A workaround to remove these MAC addresses when StrataL2 agent

is running again is to reapply the port-security configuration to the interface, then use 'clear port-security interface <intf>', then remove the configuration. (1098136)

- If "switchport port-security violation protect" is configured on an interface, and the number of MACs on the port is continuously moving above and below the port-security MAC-address maximum, a large number of ACLs will be created and removed, which may result in the Strata agent exiting with the SOCKET_CLOSE_REMOTE syslog.

A workaround is to disable MAC moves for port-security ports to reduce the number of MAC events using the "no switchport port-security mac-address moveable" command. (1105452)

- The Strata agent may have high CPU utilization for an extended period of time when subject to several continuous MAC moves between port-security protect interfaces. (1111717)
- Using a config session, when more than one ACL is applied to a SVI/L3 sub-interface along with an ACL on a different interface, any ingress/egress traffic on the SVI/L3 sub-interface will be dropped. Workaround is to reapply the ACL on the SVI/L3 sub-interface. "hardware access-list update default-result permit" could be used to prevent the problem from occurring. (1122735)

CCS-750 Series

- Switchcard switchover may disrupt PTP time synchronization and cause PTP to go into holdover state temporarily. PTP will recover and exit holdover state when time synchronization is re-acquired without any intervention. (728467)
- Ingress PTP management packets will not be dropped as expected when both 'ptp management drop' and 'ptp forward-unicast' are configured. (936957)
- When a switchover occurs, the Strata agent may restart unexpectedly (965328)
SKUs Affected: CCS-755-CH and CCS-758-CH
- Control traffic can be impacted when receiving 100% linerate traffic on an uplink port.
This can be avoided by reducing the incoming rate to 99.98% or less. (1064421)
SKUs Affected: CCS-750-SUP100 and CCS-750-SUP25

7300X3, 7320X, DCS-7060X and DCS-7260X Series

- If an interface is configured to use auto-negotiation with a 1000BASE-SX or a 1000BASE-LX transceiver and is shutdown, the interface will report continuous link flaps.

Workaround is to configure forced speed on the interface. (427843)

- Packet buffer memory corruption can result in packets getting discarded without visibility in counters.
When this condition occurs, the buffer usage reported by the "show platform trident mmu queue status" command exceeds the total number of buffers reported by the "show platform trident mmu buffers" command.
A workaround is to run the "platform trident reset" command, which results in forwarding agent restart causing the links of the system to flap momentarily. (553299)
- If there is a loop in the network we could see a lot of MAC moves, leading to high CPU utilization by the StrataL2 agent and causing problems with learning and aging. (845250) (1)
SKUs Affected: 7300X3-32C-LC, CCS-710P-12, CCS-720DF-48Y-2F, CCS-720DP-24S-2F, CCS-720DP-48S-F, CCS-720XP-24Y6-R, CCS-722XPM-48Y4-F, DCS-7060CX-32C-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7304 and DCS-7308
- Changing a static MAC to a drop MAC may not take effect in the hardware. The workaround is to remove the existing MAC first, then apply the drop MAC configuration. (938220) (1)
- If a large amount of flows are continuously being re-learned, Strata slice agent may restart. (993211) (1)
- Continuous UFT config changes can result in MAC addresses not being learned. Workaround is to restart StrataL2 (1004102)
Series Affected: DCS-7060X Series
- Interfaces with 100BASE-FX transceivers may flap during Leaf Smart Software Upgrade (Leaf SSU). (1024404) (1)
- Configuring link-debounce prior the Strata forwarding agent hitless restart or SSU, and then subsequently removing link-debounce configuration on Ethernet33 and Ethernet34, one way link up may be experienced.

If the Strata forwarding agent restarts hitlessly or SSU is performed, and then link-debounce is configured on Ethernet33 and Ethernet34, the configuration will not prevent the peer from going down during the debounce period..

The workaround is to perform forwarding chip reset using the `platform trident reset` CLI command. This will lead to a flap of all interfaces and a brief traffic outage. (1024728) (1)

- Traffic to a multihomed MAC destination may see packet drops during VTEP reachability change when putting spine devices into maintenance mode.

There is no workaround for this issue. (1044420) (1)

- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)

DCS-7060X2 Series

- If there is a loop in the network we could see a lot of MAC moves, leading to high CPU utilization by the StrataL2 agent and causing problems with learning and aging. (845250) (1)
SKUs Affected: 7300X3-32C-LC, CCS-710P-12, CCS-720DF-48Y-2F, CCS-720DP-24S-2F, CCS-720DP-48S-F, CCS-720XP-24Y6-R, CCS-722XPM-48Y4-F, DCS-7060CX-32C-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7304 and DCS-7308
- Changing a static MAC to a drop MAC may not take effect in the hardware. The workaround is to remove the existing MAC first, then apply the drop MAC configuration. (938220) (1)
- If a large amount of flows are continuously being re-learned, Strata slice agent may restart. (993211) (1)
- Interfaces with 100BASE-FX transceivers may flap during Leaf Smart Software Upgrade (Leaf SSU). (1024404) (1)
- Configuring link-debounce prior the Strata forwarding agent hitless restart or SSU, and then subsequently removing link-debounce configuration on Ethernet33 and Ethernet34, one way link up may be experienced.

If the Strata forwarding agent restarts hitlessly or SSU is performed, and then link-debounce is configured on Ethernet33 and Ethernet34, the configuration will not prevent the peer from going down during the debounce period..

The workaround is to perform forwarding chip reset using the `platform trident reset` CLI command. This will lead to a flap of all interfaces and a brief traffic outage. (1024728) (1)

- Traffic to a multihomed MAC destination may see packet drops during VTEP reachability change when putting spine devices into maintenance mode.

There is no workaround for this issue. (1044420) (1)

- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-

IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)

DCS-7260X3 Series

- Forwarding agent may restart when a 40GBASE-SRBD transceiver is inserted after the hitless reload. To work around, configure "speed forced 40gfull" on the desired interface prior to inserting the transceiver. (417311)
- After a reload, the forwarding agents may fail to initialize and the Linux dmesg will be filled with "link down" log messages.

Workaround is to reload the switch. (526172)

SKUs Affected: DCS-7260CX3-64, DCS-7260CX3-64E and DCS-7260CX3-64LQ

- If there is a loop in the network we could see a lot of MAC moves, leading to high CPU utilization by the StrataL2 agent and causing problems with learning and aging. (845250) (1)
SKUs Affected: 7300X3-32C-LC, CCS-710P-12, CCS-720DF-48Y-2F, CCS-720DP-24S-2F, CCS-720DP-48S-F, CCS-720XP-24Y6-R, CCS-722XPM-48Y4-F, DCS-7060CX-32C-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7304 and DCS-7308
- Changing a static MAC to a drop MAC may not take effect in the hardware. The workaround is to remove the existing MAC first, then apply the drop MAC configuration. (938220) (1)
- Symptom: On restarting StrataL3 agent, a traffic outage of around 5-10 seconds is observed.

Trigger: The issue is seen only when VXLAN Routing is enabled and interface is configured as Recirc-channel interface. On restarting StrataL3 agent, Routing gets disabled momentarily and Router MAC entries pointing to Recirc-channel interface go for a toss, leading to the traffic outage.

Workaround: The traffic converges automatically in 5-10 seconds (951137)

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

- If a large amount of flows are continuously being re-learned, Strata slice agent may restart. (993211) (1)
- SSU upgrade sequences starting from EOS versions:
 - 4.20.5 and later 4.20.X
 - 4.21.X earlier than the affected releaseto affected releases when using 25G/50G/100G speeds on optical transceivers may result in these interfaces to fail to link up.

The workaround is to change the speed of the affected interfaces to a different one, and then to change it back by using the "speed" command. For example, if an affected interface is configured with 100G speed, the issue can be resolved by changing the speed first to 25G and then changing it back to 100G, via "speed 25g" and then "speed 100g" commands.

Note that SSU upgrade sequences that stop at some other intermediate release(s) prior to upgrading to an affected release will still encounter the issue upon upgrade to the affected release. (1008955)

- Traffic to a multihomed MAC destination may see packet drops during VTEP reachability change when putting spine devices into maintenance mode.

There is no workaround for this issue. (1044420) (1)

- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)

7368 and DCS-7060X4 Series

- 802.3X pause is not supported (348758)
- If speed of a member interface of a port channel belonging to a VLAN is changed, then it can cause packet drops on that interface. Workaround is to unconfigure and reconfigure that port channel or kill L3 agent. (380019)
Series Affected: 7368 Series
SKUs Affected: 7368-16C, 7368-16S, 7368-4D and 7368-4P
- 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)

- Speed change on a port-channel member can cause a flap in BFD and BGP sessions for the entire port-channel (520638)
- StrataMmu agent has high CPU usage when PFC watchdog is configured with a low polling interval, such as 0.010. (636394) (1)
- If an interface is shutdown, static MACs on that interface may not be counted in number of allowed addresses for port-security (862079) (1)
- A single-event upset (SEU) in the IFP_TCAM can lead to forwarding agent restart, after which the switch forwards traffic normally. (917956)

- Changing a static MAC to a drop MAC may not take effect in the hardware. The workaround is to remove the existing MAC first, then apply the drop MAC configuration. (938220) (1)
- "show platform trident l3 shadow next-hops ifp raw" may result in unexpected restart of StrataL3 forwarding agent. (954273) (1)
- When the multicast route is pointing to a large replication set, the hardware replication table may not be programmed correctly resulting in dropping the multicast data traffic destined to the route. (966098)
- If an interface is configured with "phy link training", then a SSU reload may result in a fatal error reboot that will cause extended traffic outage. (983323) (1)
- If a large amount of flows are continuously being re-learned, Strata slice agent may restart. (993211) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)

DCS-7060X5 Series

- Interfaces with copper (CR) transceivers may require auto-negotiation to reliably link up at all rates except 10G and 40G. Auto-negotiation can be enabled using the "speed auto <speed>" interface configuration command. (736210)
SKUs Affected: DCS-7060DX5-64 and DCS-7060PX5-64E
- Some interfaces may flap intermittently due to jitter (1013413) (1)
SKUs Affected: DCS-7060DX5-64 and DCS-7060PX5-64E

7388 and DCS-7060X5 Series

- If an interface at 200g-4 is down or errdisabled prior to ASU reload, it may fail to come up after ASU reload completes.

Workaround is to shutdown the interface prior, or perform shutdown followed by no shutdown post-reload. (1064374)

CCS-710P, CCS-720DP and CCS-720DT Series

- Scaled IP Locking configuration may result in an unexpected restart of Strata forwarding agent. (678926)

SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DP-24S, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DT-24S, CCS-720DT-24S-2F and CCS-720DT-24S-F

- If there is a loop in the network we could see a lot of MAC moves, leading to high CPU utilization by the StrataL2 agent and causing problems with learning and aging. (845250) (1)

SKUs Affected: 7300X3-32C-LC, CCS-710P-12, CCS-720DF-48Y-2F, CCS-720DP-24S-2F, CCS-720DP-48S-F, CCS-720XP-24Y6-R, CCS-722XPM-48Y4-F, DCS-7060CX-32C-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7304 and DCS-7308

- On the relevant platforms with a large scale of VXLAN remote ARP entries, when SSU is performed, there could be up to 30 seconds of traffic loss for bidirectional VXLAN routed flows between local host and remote VXLAN hosts. (1002874) (1)
- An SEU in the L3_ENTRY_LP table may restart the forwarding agents resulting in momentary traffic loss. (1009663) (1)
- Port Security and Dot1x authenticated MAC address entries may forward traffic incorrectly after all interfaces in a VLAN are flapped or when MSTP is flapped. Workaround is to remove the MAC addresses and allow them to get relearned and/or get authenticated. (1036032) (1)
- In a scenario with a large number of VXLAN VLANs and VTEPs, if the flood VTEP list experiences churn—either due to flood VTEP reachability issues or VTEPs rebooting—there may be exhaustion of the IP multicast group (IPMC) and replication group (ReplGroup) lists. This situation can result in VXLAN VLAN broadcast, unknown unicast, and multicasting (BUM) traffic being blackholed. A workaround for this issue is to flap the VXLAN interface. (1061846) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- An SEU (Single Event Upset) in the L2_ENTRY_LP, VLAN_XLATE_1_LP, VLAN_XLATE_2_LP, EGR_VLAN_XLATE_1_LP, EGR_VLAN_XLATE_2_LP, MPLS_ENTRY_LP, or EXACT_MATCH_LP tables may result in an unexpected forwarding agent restart causing a brief traffic disruption. (1134953) (1)

CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series

- Egress mirroring does not encrypt mirrored packets when MACsec is enabled on the destination port. (769793)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- If there is a loop in the network we could see a lot of MAC moves, leading to high CPU utilization by the StrataL2 agent and causing problems with learning and aging. (845250) (1)
 SKUs Affected: 7300X3-32C-LC, CCS-710P-12, CCS-720DF-48Y-2F, CCS-720DP-24S-2F, CCS-720DP-48S-F, CCS-720XP-24Y6-R, CCS-722XPM-48Y4-F, DCS-7060CX-32C-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7304 and DCS-7308
- DirectFlow is not supported on MACsec enabled interfaces (861705)
 Series Affected: CCS-722XPM Series
 SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- ARP and IPv6 neighbours may be in unresolved state over VXLAN tunnels when MACsec is enabled on core facing interfaces. The workaround is to disable MACsec on core facing interfaces. (897240)
 Series Affected: CCS-722XPM Series
- StrataFlowTracker forwarding agent (StftAgent) may restart unexpectedly, when "flow tracker hardware" is turned on and off in a short interval interval. (922896) (1)
- After an SSU reload, the system may encounter an error which prevents SSU from completing. This can lead to extended traffic loss. A cold reboot is required to recover the system. (954577)
 SKUs Affected: DCS-7010TX-48, DCS-7010TX-48-DC, DCS-7010TX-48-DC-F, DCS-7010TX-48-DC-R, DCS-7010TX-48-F and DCS-7010TX-48-R
- On the relevant platforms with a large scale of VXLAN remote ARP entries, when SSU is performed, there could be up to 30 seconds of traffic loss for bidirectional VXLAN routed flows between local host and remote VXLAN hosts. (1002874) (1)
- When MACsec is configured on a VXLAN core interface, broadcast, unknown-unicast and multicast packets may get dropped after forwarding plane agent restart.
 The workaround is to toggle MACsec on the VXLAN core interface. (1007386)
 SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- An SEU in the L3_ENTRY_LP table may restart the forwarding agents resulting in momentary traffic loss. (1009663) (1)
- MACsec decrypted frames shorter than 64 bytes will be dropped upon reception. (1035641)
 Series Affected: CCS-722XPM Series
- Port Security and Dot1x authenticated MAC address entries may forward traffic incorrectly after all interfaces in a VLAN are flapped or when MSTP is

flapped. Workaround is to remove the MAC addresses and allow them to get relearned and/or get authenticated. (1036032) (1)

- In a scenario with a large number of VXLAN VLANs and VTEPs, if the flood VTEP list experiences churn—either due to flood VTEP reachability issues or VTEPs rebooting—there may be exhaustion of the IP multicast group (IPMC) and replication group (ReplGroup) lists. This situation can result in VXLAN VLAN broadcast, unknown unicast, and multicasting (BUM) traffic being blackholed. A workaround for this issue is to flap the VXLAN interface. (1061846) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- An SEU (Single Event Upset) in the L2_ENTRY_LP, VLAN_XLATE_1_LP, VLAN_XLATE_2_LP, EGR_VLAN_XLATE_1_LP, EGR_VLAN_XLATE_2_LP, MPLS_ENTRY_LP, or EXACT_MATCH_LP tables may result in an unexpected forwarding agent restart causing a brief traffic disruption. (1134953) (1)

CCS-720DF and CCS-720XP Series

- Setup with only Storm-control and scaled IP Locking configuration may cause slice agent to restart continuously. (599935)
- A single event upset (SEU) in a TCAM after a SSU upgrade may not be detected. This may cause traffic to be misclassified. A forwarding agent restart is required to forward traffic normally. (621616)
- If there is a loop in the network we could see a lot of MAC moves, leading to high CPU utilization by the StrataL2 agent and causing problems with learning and aging. (845250) (1)
SKUs Affected: 7300X3-32C-LC, CCS-710P-12, CCS-720DF-48Y-2F, CCS-720DP-24S-2F, CCS-720DP-48S-F, CCS-720XP-24Y6-R, CCS-722XPM-48Y4-F, DCS-7060CX-32C-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7304 and DCS-7308
- StrataFlowTracker forwarding agent (StftAgent) may restart unexpectedly, when "flow tracker hardware" is turned on and off in a short interval. (922896) (1)
- A hitful restart of the PhyIsland agent in the presence of ingress traffic on ports Ethernet 1-40 may cause an unexpected restart of the Strata forwarding agent. This leads to a flap of all ports and a brief traffic outage. (951318)
Series Affected: CCS-720DF Series

- On the relevant platforms with a large scale of VXLAN remote ARP entries, when SSU is performed, there could be up to 30 seconds of traffic loss for bidirectional VXLAN routed flows between local host and remote VXLAN hosts. (1002874) (1)
- An SEU in the L3_ENTRY_LP table may restart the forwarding agents resulting in momentary traffic loss. (1009663) (1)
- With EVPN multicast and `platform trident forwarding-table partition flexible exact-match <> l2-shared <> l3-shared <>` configured for supporting underlay IIF switch cases like RPT-SPT switchover, underlay ECMP link switch, etc., we may see BessMgr process spinning at around 90-100% if there another feature using UDF is configured. Work-around is to remove the other feature which uses UDF. (1016232) (1)
- Exact match table lookup is not supported with IPV6 underlay multicast enabled. As a result we see duplicate packets when we enabled platform trident forwarding-table partition flexible exact-match configuration. (1018543) (1)
- In a scenario with a large number of VXLAN VLANs and VTEPs, if the flood VTEP list experiences churn—either due to flood VTEP reachability issues or VTEPs rebooting—there may be exhaustion of the IP multicast group (IPMC) and replication group (ReplGroup) lists. This situation can result in VXLAN VLAN broadcast, unknown unicast, and multicasting (BUM) traffic being blackholed. A workaround for this issue is to flap the VXLAN interface. (1061846) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- An SEU (Single Event Upset) in the L2_ENTRY_LP, VLAN_XLATE_1_LP, VLAN_XLATE_2_LP, EGR_VLAN_XLATE_1_LP, EGR_VLAN_XLATE_2_LP, MPLS_ENTRY_LP, or EXACT_MATCH_LP tables may result in an unexpected forwarding agent restart causing a brief traffic disruption. (1134953) (1)

CCS-750 Series

- When report flooding is disabled in IGMP Snooping configuration, report packets may still be flooded on port-channel member interfaces.

The workaround is to disable and re-enable IGMP snooping. (534531)

- The StrataCes Agent may unexpectedly restart when a linecard is hot swapped. (568170)

- After a hitless restart of the Strata-Switchcard agent, configuring global ECN fails with "ECN Profiles Exhausted message".

To work around the issue, restart "Strata-Switchcard" agents using the following commands:

```
platform trident Strata-Switchcard1 reset
```

```
platform trident Strata-Switchcard2 reset
```

The operation is hitful and temporarily disrupts traffic across all the interfaces of the switch. (618006)

- Linecard slice agents may restart after supervisor switchover. The system recovers normally. There is no workaround. (623381)
- SSO with L3 VXLAN may lead to extended outage greater than 1 second. (637602)
- Inserting secondary switch card may result in unexpected restart of StrataL3 forwarding agent. (639631)
- An SEU (Single Event Upset) in the L3_ENTRY_LP table results in an unexpected forwarding agent restart causing a brief traffic disruption. (641120) (1)
- When packets are flooded on a VLAN and the ingress interface is a downlink, the ingress interface will have lower egress bandwidth for other traffic flows even though the flooded traffic does not egress the ingress interface. (649424)
- Hotswapping the active supervisor, and then issuing "redundancy manual switchover" will cause stateful switchover to fail. (718863)
Series Affected: CCS-750 Series
- Performing SSO may result in extended traffic loss. (748533)
Series Affected: CCS-750 Series
- Feature agent StrataL2 would continuously restart when the number of unique multicast replication list goes higher than 16k. Workaround is to reduce the multicast scale (784949)
- In some rare cases, after a stateful switchover, the switch could start discarding all packets. To recover from the problem state, "platform trident Switchcard1/0 reset" and "platform trident Switchcard2/0 reset" must be executed." (787513)
- When the core VLAN of a VXLAN VLAN has a LAG and needs to flood edge-to-core traffic to all of its ports, the addition of another non-LAG port to that core VLAN results in traffic not egressing the LAG. The removal of the LAG and the addition of another port to the core VLAN results in traffic not egressing this new port.

Restart the slice agent to recover from this condition. (828814)

- When IGMP Snooping is enabled, SwitchCard or LineCard restarts or resets may result in unexpected restart of StrataL3 agent. (830962)
- Known unicast traffic may not flood to all ports of a core VLAN when the underlay MAC of the nexthop to the remote VTEP hasn't been learnt.

Flap the VXLAN interface or restart the StrataL3 agent to recover. (838183)

- If there is a loop in the network we could see a lot of MAC moves, leading to high CPU utilization by the StrataL2 agent and causing problems with learning and aging. (845250) (1)
 SKUs Affected: 7300X3-32C-LC, CCS-710P-12, CCS-720DF-48Y-2F, CCS-720DP-24S-2F, CCS-720DP-48S-F, CCS-720XP-24Y6-R, CCS-722XPM-48Y4-F, DCS-7060CX-32C-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7304 and DCS-7308
- StrataL3 agent restart may result in extended traffic loss of more than 1 seconds on VXLAN VLANs. (855597)

- after a hitful reset on Trident3X4 platform, vxlan encap in the case of v6 vxlan underlay would fail. (855762)

- StrataFlowTracker forwarding agent (StftAgent) may restart unexpectedly, when "flow tracker hardware" is turned on and off in a short interval interval. (922896) (1)

- Performing a physical supervisor pull with SSO enabled, with an ECMP group containing interfaces from both supervisors and passing L3 multicast traffic, may result in extended traffic outage for the switchover. (1002219)

SKUs Affected: CCS-755-CH and CCS-758-CH

- QoS policy-maps with atleast one policer attached cannot be scaled over 195 interfaces (1003605)

Series Affected: CCS-750 Series

- With EVPN multicast and `platform trident forwarding-table partition flexible exact-match <> l2-shared <> l3-shared <>` configured for supporting underlay IIF switch cases like RPT-SPT switchover, underlay ECMP link switch, etc., we may see BessMgr process spinning at around 90-100% if there another feature using UDF is configured.

Work-around is to remove the other feature which uses UDF. (1016232) (1)

- Stateful Switchover due to active supervisor removal may cause an extended data-plane outage.

Workaround is to pull out the links from the active supervisor, wait 30 seconds, and then remove the active supervisor. (1017834)

- Performing supervisor stateful switchover on a device running Multiple Spanning Tree Protocol (MSTP) with multiple MST instances may cause STP topology to reconverge. (1042130)

- With EVPN multicast and `platform trident forwarding-table partition flexible exact-match <> l2-shared <> l3-shared <>` configured for supporting underlay IIF switch cases like RPT-SPT switchover, underlay ECMP link switch, etc., we may see BessMgr process spinning at around 90-100% if there another feature using UDF is configured.
Work-around is to remove the other feature which uses UDF. (1053562)
- In a scenario with a large number of VXLAN VLANs and VTEPs, if the flood VTEP list experiences churn—either due to flood VTEP reachability issues or VTEPs rebooting—there may be exhaustion of the IP multicast group (IPMC) and replication group (ReplGroup) lists. This situation can result in VXLAN VLAN broadcast, unknown unicast, and multicasting (BUM) traffic being blackholed. A workaround for this issue is to flap the VXLAN interface.
(1061846) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- An SEU (Single Event Upset) in the L2_ENTRY_LP, VLAN_XLATE_1_LP, VLAN_XLATE_2_LP, EGR_VLAN_XLATE_1_LP, EGR_VLAN_XLATE_2_LP, MPLS_ENTRY_LP, or EXACT_MATCH_LP tables may result in an unexpected forwarding agent restart causing a brief traffic disruption. (1134953) (1)

DCS-7050X3 Series

- An SEU (Single Event Upset) in the L3_ENTRY_LP table results in an unexpected forwarding agent restart causing a brief traffic disruption.
(641120) (1)
- Changing interface speed can cause all subsequent link up and link down notifications to be delayed by approximately one second on all ports. The workaround is to restart the forwarding agent after the speed change using the CLI command "platform trident reset". This causes all ports to flap and a brief traffic outage to occur. (810073)
- If there is a loop in the network we could see a lot of MAC moves, leading to high CPU utilization by the StrataL2 agent and causing problems with learning and aging. (845250) (1)
SKUs Affected: 7300X3-32C-LC, CCS-710P-12, CCS-720DF-48Y-2F, CCS-720DP-24S-2F, CCS-720DP-48S-F, CCS-720XP-24Y6-R, CCS-722XPM-48Y4-F, DCS-7060CX-32C-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7304 and DCS-7308
- Fru agent will crash on systems with base MAC address ending in :fd (927753)
SKUs Affected: DCS-7050TX3-48C8

- With EVPN multicast and `platform trident forwarding-table partition flexible exact-match <> l2-shared <> l3-shared <>` configured for supporting underlay IIF switch cases like RPT-SPT switchover, underlay ECMP link switch, etc., we may see BessMgr process spinning at around 90-100% if there another feature using UDF is configured.
Work-around is to remove the other feature which uses UDF. (1016232) (1)
- Exact match table lookup is not supported with IPV6 underlay multicast enabled. As a result we see duplicate packets when we enabled platform trident forwarding-table partition flexible exact-match configuration.
(1018543) (1)
- In a scenario with a large number of VXLAN VLANs and VTEPs, if the flood VTEP list experiences churn—either due to flood VTEP reachability issues or VTEPs rebooting—there may be exhaustion of the IP multicast group (IPMC) and replication group (ReplGroup) lists. This situation can result in VXLAN VLAN broadcast, unknown unicast, and multicasting (BUM) traffic being blackholed. A workaround for this issue is to flap the VXLAN interface.
(1061846) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- An SEU (Single Event Upset) in the L2_ENTRY_LP, VLAN_XLATE_1_LP, VLAN_XLATE_2_LP, EGR_VLAN_XLATE_1_LP, EGR_VLAN_XLATE_2_LP, MPLS_ENTRY_LP, or EXACT_MATCH_LP tables may result in an unexpected forwarding agent restart causing a brief traffic disruption. (1134953) (1)

7300X3, AS7326, AS7726 and DCS-7050X3 Series

- On interfaces with 40GBASE-SRBD transceivers operating at their default speed and "no transceiver qsfp default-mode 4x10G" configured, SSU may cause more than 100 seconds outage.

The workaround is to configure such interfaces with 40GBASE-SRBD transceivers to "speed forced 40g" (761971)

- If there is a loop in the network we could see a lot of MAC moves, leading to high CPU utilization by the StrataL2 agent and causing problems with learning and aging. (845250) (1)

SKUs Affected: 7300X3-32C-LC, CCS-710P-12, CCS-720DF-48Y-2F, CCS-720DP-24S-2F, CCS-720DP-48S-F, CCS-720XP-24Y6-R, CCS-722XPM-48Y4-F, DCS-7060CX-32C-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7304 and DCS-7308

- Momentary MAC Rx local faults and remote faults may not be reflected in MAC Rx fault counters and may not cause the interface state to flap. (880612)
SKUs Affected: 7300X3-32C-LC, 7300X3-48TC4-LC and 7300X3-48YC4-LC
- SSU from a release < EOS-4.24.0F may fail with a fatal error if interfaces on a single QSFP port are configured with speeds 50G and 10G, 25G or default. Such interfaces are identifiable by the speed-misconfigured errdisabled state. The workaround is to configure matching speeds on all interfaces in a QSFP port. (891375)
- With EVPN multicast and `platform trident forwarding-table partition flexible exact-match <> l2-shared <> l3-shared <>` configured for supporting underlay IIF switch cases like RPT-SPT switchover, underlay ECMP link switch, etc., we may see BessMgr process spinning at around 90-100% if there another feature using UDF is configured.
Work-around is to remove the other feature which uses UDF. (1016232) (1)
- Exact match table lookup is not supported with IPV6 underlay multicast enabled. As a result we see duplicate packets when we enabled platform trident forwarding-table partition flexible exact-match configuration. (1018543) (1)
- In a scenario with a large number of VXLAN VLANs and VTEPs, if the flood VTEP list experiences churn—either due to flood VTEP reachability issues or VTEPs rebooting—there may be exhaustion of the IP multicast group (IPMC) and replication group (ReplGroup) lists. This situation can result in VXLAN VLAN broadcast, unknown unicast, and multicasting (BUM) traffic being blackholed. A workaround for this issue is to flap the VXLAN interface. (1061846) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)

7358 and DCS-7050X4 Series

- When more than 500 access list rules are configured and applied to range of interfaces, slice agent may encounter SIGQUIT and restart when shut and no shut on the range of interfaces.
Workaround is to reduce number of access list rules. (679292)
Series Affected: 7368 Series
- When security ACL is being updated, port blocking rules are applied which drop the packets to avoid mis forwarding. This will also block traffic from reaching the control plane. This will cause the protocol like BGP, BFD etc to not work temporarily which requires control plane for processing.

As a workaround, use CLI ``hardware access-list update default-result permit`` to avoid port blocking rules taking affect. With this configuration packets are not dropped due to port blocking rules when ACL is being updated. (703753)

- StrataL3Unicast agent restart may result in the flapping of Layer3 sub interfaces. (732310)
- After parity error occurs in LTS_TCAM memories, traffic forwarding may not work as expected.
Workaround is to restart the Strata-FixedSystem agent after the parity error. (756038)

(1) This item is in two or more sections on this document

7388 and DCS-7060X5 Series

- StrataMnu agent has high CPU usage when PFC watchdog is configured with a low polling interval, such as 0.010. (636394) (1)
- When StrataL2 agent restarts, some static MAC addresses pointing to port-channel may not get programmed in hardware table resulting in flooding.
Workaround is to removing and add the static MAC address configuration. (787583)
- With the scale of VXLAN VLANs and flood VTEPs, churn in flood VTEP list may result in exhaustion of replication group used for flooding.
Workaround is to flap the VXLAN interface. (824256)
- With the scale of VXLAN VLANs and flood VTEPs, churn in flood VTEP list may result in exhaustion of replication group used for flooding. Workaround is to flap the VXLAN interface. (840624)
- When the egress port mirror-to-cpu feature is enabled on a VXLAN core port, sending packets of certain sizes may result in the some types of CPU-bound traffic getting corrupted, resulting in protocols not converging.

These VXLAN traffic may also cause parity errors in the egress data buffer of the switch.

As a workaround, disable all CPU monitoring sessions. (858716)

Series Affected: 7388 Series

SKUs Affected: 7388-16CD, 7388-16CD2, 7388-8D, 7388-SUP, 7388-SUP-D, 7388X5-SC, DCS-7060DX5-64-F and DCS-7060PX5-64E-F

- If an interface is shutdown, static MACs on that interface may not be counted in number of allowed addresses for port-security (862079) (1)

- VXLAN VTEP(s) and HER nexthop entries may not be cleared/deleted from the system ("show vxlan vtep command") , even after corresponding EVPN routes are deleted, when the system fails to allocate new multicast groups due to scaled config.

Workaround is to flap the VXLAN interface. (908722)

- Remote flood VTEP removal/addition can trigger nexthop table full scenarios resulting in blackholing of traffic.

Workaround: Restart StrataL3 agent. (916116)

- Configuring auto-negotiation on an interface (e.g., "speed auto 400g-8") may cause an unexpected restart of the forwarding agent. This leads to a flap of all interfaces and a brief traffic outage. (928889)

SKUs Affected: DCS-7060DX5-64 and DCS-7060DX5-64S

- "show platform trident l3 shadow next-hops ifp raw" may result in unexpected restart of StrataL3 forwarding agent. (954273) (1)
- StrataL3 agent may restart unexpectedly if learned remote VTEP address using EVPN is a IPv6 address.
Removing IPv6 underlay VXLAN config should resolve the issue. (994571)
- Routing may behave unexpectedly under edge case circumstances (1009562)
- Some interfaces may flap intermittently due to jitter (1013413) (1)
SKUs Affected: DCS-7060DX5-64 and DCS-7060PX5-64E
- A software correctable bit parity error in MY_STATION_2_TCAM table in hardware may remain uncorrected affecting the routed traffic. There is no workaround. (1025255)

- Continuous churning of "Unknown Unicast" or "Unknown Multicast" multicast group memberships for VXLAN VLANs may result in memory leak within StrataL3 Agent. (1037966)

Series Affected: 7388 Series

SKUs Affected: DCS-7060PX5-64E and DCS-7060PX5-64E-F

- Single Event Upset (SEU) events that occur on the EGR_VFI table are not corrected which can impact egress traffic. (1039015)
- IPv6 address profile isolation-criteria must be less than 64 bits in length. That is, the offset plus the length cannot exceed 64 bits. Configuring such a profile and applying it to an interface will cause the Strata Agent to restart, causing forwarding plane traffic to be dropped, until the IPv6 address profile is either removed from the interface, or the offset and/or length is reduced to 64 bits or less. (1042951)
- Traffic to a multihomed MAC destination may see packet drops during VTEP reachability change when putting spine devices into maintenance mode.

There is no workaround for this issue. (1044420) (1)

- Some routed packets maybe mis-forwarded by the CPU, which may lead to unexpected behavior. (1052025)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)

(1) This item is in two or more sections on this document

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)
- 100GBASE-SR4 modules may not link up or experience errors when inserted into a QSFP-DD slot or an OSFP slot using an OSFP to QSFP Adapter (OQA) on some platforms.

The workaround is to override the transceiver tuning with the command 'transceiver electrical lane 1-4 rx-output-amplitude module-default rx-output-pre-emphasis module-default tx-input-equalization module-default' (631109)

SKUs Affected: 100GBASE-SR4

- After ASU2 upgrade, some QSFP power class 8 transceivers (including 100GBASE-ZR4) may stop transmitting.

Operational functionality of the transceiver can be restored by issuing interface configuration command 'transceiver diag simulate removed' followed by 'no transceiver diag simulate removed' (755695)

- SFP-10G-MRA-T 10GBASE-T ports configured with 'speed auto 100full' or 'speed 100full' (100M port speed) may experience extended traffic loss during the Leaf Smart Software Upgrade (SSU) boot up sequence. (866682)
- Revision 20 QSFP-100G-DR and QSFP-100G-FR transceivers with serials starting with XDP may result in link flaps
In most cases, the work around is to increase the Rx output amplitude with the command 'transceiver electrical lane 1-4 rx-output-amplitude 3 rx-

output-pre-emphasis 2' (910190)

SKUs Affected: 100GBASE-DR and 100GBASE-FR

- Link may not come up on certain 100GBASE-PSM4 modules. (929400) (1)
SKUs Affected: DCS-7280DR3A-54, DCS-7280DR3AK-54 and DCS-7280DR3AM-54
- Issuing any "transceiver firmware update" or "transceiver firmware activate-standby" commands on ports that in the past have executed such commands may produce a timeout CLI message if there was an XcvrAgent restart event at some point in time between the previous instance of the firmware CLI commands being executed and the present. (1004332)
SKUs Affected: OSFP-400G-ZR, OSFP-400G-ZRP, QDD-400G-ZR and QDD-400G-ZRP
- QSFP-100G-LR4 or QSFP-100G-LRL4 transceivers with serials beginning with XHT or XMH may experience high pre-FEC BER and uncorrectable FEC errors on the peer or fail to link up.

A workaround is to adjust the tx input equalization of the transceiver with the command 'transceiver electrical lane 1-4 tx-input-equalization module-default' in config-if mode. The command is to be applied on the /1 of the QSFP interface of transmitting system with the impacted transceiver.

(1133858)

SKUs Affected: 100GBASE-LR4 and 100GBASE-LRL4

(1) This item is in two or more sections on this document

7300X3, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX and DCS-7050X3 Series

- CoS 0/1 packets coming on MLAG peer link with CoS trust configured may take incorrect Tx-queue while egressing.
Workaround : QoS maps can be changed for cos0/1 to non-default TxQ for the MLAG peer link but it would change the mapping for all the ports (407056)
- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (416694)
- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (417223)
- 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

- 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)

- Disabling static NAT access-list sharing with "no hardware nat static access-list resource sharing" may in some cases leak hardware VFP resources.

Workaround is to reload the switch. (539951)

- Some multicast NAT translations may stop working after the restart of Strata forwarding plane agent.

Workaround is to remove and re-add the affected rules. (544827)

- With GRE tunnels or IPinIP tunnels or Nexthop-group tunnels configuration along with subinterface in the tunnel core interface, tunnel encapsulated packet will not be tagged with dot1q tag as configured in the subinterfaces, instead packet will go out tagged with internal vlanId assigned for the subinterface. There is no workaround for this issue other than removing subinterface configuration in the core interface. (548327)

- When VXLAN encapsulated packet is received with an inner payload tag, the switch may decapsulate and send the packet with inner tag stripped when the egress chip is not same as the ingress chip. There is no workaround to this issue. (548492)

SKUs Affected: 7320X-32C-LC, CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F

- Some NAT translations may stop working after rebooting a switch.

Workaround is to restart the Strata Slice Agent. If the problem persists, perform a cold reboot of the switch. (563526)

- When VXLAN IPv6 underlay is configured and when core interface is a port channel on a SVI, continuous port channel interface flaps may result in persistent loss of VXLAN traffic to remote VTEPs that are reachable through that port channel. (567455)

- VXLAN encapsulated packets received with inner payload having more than one dot1q tag will result in outer dot1q tag being stripped when egressing out

on a port. (567587)

Series Affected: 7300X3 and CCS-750 Series

SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F

- Replacing new scaled configuration with old scaled config may cause StrataL3 agent to restart. Feature will function normally after agent restart. (591050)
- When a linecard with LAG members configured statically is removed, subsequently moving or adding the LAG to a different VLAN may cause BUM traffic on that VLAN to be blackholed.

To work around this issue, remove the removed LAG members from the LAG. (603388)

SKUs Affected: 7300X3-32C-LC and 7300X3-48YC4-LC

- When StrataLag and StrataL3 forwarding agents restart, Egress VLAN translation (configured using "switchport vlan mapping") configured on port channel may not work.
The workaround is to shut and un-shut the impacted Port Channel. (603906)
Series Affected: CCS-710P, CCS-720DP, CCS-720DT, CCS-720XP, CCS-750 and DCS-7010TX Series
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- On switch reboot/reload with NAT configured, the ports come up and then NAT configuration is applied. If traffic is flowing before the NAT configuration is fully programmed, then some packets might go out untranslated or dropped. Workaround is to stop traffic till NAT configuration is fully applied. (606202)
- If dynamic NAT configuration is modified or if the admin shutdown/no shutdown performed on an interface configured with dynamic NAT while NAT traffic is flowing, NAT connections may be left untranslated.

Workaround is to manually flush all the entries using the command "conntrack -F" in the corresponding network namespace bash mode. If configured on a non-default VRF (vrfName, for example), run the command "ip netns exec ns-vrfName conntrack -F". (615735)

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

A forwarding agent restart is required to forward traffic normally. (626452)

- Configuring 'switchport source-interface tx [multicast]' while the Strata-FixedSystem/Strata-Linecard agents are shutdown will result in an unexpected StrataL3 agent restart.

Workaround: Ensure the Strata-FixedSystem/Strata-Linecard agents are

running before configuring 'switchport source-interface tx
[multicast]' (635275)

- A single-event upset (SEU) in the IFP_TCAM can lead to forwarding agent restart, after which the switch forwards traffic normally. (641105)
- With EVPN multicast solution configured, first few underlay multicast <S,G> entry may not get programmed; leading to any traffic arriving with the underlay <S,G> to be punted to CPU.
Workaround is to create an extra VLAN/VRF to underlay multicast group mapping before configuring the rest of the mappings. (643866)
- During config replace with NAT configuration, TCAM churn may result in forwarding agent restart. (676117)
- With IPV6 VXLAN Underlay configuration, when routing is toggled VLAN encapsulated packets sent on routed ports may be sent incorrectly with a VLAN tag resulting in packets getting dropped.
Workaround is to flap the core underlay interface. (679980)
- SSO switchover may result in remote ARP entries getting refreshed resulting in brief traffic loss for traffic routed to VXLAN remote hosts. (724315)
SKUs Affected: CCS-755-CH, CCS-758-CH, DCS-7304 and DCS-7308
- Under heavy lookup scenario, if the host CPU performs a write to one of hardware table responsible for MPLS/VXLAN tunneling feature, while a tunneled packet is undergoing a second lookup in the same table, it can result in packet drops or mis-forwarding for the flow.

Workaround is to restart the slice agent after stopping the traffic. (753304)
Series Affected: 7300X3, DCS-7050CX3, DCS-7050SX3 and DCS-7050TX3 Series

- With a large route scale, executing the CLI command "show platform trident 13 software alpm-table routes" may result in the restart of multiple forwarding agents. There is no workaround for this issue. (762247)
- After linecard hotswap or slice agent restart, static or peer-learned MAC addresses may go missing in hardware, which can result in unexpected flooding. Further, such missing macs may incorrectly be displayed in the output for "show mac-address-table".
A workaround for this issue is to remove and readd the MAC address. (762363)
- On multichip systems, reload may cause VLAN traffic-policy rules to not be reprogrammed across all chips. Workaround would be to reconfigure the VLAN traffic-policy rule. (824297)
SKUs Affected: 7300X3-32C-LC, 7300X3-48TC4-LC, 7300X3-48YC4-LC, DCS-7304 and DCS-7308
- On devices running EVPN VXLAN multicast OISM, Overlay multicast traffic with TTL1 arriving on VXLAN enabled VLAN in a VRF with `evpn multicast` enabled

is not VXLAN encapsulated when there is a receiver in the same VLAN on remote VTEP.

Use TTL>1 or enable `redistribute igmp` under router BGP for the VLAN on the receiver VTEP. (833210)

- Executing "platform trident * reset" continuously may result in unexpected restart of StrataL3 forwarding agent. (854688)
- If StrataNat agent is killed or shutdown while NAT traffic is flowing, it may cause the StrataNat agent to continuously restart. System needs to be rebooted to recover from it.

Workaround is to stop NAT traffic before shutting down the StrataNat agent. (862801)

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

- When an EVPN ethernet-segment is configured on a port channel having members managed by multiple primary forwarding agents, removal of members of the port channel might cause the agents to restart. (864951)
- Sometimes the acls programmed for dynamic NAT rules (in the TCAM) may not be unprogrammed after the corresponding rules are un-configured.

Restarting the forwarding plane agent should force the original rules to be unprogrammed. (881869)

- When the StrataMirror agent restarts while egress port mirror-to-cpu is configured, mirrored traffic may be flooded temporarily on the VLANs of the mirror source port(s). Workaround is to use a recirculation channel with egress mirror-to-cpu. (883792)
- Global routing disable with MSS-G enabled can result in switch local IPs being unreachable.

Workaround: Enable routing globally. (890058)

- When 'reload fast-boot' is performed while egress mirror to CPU monitor session is configured, mirrored traffic may get flooded on the VLAN of egress interface. Workarounds: 1) turn off egress mirror to CPU before 'reload fast-boot' 2) if the issue happens after SSU, flap the monitor session after the reload to recover (925304)
- Changing a port belonging to a VXLAN VLAN, supporting IPv6 encapsulation, from access mode to trunk mode may result in an incorrect ingress VLAN membership for that port. (928173)
- In a VXLAN MLAG setup, reload of an MLAG peer can cause invalid hardware resource identifier to be assigned for the MLAG peer VTEP which can lead to

packets destined from one peer to the other peer to be dropped.

Workaround is to flap the Vxlan1 interface. (929921)

Series Affected: CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, DCS-7050X3 and DCS-7300X3 Series

- When mirroring truncation is enabled on a TX mirror-to-CPU monitor session, the Strata slice agent may crash, flapping all ports on the device.
Workaround is to ensure truncation is disabled (937673)

- In Active-Active NAT synchronization mode between the two peers, TCP and ICMP connection establishment might fail if the control packets are received in the other peer.

A workaround is to make sure that all the control packets for a specific connection arrive on the same peer. (954033) (1)

- Cos value in the egress packet after VXLAN decapsulation is not set correctly in a MultiChip switch , when ingress and egress ports are in different chips. Cos value is always 0 even if cos rewrite is enabled. (988967)

SKUs Affected: CCS-720XP-96ZC2-F

- ECN bit may not be set correctly in VXLAN Encapsulated packet in a MultiChip switch , when ingress and egress ports are in different chips. (992398)

SKUs Affected: CCS-720XP-96ZC2-F

- If NAT fragmentation settings are changed where there are active NAT connections/traffic, NAT translations might stop working.

Recommendation is to change NAT config when there is no NAT traffic.

Work-around is to restart StrataNat agent. (992522)

- If a large amount of flows are continuously being re-learned, Strata slice agent may restart. (993211) (1)
- In an EVPN VXLAN active-active multi-homing setup, if the designated forwarder (DF) election result changes rapidly or the interfaces in local ethernet segments flap rapidly, then multi-homed hosts connected to the impacted ethernet segments could start receiving duplicate BUM packets. (1008046) (1)
- In scenario where MAC table has same MAC with multiple VLANs, Per-MAC ACL feature will not update all MACs with allocated ACL ID. Also It might not cleanup ACL ID from MAC table on removing configuration. There is no known workaround for this case other than removing affected MAC and programming it again. (1008138)

- IP directed-broadcast source host filtering rules match on switched traffic in addition to routed traffic. (1017913)
- Traffic sourced from the switch may result in drops if "segment policy policy-drop-all default" is configured.

Workaround: Remove "segment policy policy-drop-all default" configuration. (1033038)

Series Affected: 7300X3 Series

SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F

- VXLAN IPv6 underlay encapsulation flows (tunnel initiation) may experience a few milliseconds traffic loss when the new members are added to the underlay ECMP. (1037644)
- Traffic to a multihomed MAC destination may see packet drops during VTEP reachability change when putting spine devices into maintenance mode.

There is no workaround for this issue. (1044420) (1)

- With IPv6 VXLAN underlay , if underlay to reach remote VTEP(s) uses ECMP and egress interfaces are Port-channel(s) then VXLAN MEMBER tunnels may remain down upon port-channel interface flap.
Workaround is to flap all underlay port-channel interfaces or replace port channels with individual routed interfaces to resolve the issue. (1050592)
- In a scenario with a large number of VXLAN VLANs and VTEPs, if the flood VTEP list experiences churn—either due to flood VTEP reachability issues or VTEPs rebooting—there may be exhaustion of the IP multicast group (IPMC) and replication group (ReplGroup) lists. This situation can result in VXLAN VLAN broadcast, unknown unicast, and multicasting (BUM) traffic being blackholed. A workaround for this issue is to flap the VXLAN interface. (1061846) (1)
- With per-MAC ACL configured, 802.1x supplicants that continuously authenticate and try to attach to a unique ACL may result in out of memory situations where the OOM killer is invoked and starts killing other agents (1108413)
- If the device is restarted when overload mode dynamic network address translation is configured and the NAT traffic is running then the NAT translations may brake and some or all of the flows subject to NAT are being forwarded without translation after the device completes the restart.

Workaround is to manually flush all the NAT entries using the command "conntrack -F" in the corresponding network namespace bash mode. If configured on a non-default VRF (vrfName, for example), run the command "ip netns exec ns-vrfName conntrack -F". (1122152) (1)

7358 and DCS-7050X4 Series

- StrataMmu agent has high CPU usage when PFC watchdog is configured with a low polling interval, such as 0.010. (636394) (1)
- Traffic loss may be seen during StrataL2 restart (745553)
- MAC Moves may happen for very small duration when forwarding agent restarts on any of the peers of a MLAG. In case of VXLAN BGP EVPN, this may result in the peer learned MACs being locally learned and being advertised back to remote VTEPs. (749358)
- When a scaled QoS policy-map config with policers is applied in ingress direction, forwarding agent(s) may unexpectedly restart. (751937)
- Ipv4 microBfd packet sent to non-routed port with interface destination mac will be treated like regular L2 traffic. (769736)
- When BFD is configured but the session is down, BUM traffic can cause loops in a LAG. (778369)
- Simultaneously restarting multiple forwarding-plane agents might result in a small memory leak in AsicResourceMgr agent. (784711)
- If QoS Policy-maps are using some ACLs with rules having many L4 port ranges, in some cases, it is possible that rules are not programmed correctly in hardware. There is no error reported while programming the rules and it will have unexpected behavior in terms of traffic flow. (786011)
SKUs Affected: 7358X4-SC, DCS-7050DX4-32S, DCS-7050DX4-32S-F, DCS-7050PX4-32S and DCS-7050PX4-32S-F
- BGP 'update wait-install' command is not supported on this platform. If this command is present in the startup-config, BGP routes are not advertised to peers.

To workaround this issue, remove 'update wait-install' command from startup-config and replace the config using 'config replace flash:startup-config' command or reboot the switch. (796360)

- Simultaneously restarting multiple forwarding-plane agents might result in a small memory leak in AsicResourceMgr agent. (807391)
- Scaling down static MAC address configuration while traffic with the same source MAC address as one of the removed MAC addresses is running might result in the MAC address not being learnt.
A workaround is to execute the "clear mac-address-table dynamic" command to prompt relearning of such MAC addresses. (852039)

- Scaling down static MAC address configuration while traffic with the same source MAC address as one of the removed MAC addresses is running might result in the newly learnt dynamic MAC address still being considered a static MAC address.

A workaround is to execute the "clear mac-address-table dynamic" command to prompt relearning of such MAC addresses. (852203)

- If an interface is shutdown, static MACs on that interface may not be counted in number of allowed addresses for port-security (862079) (1)
- Encapped VXLAN packets can be dropped after multiple StrataL2 and StrataL3Unicast restart at the same time. The workaround is to restart StrataL3Unicast alone (949366)
- In an EVPN VXLAN active-active multi-homing setup, if the designated forwarder (DF) election result changes rapidly or the interfaces in local ethernet segments flap rapidly, then multi-homed hosts connected to the impacted ethernet segments could start receiving duplicate BUM packets. (1008046) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1072368)

(1) This item is in two or more sections on this document

Limitations and Restrictions in 4.29.9.2M

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-7260CX3-64, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F and DCS-7260CX3-64E-R
- Aborting 'reload fast-boot' or 'reload fast-boot now' command with Ctrl-C is unsupported. It may lead to the incorrect reload cause in 'show reload cause' in the subsequent reload. (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)
- Performing SSU to boot into an image with this limitation may require that the currently running image is deleted to ensure that there is sufficient space available to perform the upgrade. (525870)
- Performing SSU reload with a transceiver that is unsupported in the source image but supported in the destination image may cause the port not to link up after SSU reload completes.

The workaround is to hotswap the affected ports after SSU reload. (955515)

BGP EVPN L2/L3 VXLAN/MPLS

- 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)
- VRRP+VXLAN is not supported in deployment with EVPN VXLAN. (549525)
- In an EVPN/VXLAN environment with VLAN-Based MAC-VRF, type 2 and type 3 routes are incorrectly imported when the route-targets match even if the VNIs are different.

The workaround is to use different route-targets for the VLAN-Based MAC-VRF's to prevent the routes from being incorrectly imported. (631589)

- For an EVPN gateway to advertise received IP-PREFIX routes with next-hop-self, the corresponding IP-VRF must be active, which requires an interface to be configured in the VRF or the configuration of Router ID. BGP IP-VRF status is shown via "show bgp instance vrf VRF". (703768)
- IP ARP proxy is not compatible with EVPN VXLAN enabled VLANs.

Disable IP ARP proxy on VLANs with EVPN VXLAN enabled. (710344)

- In a EVPN DCI deployment, Symmetric Integrated Routing and Bridging is not supported in the default VRF. (733143)
- In a EVPN VXLAN deployment with IP prefix routes with a gateway IP, address resolution may be unevenly distributed across various VXLAN endpoints, VNIs and router MAC addresses, each destination (VTEP/VNI/router MAC) as identical VIAs are not considered when forming ECMP. (810595)

- In networks using EVPN Active/Active multihoming with the modulus designated forwarder election algorithm, EOS versions prior to 4.29.0 consider all of the VLANs configured in a vlan-aware-bundle MACVRF when calculating the DF. EOS versions starting with 4.29.0 only consider VLANs configured in both the MACVRF and on the device itself. This can lead to disagreements on the elected DF if the range of VLANs configured in a bundle MACVRF is wider than the range of VLANs actually configured on the device, and the devices in the ethernet segment are running disparate EOS releases on either side of version 4.29.0. If different devices in the same ES elect different DFs, traffic loss can ensue.

As a workaround, either make sure all devices participating in an EVPN active/active multihoming ethernet segment are running the same version of EOS, or alternatively, before upgrading any device, ensure that the lowest VLAN configured in a vlan-aware-bundle is also configured on the device itself. This will ensure that the VLAN considered as the lowest VLAN in the MACVRF is the same pre and post-upgrade. (868750)

- In EVPN setups, the MAC duplication detection logic helps prevent control plane churn by adding MACs that move over a certain number of times over a certain rolling period of time. This is typically caused by an L2 loop. While control plane churn is mitigated, the data plane traffic remains affected from the L2 loop.

An event-handler such as one attached to the bug can be deployed as a workaround. The event-handler is triggered when a syslog is thrown in response to a host being denylisted, and it removes the VLAN corresponding to the host from the offending interface. Since the syslog does not contain the offending interface, the interface lookup is done in a separate call. If the host is removed before the interface could be obtained, the event-handler would not take any action (1117596)

Containerized EOS

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

vEOS Router / CloudEOS

- The /mnt/flash/cloud_ha_config.json based configuration for the CloudHa agent is not supported from this release. If you are upgrading from an older image (< 4.20.5) please reconfigure HA using EOS CLI. Please see information on converting json file based config to CLI commands in the vEOS Router Configuration Guide. (264660)
- While vEOS instances in an AWS cloud can be created using either a Bring Your Own License (BYOL) or Pay As You Go (PAYG) format, there currently is

no show command in vEOS to easily help the user or support determine which mode this instance uses. (272649)

- Tx/Rx ring parameters of an "Elastic Network Adapter (ENA)" interface cannot be set through config.json in vEOS. (276382)
- vEOS supports up a maximum of 8 physical cores. With hyperthreading enabled it is 16 vCPUs (278286)
- In dpdk based vEOS(MODE=sfe in /mnt/flash/veos-config), even a packet matches a data plane ACL rule with log action, the log message might not show up if there is a similar rule with more exact ttl/dscp value match before that rule. (293641)
- In dpdk based vEOS(MODE=sfe in /mnt/flash/veos-config), SR-IOV mode doesn't work with Intel x520/82599 on ESXi. As a workaround use mixed mode instead of the SRIOV only mode in ESXi (296718)
- Hot detaching network interface instances is unsupported. Reboot the instance after detaching a network interface. (306132)
- CloudEOS does not support more than 10k /30 routes. The SFE agent will restart if more than 10k /30 routes are added to the system. The workaround is to reduce the number of /30 prefixes. (540204)
- NIC i40en driver version 2.2.4.0 is required on ESXi 7.0 host for SRIOV virtual function trust mode on/off functionality. (680174)
- NAT interface configuration mode is not supported in SFE. This needs to be kept in mind when migrating to SFE mode from Sfa.

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

- The command "show flow tracking hardware flow-table" may trigger OOM events when there are over 1 million flows present in the system. This could occur as a result of many short lived flows and the need to process many adds or deletes.

The flow-tracking information can be viewed correctly on Cloud Vision As A Service (721460)

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)
- In a Macro-Segmentation Service (MSS-FW) deployment with Palo Alto Networks Panorama, security zones configured on the Panorama are not reported in the API MSS uses.

As a workaround, configure the zones directly on the firewalls. (363371)

- The MacroSegmentation Service (MSS) allows users to configure multiple tags on CVX that can be used in firewall policy definition, to identify policies MSS should work on. It is however not possible to delete one of the multiple tags configured.

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

- In Macro-Segmentation Service (MSS) for L3 Firewalls , Virtual SVI IP addresses cannot be treated as intercepted hosts and thereby cannot be

added to any redirect policy. However, these IPs can still be part of offload policies, even though implicit redirect will not work for them. (372504)

- If the Macro-Segmentation Service (MSS), working with a L3 firewall, is configured to intercept traffic from a host generating multicast traffic, the traffic can be black-holed.

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

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

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

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

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

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

- The MssPolicyMonitor agent might unexpectedly start running and eventually exit when the CVX functionality is disabled on the master CVX instance. This does not have any side-effect on the operation of the Macro-Segmentation Service (MSS). (378077)

- When the Macro-Segmentation Service (MSS) is enabled and works with a L3 firewall, there can be a momentary traffic outage on reloading a MLAG peer if the CVX to VTEP (MLAG switch) connection is not re-established before the MLAG reload delay expires at the reloaded peer. (418560)
- 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)
- Macro-Segmentation Service (MSS) running in consistent enforcement mode may emit "CVX-3-MSS_HOST_UNPROGRAM_ERROR" syslog messages for the hosts attached to only one MLAG peer when that peer is reloaded.

These messages can be ignored. "CVX-3-MSS_HOST_RECOVER" messages will be observed as soon as these hosts are learned again at the reloaded MLAG peer afterwards. (524705)

- For Macro-Segmentation Service (MSS) deployment running in consistent enforcement mode, "CVX-3-MSS_HOST_UNPROGRAM_ERROR" or "CVX-3-MSS_RULE_UNPROGRAM_ERROR" syslog messages may be seen when a switch is reloaded (planned reload).

Use "clear cvx connections client <switch-id>" command to clear the switch from CVX registration. (524717)

- In a Macro-Segmentation Service (MSS) deployment in a multi-VRF environment with EVPN as control plane, when an MLAG peer is reloaded, CVX-3-MSS_HOST_PROGRAM_ERROR message is logged for the intercept hosts followed by CVX-3-MSS_HOST_RECOVER message. (525851)
- In a Macro-Segmentation Service (MSS) deployment with a Palo-Alto Firewall, when using an aggregator interface in a virtual system, all its member interface must be in the same virtual system. (528511)
- In a Macro-Segmentation Service (MSS) deployment with a Palo-Alto firewall, a policy with multiple source or destination zones will be enforced over only one source zone and/or one destination zone.

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

- In a Macro-Segmentation Service deployment with Layer 3 firewall (MSS-FW), when a policy matching a subnet with a mask less than /16 is tagged as a group-based redirect or a regular offload policy, the service may become non-operational. However, any subnet mask length is supported for all variants of verbatim policies. (716283)

Enterprise routing

- Any config change in DSCP value configured via AVT policy/Application profile is not updated for existing flows. New flows created after the config change inherit the updated DSCP value. (873846)

EOS SDK

- In the "management api eos-sdk-rpc" mode the default values for a configured "transport grpc" do not show up when running "show active all". (1084076)

General

- CLI does not allow access to files whose names contain spaces. (539)
- 'speed forced 1000full' is not supported on interfaces with the "Management" prefix such as ma1. Use speed autonegotiation instead. (1506)
- Ports configured as mirroring destinations will be shown as active even when they are shutdown. (65221)
- 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)

- On the DCS-7150 series, IEEE 1588 PTP transparent clock is not supported on 40G interfaces (182897)
- SSO is hitful with subinterfaces. Protocols running on subinterfaces may experience session flaps during SSO. (188070)
- ACL drop counters cannot be cleared. (202905)
- when ACLs are configured with stats over port-channels, these stats are kept at the physical port level not at the port-channel level. (203131)
- With international and HW-REV-01 images and the introduction of "ip access-group <access-list name>" other configuration commands beginning with 'ip' that are configurable in the global config mode are not accessible in "router bgp" mode.

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

- Changes made to running-config after a config session is created might be reverted when the config session is committed which contains changes in similar or related areas. Use 'show session-config diff' inside the config session before committing to verify what changes will be made. (226850)
- The uptime in the output of "show module" may not match the uptime as shown by "show version" or "show uptime". The uptime in "show module" represents the time the module has been up since the last change in the "Status" column in the output of "show module", which differs from the meaning of uptime in "show version" and "show uptime". (248230)
- To avoid interoperability problems when using strong SNMPv3 encryption algorithms, follow the following minimums:
 - When using AES192 for privacy, use a minimum of SHA224 for authentication.
 - When using AES256 for privacy, use a minimum of SHA256 for authentication. (268290)
- Mirroring to an MLAG port channel is not supported. (276973)
- Copy commands with sftp: or scp: URL do not work via eAPI by passing password through the "input" parameter. SSH keys have to be setup to run the commands without a password. (283716)
- EOS support alphanumeric VLAN names, however ODL gives an error when reading the leaf /interfaces/interface/routed-vlan/config/vlan if the name contains alphabetic characters.

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

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

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

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

SKUs Affected: DCS-7260CX-64-F, DCS-7260CX-64-R, DCS-7260CX-64-SSD-F, DCS-7260CX-64-SSD-R, DCS-7260QX-64-F, DCS-7260QX-64-R, DCS-7260QX-64-SSD-F, DCS-7260QX-64-SSD-R, DCS-7280CR-48-F, DCS-7280CR2-60-F, DCS-7280CR2A-60-F, DCS-7280CR2K-60-F, DCS-7280QR-C72-F, DCS-7280QR-C72-M-F, DCS-7280QR-C72-M-R and DCS-7280QR-C72-R

- For static tunnel interfaces, duplicate or conflicting GRE tunnels are tunnels that have the same source, destination, tunnel mode (GRE) and key configurations as an existing tunnel. A duplicate tunnel is always in the down state. When there are duplicate static tunnel interfaces and a config-replace is executed, it is non-deterministic which one of the conflicting tunnels will be in the up state. (373135)
- When 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)
- Egress MAC ACL logging is not supported (433169)
- 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)
- SSH may fail when sFlow and Management SSH have the same QoS DSCP value configured, and sample rate for sFlow is set to dangerous. (434763)
- Mirror-On-Drop does not monitor IP packets with an invalid IP version. (443839)
- 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)
- "show mac/ip/ipv6 access-lists summary" now excludes remarks from "total rules configured". (452368)
- 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)
- EOS-2GB.swi is no longer available. (475680)
- Installed licenses without corresponding licensing features configured might not be visible from the CLI (504793)

- CVE-2019-16905: The version of openssh shipped in EOS does not enable support for XMSS, so EOS is not affected by CVE-2019-16905. (511376)
- CVE-2020-1968: EOS is susceptible to CVE-2020-1968. The workaround is to only use ephemeral (DHE) ciphersuites. (519024)
- If Octa, OpenConfig or TerminAttr are started under low memory conditions, they can cause a spike in CPU usage before unexpectedly restarting. (534841)
- Configuring NTP to serve clients in more than 500 VRFs fails with a syslog message "ntpd: Too many sockets in use". (543116)
- When the "tunnel destination" address is configured as a broadcast address for static tunnel interfaces then there could be packets drops during decapsulation at the tunnel endpoint. (544032)
- For a given interface, multiple simultaneous MAC address flush requests on different VLANs may result in clearing all MAC addresses on that interface. This is an optimization to support a higher scale of VLANs and interfaces. To disable this optimization, use the "mac address-table flushing interface vlan aggregation disabled" CLI configuration. (544712)
- The "dot1x re-authenticate" command is not supported for failed supplicants. Workaround is to flap the interface. (567028)
- Removing "vlan" key-field from custom TCAM profile for ingress security ACL feature causes ACL application on SVIs to fail.

To workaround this issue, add "vlan" key-field to custom TCAM profile for ingress security ACL features. (570680)

SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F

- In SSO redundancy mode, the switch may take longer (more than 20 minutes) to be declared switchover ready. This is to accommodate for all agents to become ready for switchover. (572615)
- When unsupported transceiver support is unlocked using a key, the ports that are in disabled state before unlocking will continue to be disabled.

The workaround is to remove such transceivers and re-insert them after unlocking the unsupported transceiver support. (574146)

- The 'logging level all' CLI expands to individual commands in the configuration, preventing dynamic updates of the logging list as part of EOS upgrades. (579047)
- TACACS+ authentication and authorization through SSH does not work with Pulse Secure server because EOS sends different port names in authentication ("ssh") and authorization (actual vty) requests which are rejected by the server. (583678)

- Sampled flow tracking exports IPFIX flow records with egress interface as 'discard' for flows on pseudowire patched interfaces. (603916)
- A CPU internal error (BIT30 SyncFlood BIT11) may cause an unexpected reboot (623392)

SKUs Affected: DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7050SX3-48C8-F, DCS-7050SX3-48C8-R, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050SX3-96YC8-F, DCS-7050SX3-96YC8-R, DCS-7050TX3-48C8-F, DCS-7050TX3-48C8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-36S-F, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-36S-F, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F, DCS-7280PR3K-24-F, DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R, DCS-7280SR3K-48YC8-F and DCS-7280SR3K-48YC8-R
- IP Locking does not support supervisor stateful switchover (SSO). (631563)
- Dot1x agent incorrectly accepts larger than advertised values for idle timeout and session timeout RADIUS response attributes (639283)
- /mnt/flash cannot be used as the log archiver destination on systems formatted as VFAT. To check if a product is formatted as /mnt/flash, from bash, run "mount | grep flash", and look for "type vfat" or "type ext4". If you see "type vfat", then that product's /mnt/flash storage is formatted as VFAT. (642415)
- locator-led fantray & locator-led powersupply CLI commands do not operate successfully on Modular Systems and may report "This LED cannot be used as locator-led target". (652512)

Series Affected: 7368, 7388, CCS-750, DCS-7300X, DCS-7300X3, DCS-7500R, DCS-7500R2, DCS-7500R3 and DCS-7800R3 Series
- "mac security profile" is not supported on L2 subinterfaces. (680967)

SKUs Affected: 7800R3-48CQM2-LC, 7800R3K-36DM-LC, DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280CR2M-30, DCS-7280CR2M-30-F, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- When Path MTU discovery is enabled under static tunnel interfaces using "tunnel path-mtu-discovery", the discovered MTU is only used for software forwarding in the slow path. When MTU is configured under static tunnel

interfaces using "mtu" config, the configured MTU value is used for hardware forwarding. (681101)

- OpenConfig does not support BFD configuration for VLAN interfaces. BFD configuration under VLANs must be removed for any gNMI set operation to succeed. (682920)
- If RFC2544 test is started on a sub-interface and the parent interface has "traffic-loopback source system device mac" configured, the test packets will get into infinite loop.

Unconfigure traffic-loopback on the parent interface using "no traffic-loopback source system device mac" to stop the traffic loop. (683129)

- If the AVA switch sensor ("monitor security awake") is enabled along with the TerminAttr (CloudVision), the TerminAttr agent will also process IPFIX packets sent to the CPU, and the "Traffic flows" feature on the CloudVision portal is automatically enabled and can cause additional CPU usage. If "Traffic flows" feature is not required on CloudVision, disable TerminAttr IPFIX processing by adding "-ipfix=false" argument to "exec /usr/bin/TerminAttr" under "daemon TerminAttr" configuration. (688791)

Series Affected: CCS-720XP Series

- EOS images are optimized during installation on selected platforms. The optimized EOS image saved on non-volatile storage may be incompatible with other platforms. An attempt to install an incompatible SWI will result in the following warning: "Next image does not support required optimization". Only install full EOS images or compatible optimized EOS images. (691810)
- From EOS release 4.29.2F, the 64-bit image of EOS supports "monitor security awake" or AwakeSensor.swix/NDRSensor.swix on platforms with at least 8GB memory. EOS releases prior to 4.29.2F, the 64-bit image of EOS does not support AwakeSensor.swix on any of the platforms. Use the 32-bit EOS version instead. (719835)
- System will not boot if a cable connected to an ethernet switch is connected to the console interface of the system. All system LEDs may remain red and the console cable will subsequently be unresponsive, until the cable connected to an ethernet switch is removed from the console interface of the system. (724181)

SKUs Affected: DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7050SX3-48C8-F, DCS-7050SX3-48C8-R, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050SX3-96YC8-F, DCS-7050SX3-96YC8-R, DCS-7050TX3-48C8-F, DCS-7050TX3-48C8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-36S-F, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-36S-F, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R,

DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F,
DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R,
DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F,
DCS-7280PR3-24-M-F, DCS-7280PR3K-24-F, DCS-7280SR3-40YC6-F,
DCS-7280SR3-40YC6-R, DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R,
DCS-7280SR3E-40YC6-F, DCS-7280SR3E-40YC6-R, DCS-7280SR3K-48YC8-F,
DCS-7280SR3K-48YC8-R, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R

- Inbound packets destined for local agents on the switch (e.g., LACP, STP) will not be decoded properly when the "tcpdump" command listens on the "any" interface.

The workaround is to direct the "tcpdump" command to listen on the interface on which the packets are being received. (730292)

- RAIL is not supported on Trident3X1 and Trident3X2 chipset based products (736685)
- PTP traffic cannot be captured using a monitor session on PTP enabled ports for 750X devices configured as a PTP Boundary Clock (738554)
- Smbus card agent may contain assertion errors when restarting due to card removal. There is no effect on service. (740093)
Series Affected: 7368 Series
- Removing one or more rule(s) with sequence numbers which causes a hole in ACL sequence number, i.e., removing rules from middle of an ACL and, re-adding ACL rules (with auto-generated sequence numbers) can cause sequence numbers to be re-used.
This series of ACL operations will cause an "Duplicate sequence number" error to be emitted in CLI.

To workaround this issue, if a rule in middle of ACL needs to be removed then,
1. Copy the rules which appear after the rule of interest,
2. Remove all rules from the rule of interest till the end of ACL,
3. Re-add rules copied in step 1. (743473)

- * network-instances/network-instance/policy-forwarding/policies/
policy[policy-id=POLICY-ID]/rules/rule[sequence-id=SEQUENCE-ID]/ipv4/config/
source-address
* network-instances/network-instance/policy-forwarding/policies/
policy[policy-id=POLICY-ID]/rules/rule[sequence-id=SEQUENCE-ID]/ipv4/state/
source-address
configures/displays at most 1 source-address prefix per rule, despite EOS CLI supporting >1 source-address prefix per rule.

As a workaround, either use OpenConfig exclusively for configuration/query/telemetry on IPv4 source-address prefixes, OR do not use OpenConfig for configuration/query/telemetry on IPv4 source-address prefixes if >1 IPv4

source address prefixes per rule have already been configured via EOS CLI.
(747080)

- * system/logging/remote-servers/remote-server/selectors/selector/facility
* system/logging/remote-servers/remote-server/selectors/selector/config/
facility
do not filter syslogs sent to the remote-server based on RFC5424 facility.
Instead, these two paths configure the RFC5424 facility for all syslogs,
regardless of original facility, sent to the remote-server (filtered by
having severity >= configured severity). (747096)
- AAA RADIUS servers configured with shared-secret profiles can only be used to
authenticate 802.1X supplicants. They can't be used to authenticate control-
panel login or CLI enable access. (750835)
- On dual supervisor products, supervisor card 1 is not always the active
supervisor when executing "reload all" or powering off/on the chassis.
(768242)
- JSON does not specify the precision of a Number in the specification. eAPI
can generate numbers up to 6-bit, which can be truncated by some JSON
clients with lower limit (e.g., some implementations have a max value of
 $2^{53}-1$). The only workaround is to use a different JSON deserializer (769194)
- TACACS+ request has a limit of 8150 bytes in body size. If a CLI command is
too long, authorization might get rejected. (770296)
- Workaround for the restrictions on the list
* /system/logging/remote-servers/remote-server[host]/selectors/
selector[facility severity]
is to update/replace/merge all /system/logging/remote-servers/remote-
server[host]/selectors/selector[facility severity] list items at once, such
that the end configuration is that all /system/logging/remote-servers/remote-
server[host]/selectors/selector[facility severity] list have (a) exactly 1
list item with identical facility and severity key values, or (b) no list
items. (774251)
- When a fifth non-unique TTL value is configured for Static Interface Tunnels
on Hardware Forwarded Platforms using the "tunnel ttl" config command, the
Static Interface Tunnel pipeline programming may not be completed due to
hardware limitations and the traffic may not flow through the Tunnel. Only 4
unique TTL value are supported across all the Static Interface Tunnels on
Hardware Forwarded Platforms. If TTL value is not configured for Static
Interface Tunnels using the "tunnel ttl" config command, then it implies
that the TTL value is zero. This zero TTL value is also a unique TTL value.
For Static Interface Tunnels to be operational, configure 4 unique tunnel
TTL values across all the Static Interface Tunnels on a Hardware Forwarded
Platform (782720)

- Sampling of flows through untagged L3 subinterfaces in the non-default VRF for features like IPFIX and Sflow is unsupported. (791426)
- EOS with version 4.29.0 or later generate a different user id from the same username. As a result, files on flash may show random user id instead of the expected user name. For example, a file created by the "admin" user under an older version may show that it is now owned by user "106904808" after upgrading to a newer version. The workaround is to run the following command from bash:

```
find /mnt/flash -uid 106904808 -exec sudo chown admin {} \; (798231)
```

- Private keys in PEM format should not be copied into EOS as a combined x509 certificate and key PEM file. The key will be ignored, without error message, and the complete file will be treated as a certificate and displayed via show commands. This will leak the private key when shown. (798664)
- Across all flow tracking features, when the number of configured exporters for a given tracker exceeds the maximum supported for the feature on the platform, the active exporters may differ after a reload or shutdown/no shutdown of the flow tracking feature.
Reduce the number of configured exporters to maximum supported for the flow tracking feature on the platform. (806831)
- A switch configured with both MACsec L2 protocol forwarding and configurable EAPOL DMAC will send the EAPOL packets with the configured DMAC to CPU instead of transparently forwarding it in the dataplane. (824887)
- When using "flow tracking mirror-on-drop" to generate sFlow datagrams containing dropped packets, the "drops" field is used to denote the number of dropped packets, instead of the number of drops by the policer or sFlow agent as is intended for that field. (834493)
- When PTP boundary clock is configured, incoming PTP packets on a PTP-enabled access switchport that contain a VLAN other than the one configured on the access port will be sent to CPU. (837200)
Series Affected: CCS-750 Series
- In a MLAG VXLAN deployment with anycast gateways, ARP synchronization between the MLAG peers may result in packet drops in the CoPP "copp-system-selfip" queue and adversely affect any application traffic over that queue. (881798)
- SSH login may fail with broken pipe error when "authentication protocol password keyboard-interactive" is configured and empty password login is allowed at the same time.
Disallowing empty password with either "no aaa authentication policy local

allow-nopassword-remote-login" from the config mode or "authentication empty-passwords deny" from the management ssh submode will fix this issue. (899698)

- L3 routes resolved over multiple next-hops (ECMP/UCMP) may remove otherwise reachable next-hops from their via set when more than one next-hop in the group is resolved over the same L3 interface/are in the same subnet, but only one peer in the subnet/in the VLAN has a BFD session with the router. Workarounds include removing the BFD session with that singular peer, or adding an additional BFD session with one or more other peers on that L3 interface/subnet. (901915)
- Arista Networks Product Bulletin for EOS Release 4.29.1F incorrectly included DCS-7280DR3A-36-F, DCS-7280DR3AM-36-F, DCS-7280DR3AK-36-F in the New Platform Capabilities section. DCS-7280DR3A-36-F, DCS-7280DR3AM-36-F, DCS-7280DR3AK-36-F were introduced in 4.30.0F. The Arista Networks Product Bulletin for EOS Release 4.29.1F was revised on January 15, 2024 to correct the documentation. (902473)

Series Affected: DCS-7280R3 Series

- If a VRF selection policy and a match rule is created using CLI first, and the rule is updated via gNMI, the resulting VRF selection policy configuration may have unexpected additional rule. Workaround is to use only either CLI or gNMI instead of both. (910953)
- SNMPv2/v3 requests for data in VRFs as a context is limited to VRFs with maximum length 32.

The workaround is to limit such VRF names to a maximum length of 32 characters. (924461)

- In a VXLAN routing setup, the VxlanSwFwd agent restart time may be high (1 to 2 minutes). During this time, new ARP or IPv6 neighbor entries may not be learned over VXLAN. (924494)
- In an EVPN VXLAN deployment using IPv6 overlay routing, if a anycast gateway is configured with 'ipv6 address virtual' and a host generates a DAD solicitation for the link local corresponding to the switch virtual MAC address configured for the anycast gateway, the switch won't respond to defend the link local address. This may result that the link local address corresponding to the switch virtual MAC to get advertised by the host. (953497)
- On a 7368 system with a 7368X4-SC switchcard, a combination of 4 plugged power supplies and 2 or fewer linecards, can lead to the system not powering back on after reload. It might require removal/insertion of the power supplies for it to power on.

To avoid running into this, the recommendation is to use a maximum of 2 PSUs on any 7368 chassis running just with 1 or 2 LCs. (957422)

Series Affected: 7358 and 7368 Series

SKUs Affected: 7289

- OpenConfig gNMI subscriptions to paths that contain wildcards and module-prefixed identityref list keys will return no data for those paths. The workaround is to not module-prefix identityref keys in subscription paths that contain wildcards. (975123)
- BGP extension may be missing with egress sFlow if the ingress interface is a Port-channel (1041816)
- "SSL_PROFILE_FEATURE_UNSUPPORTED" syslog messages can continue to be generated when an application is associated with an SSL profile that has configuration the application does not support even after the SSL profile is disassociated from the application.

After some time the syslog messages will stop. However, as an immediate workaround the SSL profile can be deleted. (1066555)

IPSEC

- We do not support NAT on Ipsec autoVPN in this release. (564431)
- When `ethernet untagged allowed` is configured under `ip security` config, TCP MSS clamping may not apply to packets encapsulated over a VXLANsec tunnel. (685484)
Series Affected: DCS-7280CR3 Series
- When a tunnel interface is shut and IPsec profile is removed from the tunnel interface, ipsec profile can't be deleted from the config. The workaround is to delete the tunnel and delete the ipsec profile. (994463)
- Changing the MTU of a IPsec tunnel interface to a value higher than 1982 causes the tunnel interface to go down and stay in the "down" state even after lowering the MTU.

After setting the MTU to a value equal to or lower than 1982 a "shut/no shut" sequence can be performed on the tunnel interface in order to bring it back to the "up" state. (1008426)

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

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

- MACsec proxy for front panel ports requires replay protection to be disabled on the MACsec profile, otherwise some packets will be dropped due to sequence error. (339533)
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
- Clause 37 (BASE-X) auto-negotiation is not supported on SFP25 ports. (492271)
SKUs Affected: DCS-7010TX-48-F and DCS-7010TX-48-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
- InPktsNoSCI counter increments for a MACsec frame with AN that is not in use by the receiver. InPktsSANotInUse should have incremented instead. (532869)
SKUs Affected: 7500R2AM-36CQ-LC, 7500R2M-36CQ-LC, 7800R3-48CQM-LC, DCS-7280CR2M-30, DCS-7280CR2M-30-F, DCS-7280SRAM-48C6, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R
- Interfaces with BCM82391 PHY do not support auto-negotiation. (537530)
SKUs Affected: 7500R-8CFPX-LC, 7500RM-36CQ-LC, DCS-7050CX3M-32S, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- QSFP to SFP adapters are not supported on front panel interfaces ethernet25/1-36/1. To work around this, use interfaces ethernet1/1-24/1 instead. (625046)
SKUs Affected: DCS-7280CR3-36S-F and DCS-7280CR3K-36S-F
- MACsec frames may be dropped when the sci setting in the "mac security profile" does not match the peer's sci setting. (664596)
SKUs Affected: 7388-8D, 7800R3-48CQM2-LC, 7800R3K-36DM-LC, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

- MACsec frames with SL less than 5-bytes are silently dropped on ingress (714497)
 SKUs Affected: 7800R3-48CQM2-LC, 7800R3-48CQMS-LC, 7800R3K-36DM-LC, DCS-7170B-64C-F, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R
- MACsec frames with SL less than 24-bytes are silently dropped on ingress (714597)
 SKUs Affected: 7500R-8CFPX-LC, 7500RM-36CQ-LC, DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- Cable diagnostics test (TDR) on a BASE-T port may report inconclusive result if link partner is configured at forced 100M speed. (716208)
 Series Affected: CCS-710P, CCS-720XP, CCS-750 and DCS-7050TX3 Series
- The MACsec inPktsNoTag counter includes the non-standard MACsec inPktsCtrl counts. inPktsCtrl (non-standard) counts packets that are parsed as control packets and hence forwarded to the uncontrolled port. For instance, EAPOL (MKPDU) and possibly LLDP or Pause/PFC if bypassed. (735450)
 SKUs Affected: 7388-8D, 7800R3-48CQM2-LC, 7800R3-48CQMS-LC, 7800R3K-36DM-LC, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R
- The MACsec inPktsNoTag counter includes the non-standard MACsec inPktsCtrl counts. inPktsCtrl (non-standard) counts packets that are parsed as control packets and hence forwarded to the uncontrolled port. For instance, EAPOL (MKPDU) and possibly LLDP or Pause/PFC if bypassed. (735455)
 SKUs Affected: 7500RM-36CQ-LC, DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- 50G-2 autonegotiation is not supported on interfaces with the CRT50216 PHY. (755659)
 SKUs Affected: 7358-16C, 7368-16S, 7368-4D, 7368-4P, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3K-48CQ-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4A, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4A and DCS-7280CR3K-96
- 100G-4 autonegotiation with FEC disabled is not supported on interfaces with the CRT50216 PHY. (796531)
 SKUs Affected: 7368-16C, 7368-16S, 7368-4D, 7368-4P, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3K-48CQ-LC,

DCS-7060DX5-64S-F, 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 and DCS-7280CR3K-96-F

- ``show interface interactions`` will not display any information for interfaces with PHY of type BCM56070-TSCM when there is no external PHY on the L1 topology. (907292)
SKUs Affected: CCS-750X-48ZP-LC
- BASE-T management ports do not support operation in non-negotiated 100M/10M mode, despite accepting the associated configuration commands. (972448)
- The link may flap once before coming up through errdisable recovery. (998012)
SKUs Affected: CCS-720XP-24ZY4-F, CCS-720XP-96ZC2-F and DCS-7010TX-48

Layer 2

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

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

- An MLAG switch cannot be used as an EVPN multi-homing provider edge. When a switch has MLAG enabled, MAC addresses from ports with an ethernet segment configured are advertised through BGP as if they were single-homing, and local ethernet segments are not advertised through BGP. (397867)
- Even with fallback to unprotected traffic configured, traffic loss in the order of milliseconds may be seen during initialization (491232)
- Upon programming a large number (~250) of VTEPs in a flood set at once, the L2Rib agent may restart. Workaround is to program them in batches, waiting a few seconds between each batch. (544051)
- 802.1X is supported on LAG member ports using EAPOL authentication. (544556)
- The SyncE wait-to-restore (WTR) timer is armed when administratively shutting / unshutting an Ethernet interface. This is inconsistent with disabling / enabling SyncE on the interface in which case the WTR timer is not armed.

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

- When an EEC has redundant links traceable to the same reference (e.g. ITU G. 781 section 5.13.2 "'bundles' of ports"), and the quality level of the reference transitions from a higher quality to a lower quality (e.g. QL-PRC to QL-SSU-A), the EEC may briefly enter a holdover condition for a few seconds. (687326)
- Synchronous Ethernet is not supported for 10G BASE-T transceivers. (689493)
- Dot1x session move is not supported if the supplicant changes authentication methods during the move. (705466)
- There is no Accelerated Software Upgrade(ASU) support for 802.1X dynamic ACL assignment feature. (725100)
- Loop protection is not intended to be used along with Spanning Tree Protocol (STP) (745522)
- Idle-host idle timeout only applies to supplicants which have been authenticated. (747150)
- The following commands

```
"aaa authentication dot1x default group <GROUP1> group <GROUP2> ... "
```

```
"aaa accounting dot1x default start-stop group <GROUP1> group <GROUP2> ... "
```

takes redundant groups without any warnings, i.e,
"aaa authentication dot1x default group group1 group group1 group
group1 ... " is considered as a legal CLI command, but internally we would
be configuring only one group i.e group1 for 802.1X authentication. This is
a limitation of the CLI and does not affect the functionality in any manner.
(773613)

- LAG interface counters will have constant counts discrepancy compared to the aggregate interface counts of its member interfaces when there is a link flap event on LAG interface. (785073)
- In VXLAN routing deployment, when using EVPN Symmetric IRB to reach a destination host, the remote router IP address in the traceroute address may be arbitrarily set to any interface IP address in the VRF of the destination VTEP. (815530)
- If a RADIUS server's hostname is long and crosses a certain character limit, the output of 'show dot1x radius' will be misaligned. (820586)
- IP Phones doing untagged EAPOL authentication and sending tagged traffic may not work properly when 802.1X MAC based authentication is disabled. (841592)
- An idle MBA supplicant that is in DISCONNECTED-CACHED state may not be cleaned up when caching is disabled and link comes back up.
A workaround for this is to configure "dot1x timeout idle-host TIMEOUT seconds". This config would cleanup the idle supplicant after the configured timeout. (854465)
- LLDP packets from unauthenticated supplicant do not trigger MAC based authentication on 802.1X configured interface with "dot1x mac based authentication" enabled. (860140)
- Disabling EAPOL using "dot1x eapol disable" command doesn't disable the fallback mechanism and its default timeout as configured using "dot1x eapol authentication failure fallback mba" configuration command.
Possible workarounds:
 1. Use the fallback CLI with timeout 0. eg "dot1x eapol authentication failure fallback mba timeout 0"
 2. Use MBA always option, ie "dot1x mac based authentication always", which supersedes "dot1x eapol authentication failure fallback mba" CLI. (918141)
- Traffic for a Dot1x authenticated supplicant received on a non-Dot1x interface will cause the MAC address learnt for that traffic on the Dot1x interface to not get cleared due to idle-timeout until the traffic stops. (938650)
- If a supplicant receives a dynamic VLAN from AAA server and this VLAN is outside the allowed dynamic VLAN range configured on the switch, the switch will proceed to successfully authenticate this supplicant. However, the

supplicant traffic will be dropped, since the requested VLAN is not available on the switch. (948547)

- If Port Security is enabled on a 802.1X port, then the existing authenticated supplicants on the port will be logged off and they will be re-authenticated post logoff. (958954)
- If Port Security is enabled on a 802.1X port, then the existing authenticated supplicants on the port will be logged off and they will be re-authenticated post logoff. The delay between the logoff and the next successful re-auth (log in) might be as much as 1 minute. (959675)
- In a IPV6 VXLAN deployment, a system log entry of type IPV6_UNDERLAY_UNSUPPORTED may be emitted after the device is reloaded while having remote IPv6 VTEP configuration.

These messages can be ignored. The configured IPv6 addresses will be learned shortly after the IPV6_UNDERLAY_UNSUPPORTED syslog message appear. (961492)

- When there is a continuous ARP churn on L3 servers, the Dot1x agent will be busy processing all the updates and may send IPs from stale ARP entries on the L3 server in Framed-IP-Address RADIUS attributes (992621)
- There is no support for MVRP with Rapid-PVST. It is not a supported configuration and may lead to undefined behaviour. (1035620)
- When a large number of EAPOL logoff packets are received on a single interface, the switch may drop some packets and not process the logoff for the corresponding supplicants. (1079779)

Layer 3

- ECMP is not supported for routes learned via RIP. (34773)
- DHCP relay is not supported when running virtual machines on EOS (101754)
- CLI does not detect TCAM exhaustion and automatically revert the change when a PBR policy is applied to an L3 interface in the "shutdown" state, since the actual hardware programming happens when the interface comes up. (122651)
- OSPFv3 support for multiple address families (RFC 5838) in EOS introduces support for IPv4 routing with OSPFv3. OSPFv3 for IPv4 AF in EOS requires configuring neighboring OSPFv3 IPv4 routers on the same link with primary IPv4 addresses in the same subnet. Adjacency is established in OSPFv3 IPv4 AF even if the neighbor's primary IPv4 address is in a different subnet but routes through

the neighbor on a different subnet will however not be installed. This limitation may be removed in a future release. (140234)

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

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

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

- IPv6 addresses are not supported as tunnel source and destination. (283912)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289217)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289224)
- When BFD is configured with per-link mode rfc-7130 for an IPv6 link-local address on a Port-Channel and the Port-Channel transitions to the Down state, the Port-channel may not transition back to the UP state.

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

- On Macro-Segmentation Service (MSS) deployment with L3 firewalls, a flow entry that fail to get programmed due to lack of TCAM space, is not programmed when TCAM space becomes available at a later time. The workaround is to "shut" and "no shut" on the direct flow config. (372232)
Series Affected: DCS-7050TX, DCS-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)
- Layer 3 protocols incorrectly see Pseudowire local connectors as layer 3 interfaces. As a result, reachability to interface's attached IP networks may be advertised, depending on the routing protocol configuration, irrespective of pseudowire connector's operational state. (429818)
- Destination (TEP) or include hop entry for a RSVP tunnel can be TE router ID of a router or IP address on any interface that has an IGP neighbor. It cannot be any other interface address, like a loopback interface address which is not TE router ID or a stub interface address. (453019)

- Nexthop group tunnel counters does not work, when tunnel counters is enabled using 'hardware counter feature mpls tunnel' command. (474078)
- VXLAN features are not compatible the 'ip hardware fib next-hop sharing' configuration. Configuring VXLAN routes may result in identical FECs not being shared and potential misforwarding after changes to FECs are made. (477471)
- The config 'ip hardware fib next-hop multi-path maximum' does not restrict the size of installed next-hop group ECMP FECs (498484)
- CLI configuration of a static route with nexthop pointing to self ip address is not rejected. (514894)
- NAT dynamic profiles are not supported on multi-chip and modular system. (525213)
Series Affected: 7300X3, CCS-720XP and DCS-7300X Series
- BFD over Vxlan does not work for IPv6 underlay. (526897)
- Rsvp autobandwidth feature does not work for single hop tunnels. The workaround is to configure Rsvp explicit-null on the egress router for single hop tunnels. (571177)
- VARP VTEP IP configuration is not supported on AP aggregation VTEPs

It is recommended that either EVPN VXLAN symmetric or asymmetric IRB be used for distributed VXLAN routing in campus environments. (571517)

- IS-IS Extended administrative group (EAG) sub-TLV is not supported. If a system in the network advertises EAG, Arista router will ignore the sub-TLV. (628807)
- The 'show ip hardware fib ecmp resilience' CLI may sometimes incorrectly report an unordered resilient ECMP as ordered. (685318)
- If a preferred IPv6 source address for an IPv6 kernel route is obtained from a local Port-Channel interface that is used in a "software forwarding hardware offload route local interface INTERFACE" command, and the interface has a IPv6 BFD neighbor configured, and the Port-Channel interface is subsequently shut down, the preferred IPv6 source address of the kernel route is not updated as expected.

The workaround is to remove the "bfd ipv6 neighbor" command from the Port-Channel interface and re-apply the command. (698716)

- If a "software forwarding hardware offload route rcf FUNCTION" command is configured under "router kernel/address-family" sub-mode and FUNCTION returns True for a connected route, then kernel traffic to that prefix is

forwarded in hardware as expected. But if the connected route later goes away for some reason (e.g., the interface is reconfigured) and is replaced in the FIB by a route for the same prefix but learned from a routing protocol peer, the kernel route is programmed for software forwarding instead of hardware forwarding.

Restarting the KernelFib agent should cause the affected kernel route to be reprogrammed such that forwarding of kernel traffic to the prefix is offloaded to the hardware. (701746)

- Resilient ECMP may log erroneous "RECMF_CAPACITY_EXCEEDED" syslog during interface flap. (722339)
- CLI command show ip route may incorrectly display an unprogrammed route as programmed. (731198)
- BFD running over L3 interfaces can reduce the time it takes to shrink an ECMP when one of the ECMP member links fail. If all the ECMP paths use the same L3 interface (such as when using SVI interfaces) then the time to shrink the ECMP may take longer. (746361)
- The `show ip hardware fib ecmp resilience` command may display `VRF <name> not found` when given a non-default VRF. (764757)
- 'show tech extended dhcp-relay' might dump a partial output in some cases (774250)
- When an interface is inactive due to DAD (duplicate address detection) failure, disabling DAD using the 'ipv6 nd dad disabled' CLI configuration does not clear the failure state or allow the interface to become active. As a workaround, removing and re-adding the duplicate address after disabling DAD will clear the failed state. (774554)
- 'show ip hardware fib ecmp resilience' and 'show ip hardware fib ecmp resilience summary' commands may show status of unprogrammed resilient routes that have already been removed from FIB. (798614)
- IPFIX and sFlow BGP nexthop value cannot be accurately determined for ECMP flows going over the same egress interface that are encapsulated with implicit-null labels. There is no workaround for this issue (816528)
- When "ip virtual-router mac-address mlag-peer" is enabled in the MLAG environment, BFD is not supported on the MLAG peer interface. (818124)
- gNMI replace on /routing-policy/defined-sets/bgp-defined-sets/community-sets/community-set/config/community-member from one representation of standard community number to another might result in redundant CLIs even though the two representations are effectively the same (826724)

- "MIXED" prefix-sets are not supported in OpenConfig in policy constructs. The workaround is to use separate prefix-sets for IPv4 and IPv6 prefixes. (858818)
- Deleting an interface with a large number of subinterfaces can result in "MTU less than device minimum" log messages. There are no side effects from this, and functionality is unaffected. (859295)
- ECMP path shrink or fallback to the backup path may be slow when BFD session state goes down if there are multiple next hops with different address families within the ECMP group. (867829)
- Setting accumulated IGP (AIGP) metric of "-0" using the "set aigp-metric" route-map clause is not supported and results in a traceback in the EOS CLI.

Workaround is to use "0" (887974)

- The show ip fib hardware ecmp resilience command may continue to display unprogrammed routes with reason "Route not programmed". These are routes that were previously resilient, but were subsequently deleted/unprogrammed. (903779)
- The next hop in the extended router data of an sFlow datagram will not be populated if the destination IP address of the sampled packet is of a different address family than the next hop IP address (939695)
- 'show ip dhcp relay counters' and 'show ipv6 dhcp relay counters' count both IPv4 and IPv6 DHCP packets (952513)
- Added support for Extended Administrative Groups in the 0-511 range (992822)

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)
- IP ACLs are not supported for control plane route filtering for VRF leaking (536329)
- Match dynamic prefix-list clauses in route maps are only supported for BGP inbound and outbound route maps. (564552)
- CVP will not monitor multiple BGP peers that are configured using the same peering IP address. (588954)
- BMP, SNMP and EosSdk will not report multiple BGP peers configured with the same BGP peering address. (602657)
- When default-originate is configured, default route will not be advertised to a neighbor if a prefix-list configured directly for this neighbor denies the default route. (648587)
- When multiple single-hop EBGp sessions are configured using point-to-point physical interface IP addresses between the same pair of switches, some of the BGP neighborship may get established using the other peer's IP address. When any BGP neighborship comes up using such wrong IP address from cross-link, the BGP neighbor configured on the other link will have connection establishment failure with Notification error: 'connection collision resolution'.

To workaroud this issue, either use 'update-source' option for the

neighbor or explicitly enable 'no neighbor <addr> route-to-peer' for all the single-hop EBGP peers using point-to-point physical interface IP address. (686390)

- BGP session tracker may not detect quick BGP peer flaps and may not error disable interfaces (687998)
- For the RT membership AFI/SAFI, RT membership routes advertised by a given peer are also reflected back to the same peer. i.e. there is no echo suppression of the RT membership NLRI. (717330)
- BGP TCP-AO is not supported on cEOS. (762218)
- With "next-hop resolution v4-mapped-v6 translation", IPv4-mapped IPv6 next hop address of IPv6 Unicast NLRIs is translated to IPv4 address for next-hop resolution. However, such paths are not advertised to outbound peers if the next hop is unchanged.
A workaround is to explicitly configure outbound route map or RCF for the peer that sets the next hop to be path's next hop. (790764)
- BGP-LU routes directly redistributed into IS-IS using command 'tunnel source-protocol bgp ipv4 labeled-unicast' under 'router isis' mode, uses MED as metric instead of AIGP.

To workaround this issue, IS-IS metric value can be derived from AIGP metric using RCF policy. Note that the 'aigp.metric' attribute value, used by RCF at redistribution point of application, is same as AIGP metric path attribute value received from the peer and doesn't account for the next-hop's IGP metric cost. (791558)

- The CLI "show bgp debug policy" will indicate when an RCF does not explicitly return a value causing a route to be denied. (861041)
- Routes where the destination prefix covers the next hop of the route are not installed in RIB and consequently in the FIB. (898293)
- Debugging RCF policy with the "show bgp debug policy" command may not work with additional paths. (921294)
- When next hop is updated using either outbound route-map or outbound RCF policy for the labeled-unicast peer, the advertised LU label is not updated correctly.
To workaround this issue, configure 'neighbor <peer-address> next-hop-self' for such labeled-unicast peer. (1001447)
- If a link bandwidth extended community is being introduced outbound, route-map or RCF policy may not match on the link bandwidth value being advertised over the wire.

As a workaround, match on the link bandwidth value using inbound policy on the received device. (1055916)

- In multi-agent mode, the conditional default route originate feature is limited to specific types of routes such as BGP, IS-IS, RIP, static and connected routes. For other sources of routes such kernel routes, DHCP routes etc. the conditional default route originate functionality does not work in multi-agent mode. (1098356)

Rib

- PBR recursive resolution for nexthops is not supported for packets routed in software (like packets with IP options) (82077)
- 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)

- Ribd agent may restart while running with Internet scale routes and with a large number of (50+) Rib-Out groups. Such a high scale is not supported with single-agent mode.

To workaround this issue, either reduce the number of Rib-Out groups by checking the output of 'show bgp update-groups' or upgrade to multi-agent mode which can support very high scale. (558044)

- Static routes with administrative distance of 255 are installed in the RIB and eligible for redistribution into other protocols with "redistribution config". The routes are also installed in the FIB for programming in hardware.

To deny FIB programming of routes, the RIB route control configuration policy in the form of "rib ipv4|ipv6 fib policy" can be used.

To prevent RIB redistribution of routes, the routes can be filtered out as part of the "redistribution" policy into the target protocol. (663662)

- There may be a behavior difference in OSPF between single agent and multi agent mode when both OSPF and BGP are configured in a non-default VRF. In

single agent mode, a route that would come from an external type 5 LSA might be dropped where in multi agent mode, the route would not be dropped.

The workaround to prevent loops is to follow RFC4577 (use of the DN-bit) (774087)

MetaMux

- MuxIn Interfaces from non-EB profile, do not get enabled when sourcing directly from different MetaMux instance.
Atleast one another MuxIn Interface in that non-EB profile, must be sourced from either front panel interface or, switch interface from J2C, to make it enabled (982810) (1)
SKUs Affected: DCS-7130LBR-48S6QD, DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R

(1) This item is in two or more sections on this document

MetaWatch

- The ``cable-delay`` command under the ``config-hardware-clocking-port-pps-in1`` CLI configuration mode does not offset the MetaWatch timestamps.

To apply a delay/offset, use the ``timestamp offset`` command under ``config-metawatch`` CLI mode as a workaround. (618981)

- Disabling the time-sync functionality will not stop MetaWatch from passing traffic or synchronizing time.

Workaround for stopping traffic is to disable MetaWatch, and to stop the synchronization action is to switch the time-sync reference to 'free-running'. (743245)

- When using deduplication, any inserted VLAN tag (in the original packet) that is not removed prior to hashing will disable the IPv4 TTL and header masking. (935469)
- An Arista FDK application with Tscore cannot be configured alongside MetaWatch, even if the configurations are disabled.

To use an Arista FDK application with Tscore, all MetaWatch instance configurations must be removed (not just disabled).

To use MetaWatch, any Arista FDK applications with Tscore must be disabled. (1089489) (1)

MLAG

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

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

- TCAM profile switch in a MLAG switch may result in some packet duplication for broadcast or multicast packets (341353)
- With MLAG egress filter interlock configured, duplicate BUM packets may be seen on MLAG interfaces along with "%MLAG-4-PEER_ACL_STAGE_TIMEOUT" syslog when the MLAG peer-link state changes.

If duplicates are seen while exiting MLAG dual-primary detection, the workaround is to configure a large enough MLAG dual-primary recovery delay such that PEER_ACL_STAGE_TIMEOUT syslogs are not observed. If duplicates are seen when reloading a MLAG peer, ensure the MLAG reload delay is configured properly. (724152)

- On an MLAG device VRRP sessions configured on VLAN interfaces may briefly transition to the master state after device reboot even though the VLAN is not associated with any active L2 ports. The VRRP session will become Stopped once the MLAG reload delay timer has expired. Beyond unexpected log messages there is no impact of device operation. (799723)
- When MLAG subinterfaces are configured and a MLAG interface flaps, there may be traffic loss. Fast MAC redirection is not supported on MLAG subinterfaces. (1018070)

MPLS

- The MPLS agent may continuously restart on a router with high BGP LU Tunnel scale. (377939)
- BGP LU Tx MPLS routes corresponding to BGP LU advertisements with multiple labels are not supported. (395029)
- The control word configuration does not take effect for one-hop pseudowire tunnels (406396)
- When RSVP is configured with split-tunnels, the `show traffic-engineering rsvp tunnel` command may display a last sampled bandwidth value for the tunnel that is not exactly equal to the sum of the last sampled bandwidth value of all sub-tunnels. This can occur if the command is used while the values are actively changing as it is possible that some values in the output have not yet been updated while others have. (625866)

Multi-domain Segmentation Services

- Interface traffic-policies cannot match on VXLAN decapsulated bridge packets. (913467) (1)

(1) This item is in two or more sections on this document

Multicast

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

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

- After ASU reload, if the first PIM hello packet received from the neighbor gets lost and If the neighbor has a high hello query-interval, it will result in traffic loss. Workaround is to use default PIM hello query-interval and increase the PIM hello query-count. (341447)
- With scale greater than 5000 multicast IPV6 routes, show platform sfe mroute ipv6 <vrf> commands might not display the right number of routes when performing multiple addition/deletion operations across VRFs. This is applicable only if user level multicast forwarding agent is configured to be used instead of kernel. (378841)
- Bessd might unexpectedly restart on config replace with scaled setup over 30 VRFs and large number of PIM enabled interfaces. Work around is to reduce the scale. (398185)

- Pim "non-dr install-oifs" feature becomes disabled if the non-dr needs to route multicast traffic for any reason to the interface on which the feature was originally configured. This is how the feature is expected to work. The workaround is to setup the network in a way that the DR always becomes the assert-winner. However, if the non-DR becomes assert winner for some routes on an interface and feature becomes disabled, flap the pim sparse mode configuration on non-DR on that interface to attempt to get out assert winner state. (401672)
- Multicast routes are not created when a source is transmitting only fragmented multicast packets.
Switch to using SFE for multicast software forwarding. (605326)
- In certain EVPN Multicast deployments, the PEG designated router (DR) might send PIM Registers to the RP for sources located in the external PIM domain. Once the PEG DR receives register stops, it ceases to forward registers. This behavior can initially result in a few packets being duplicated. (630762)
- A device with 8GB RAM might be able to support a maximum of 20 VRFs or lesser depending on the memory requirements of other features configured on the device. A workaround is to reduce the VRF scale or use a switch with higher RAM for such a scale requirement. (707709)
- Multicast scale is limited on devices running with 4GB memory. Running `software-forwarding sfe` on these platforms can result in BessMgr agent crashing. (786687)

NAT

- SNAT will not be performed if the original source IP is not fully resolved, and DNAT will not be performed if the translated destination IP address is not fully resolved. For routed ports, a static or dynamic ARP entry must be present for the relevant IP address. For SVIs, both an ARP entry and a MAC table entry for the MAC address associated with the IP address must be present. (756211)
Series Affected: DCS-7130 Series
- Changing the forwarding profile between nat-vxlan and nat may cause the Nat agent to unexpectedly restart. (792534)
Series Affected: DCS-7170 Series
- When an interface configured with static NAT translations is taken out of the shutdown state, untranslated traffic will flow briefly. (805141)
Series Affected: DCS-7130 Series
- If static source NAT is configured on the OUTSIDE interface of an INSIDE/OUTSIDE pair of interfaces, configuring static destination NAT on the INSIDE interface will cause the static source NAT translation to fail. (821248)

SKUs Affected: DCS-7130-16G3-F (HwRev:C1), DCS-7130-16G3-R (HwRev:C1), DCS-7130-48G3-F (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48G3-R (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48L-F (HwRev:A2), DCS-7130-48L-F (HwRev:A3), DCS-7130-48L-R (HwRev:A2), DCS-7130-48L-R (HwRev:A3), DCS-7130-48LB-F (HwRev:A3), DCS-7130-48LB-R (HwRev:A2) and DCS-7130-48LB-R (HwRev:A3)

- NAT synchronization is not supported together with NAT overload. (862459)

SwitchApp

- IGMP snooping filters are not supported when SwitchApp over forwards IGMP reports. Reports are over forwarded if the "igmp flooding" CLI is enabled through the FPGA instance CLI, or if the FPGA profile does not support correct forwarding of IGMP reports. (601917)
- When a MLAG with DCS-7130 LB series devices is running SwitchApp application with L3 mid-latency profile, because of the pipeline bandwidth constraints, if the peer-links are in the same pipe as other ports and the receive traffic rate across all ports is high, MLAG control traffic across peers may be dropped causing MLAG relationship to fail.

The recommendation when using this profile is to have the peer-link ports from a dedicated pipeline.

The port to pipeline mapping is available via "show fpga switch mapping" command. (715712)

- Multicast boundaries do not filter unicast PIM encapsulated packets sent to the RP. This may result in unicast PIM traffic continuously reaching CPU in certain network topologies. (718224)
- When running the SwitchApp application as a PTP grandmaster, it will distribute its internal free running counter instead of system time, which corresponds to the time the SwitchApp has been started since 1 January 1970. (736982)
- When using SwitchApp with routing and MLAG, if link failure causes more than 64 ARP entries to be needed on the peer link then traffic may be dropped.
Design the network with care to ensure that this limitation does not occur for the relevant failure scenarios. (765647)
- MakoSharedResources agent may restart hitfully if MakoL3Unicast agent performs a hitful restart due to crashing during a hitless restart. The system will recover to normal operation without intervention. (800466)

- SwitchApp platforms are limited to 64 nexthops per L2 port. This includes the inter MLAG peer link. As a result MLAG systems where more than 64 hosts are reachable via MLAG ports may hit this limitation during failure scenarios where traffic for more than 64 hosts need to cross the peer link resulting in traffic loss. It is therefore recommended to design configurations to ensure that this limitation is not exceeded in the event of failures. (811102)
- When the hardware resources become full, new NAT rules will not be programmed in hardware, and the system will log:

```
`%NAT-3-HW_RESOURCE_FULL: Hardware resources are insufficient to program all NAT rules`
```

Those rules will remain pending, even if hardware resources subsequently become available.

The workaround is to remove any pending rules (``show ip nat translation pending``) and re-add them to use the available hardware resources. (848105)

- A router running SwitchApp may emit the following error "MakoL3Unicast: %ROUTING-3-FEC_RESOURCE_FULL: Not all FECs are programmed in hardware" followed soon after by "MakoL3Unicast: %ROUTING-3-FEC_RESOURCE_NORMAL: All FECs are programmed in hardware." This is expected in scenarios where a number of IP next hops are learned in the routing table before their MAC addresses have been learned in the MAC address-table.

There is no workaround required - the "..._RESOURCE_NORMAL" message indicates that the condition has cleared and the router is now forwarding correctly to all destinations. (1009531)

- Routes can fail to be programmed in hardware and HW_RESOURCE_FULL errors can occur even though LPM table utilizations are as low as 25% of capacity. The limit will vary with the prefixes in use and may also vary between reloads of the switch.

Workaround: Reducing the number of prefixes at the relevant prefix lengths until the utilization below 25% will resolve the problem. (1049134)

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
- For frames that are routed and VXLAN encapsulated, underlay MTU is not used for MTU enforcement. Therefore the packets will be routed, encapsulated and transmitted without fragmentation. (254767)

- 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)
- The number of L2 subinterfaces which can be supported on a parent interface is 63.
On a given parent interface, there cannot be more than one L2 subinterface with the same forwarding vlan. (357946)
- 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)
- 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
- When sFlow and monitor session are both enabled in the system, the mirrored traffic may be incorrectly throttled to the sFlow copp policy rate. (604224)
- The 32-bit image of EOS does not support the 7170B series switches. Use the 64-bit version instead (646589)
- IP Multicast packets that are software forwarded are not subjected to egress static NAT. (771707)

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-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM and CCS-750 Series

- Configuring Authentication Failure VLAN to same value as 'Internal VLAN' or 'AAA-server assigned successful VLANs', may lead to un-desirable behaviour (376571)
- When hardware flow tracking is enabled, tcpControlBits information element in IPFIX flow records for VXLAN encapsulated TCP flows is incorrect. (416079)
- On receiving COA-REQ for changing VLAN for EAPOL supplicant, the NAS deliberately fails authorization for the first time in the new vlan. (449890)
- The per-interface configuration to ignore reauthorization timeouts from the AAA server ("dot1x timeout reauth-timeout-ignore") is not designed to work if manual reauthorization is requested from the command line. The reason for this is that manual reauthentication is designed to force the supplicant to contact the AAA server to retrieve the latest credentials. (453044)
- CoA request for VLAN change for EAPOL supplicant returns CoA-NAK. The VLAN does get changed when device is reauthenticated and the AAA server sends the new VLAN in an Access-Accept response. (459726)
- Hitless FPGA upgrades are not supported due to hardware limitations. Therefore FPGA images will not be upgraded during SSU. (474978)
 SKUs Affected: CCS-720XP-24Y6-F, CCS-720XP-24Y6-R, CCS-720XP-24ZY4-F, CCS-720XP-48Y6-F, CCS-720XP-48Y6-R and CCS-720XP-48ZC2-F
- Negation operator (!) and 'assigned' keyword are not supported in Radius NAS-Filter-rule(attribute-id 92) Attribute. (475207)
- Ip options qualifiers are not supported in the NAS-FILTER-RULE attribute if present in the Access-Accept or COA packet from the AAA server. If Ip options are present in NAS-FILTER-RULE, it may cause supplicant to fail authorization. (480659)
- WOL feature is currently only available for phone trunk port, but not for regular trunk port. (481468)
- Fallback to MBA authentication feature may not work as expected with Caching Auth feature. (482333)
- 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. It is recommended to configure the hold period to be lower than the timeout value for MBA fallback. (483995)

- In Mac-Based-Authentication in Dot1x, once a host's mac address gets authenticated with one vlan, it cannot change its vlan to get authenticated with a new vlan. (490696)
- 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)
- CoA request to change VLAN of multiple supplicants identified with NAS-PORT-ID may not work as expected. (495905)
- MBA for passive device can not be triggered by ARP/ICMP if IP locking is enabled (500439)
- If a switchport loses power, the credentials of the attached phone and PC will remain cached when authentication caching is enabled. However, when power is restored, the EAPOL enabled PC may not be reachable within the default reauthorization limit of 3 attempts and the supplicant will be removed as a result.
The solution in this situation is to increase the reauthorization limit using the per-interface dot1x command: dot1x reauthorization request limit 10 (500488)
- Limitation - 802.1X Web Authentication Implicit ACL feature doesn't work when IP Locking is configured on the interface. The workaround is to use 802.1X ACL based Web Authentication. (557986)
- PFC and pause generation are not supported when MACsec is enabled. (577310)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- Mirrored packets sent to a MACSec enabled destination interface are sent out unencrypted. (649877)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- Command "ptp free-running source clock system" is not supported. Although it may complete without errors, it will not take effect. (708316)
Series Affected: CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series
- PTP will not synchronize over interfaces configured with half-duplex link speeds. (740727)
- Phone sending LLDP packets before getting authenticated (by sending non-LLDP packets like DHCP) may not get classified as phones.

The workaround is to configure "dot1x protocol lldp bypass" (742933)

- "VTEP to VTEP bridging" is not supported. (791265)
SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F
- The system will not boot if a cable connected to an ethernet switch is connected to the console interface of the system. To recover the system, remove the cable from the console port and power cycle. (934172)
Series Affected: CCS-720DF, CCS-720DP and CCS-720DT Series
SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DF-48Y-2, CCS-720DF-48Y-2F, CCS-720DF-48Y-F, CCS-720DP-24S, CCS-720DP-24S-2, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-48S, CCS-720DP-48S-2, CCS-720DP-48S-2F, CCS-720DP-48S-F, CCS-720DT-24S, CCS-720DT-24S-2, CCS-720DT-24S-2F, CCS-720DT-24S-F, CCS-720DT-48S-2, CCS-720DT-48S-2F and CCS-720DT-48S-2R

CCS-750 and DCS-7060X4 Series

- When re-enabling a disabled power supply using `power supply enable module PowerSupply<n>`, the power supply may not re-enable.

The power supply can be recovered by removing and re-applying the input power to the power supply. (857791)

SKUs Affected: PWR-1511-AC-BLUE, PWR-1511-AC-RED, PWR-1511-DC-RED, PWR-400-DC-RED, PWR-500AC-F, PWR-500AC-R, PWR-501AC-F, PWR-511-AC, PWR-511-AC-BLUE, PWR-511-AC-RED, PWR-511-DC, PWR-511-DC-RED, PWR-745AC-F, PWR-747AC-F and PWR-747AC-R

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)
- Ingress ACLs on IPv6 packets do not work when IPv6 uRPF is configured and AlgoMatch is enabled. (512394)

DCS-7130 Series

- This platform only supports 64-bit versions of EOS. (615075)
- Several link flaps may be seen as the links are coming up. Debounce time CLI do not apply to Ethernet Interfaces on dropbear products. Link monitoring, Debounce and EOAM rely on PCS link status and not enabled on these since these interfaces use network PMA to set link status. The failure this reported here is a test issue. (633947)
- In the event of power loss, the switch will not report the correct reload cause. If power is interrupted briefly, the switch will report the reload as being caused by a watchdog. If the power loss lasts more than a few

seconds, the switch will report the reload as cause unknown. (700291)

Series Affected: DCS-7130 Series

- CL28 Autoneg is not supported. This means auto-negotiation on copper cables at speeds 10G or lower is not supported. (719438)

Series Affected: DCS-7130 Series

- When link loss is experienced on a MetaMux input interface, MetaMux may erroneously increment the error count for that MetaMux input interface. (723415)

Series Affected: DCS-7130 Series

- SwitchApp's published routing scale limits are approximations and simplifications so it is possible for a system to run out of internal resources for certain prefixes even though the overall scale is within the published limits. (764479)

- Maximum packet size supported by MetaWatch could be limited to 9214 bytes depending on platforms (775872)

SKUs Affected: DCS-7130LBR-48S6QD, DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R

- Use of unidirectional L1 configuration for a front panel Ethernet interface populated with a BASE-T type transceiver is not supported. Setting unidirectional configuration to such an interface may produce unexpected results, including traffic loss. Workaround is to use unidirectional configuration with non BASE-T type transceivers. (860391)

Series Affected: DCS-7130 Series

- Unidirectional link configuration is not available for front-panel Ethernet interfaces. This may lead to ambiguity in link signal status when the Ethernet interface is connected to a Switch interface that has unidirectional configuration applied. (902809)

SKUs Affected: DCS-7130LBR-48S6QD

- LLDP transmit does not work if the front panel Ethernet interface is layer1 sourcing an interface (except a Switch interface on the Switch FPGA application). Transmit Discard counters may increment on the interfaces as these packets are dropped. (941796)

SKUs Affected: DCS-7130-16G3, DCS-7130-16G3-F (HwRev:C1), DCS-7130-16G3-R (HwRev:C1), DCS-7130-16G3S, DCS-7130-16G3S-F (HwRev:D1), DCS-7130-16G3S-R (HwRev:D1), DCS-7130-48E, DCS-7130-48E-F (HwRev:A4), DCS-7130-48E-R (HwRev:A4), DCS-7130-48EH, DCS-7130-48EH-F (HwRev:A4), DCS-7130-48EH-R (HwRev:A4), DCS-7130-48EHS, DCS-7130-48EHS-F (HwRev:B1), DCS-7130-48EHS-F (HwRev:B2), DCS-7130-48EHS-R (HwRev:B1), DCS-7130-48EHS-R (HwRev:B2), DCS-7130-48G3, DCS-7130-48G3-F (HwRev:B1), DCS-7130-48G3-F (HwRev:B2), DCS-7130-48G3-R (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48G3S, DCS-7130-48G3S-F (HwRev:C1), DCS-7130-48G3S-R (HwRev:C1), DCS-7130-48L,

DCS-7130-48L-F (HwRev:A2), DCS-7130-48L-F (HwRev:A3), DCS-7130-48L-R (HwRev:A2), DCS-7130-48L-R (HwRev:A3), DCS-7130-48LA, DCS-7130-48LA-F (HwRev:A2), DCS-7130-48LA-F (HwRev:A3), DCS-7130-48LA-R (HwRev:A2), DCS-7130-48LA-R (HwRev:A3), DCS-7130-48LAS, DCS-7130-48LAS-F (HwRev:B2), DCS-7130-48LAS-F (HwRev:B3), DCS-7130-48LAS-R (HwRev:B2), DCS-7130-48LAS-R (HwRev:B3), DCS-7130-48LB, DCS-7130-48LB-F (HwRev:A2), DCS-7130-48LB-F (HwRev:A3), DCS-7130-48LB-R (HwRev:A2), DCS-7130-48LB-R (HwRev:A3), DCS-7130-48LBA, DCS-7130-48LBA-F (HwRev:A2), DCS-7130-48LBA-F (HwRev:A3), DCS-7130-48LBA-R (HwRev:A2), DCS-7130-48LBA-R (HwRev:A3), DCS-7130-48LBAS, DCS-7130-48LBAS-F (HwRev:B2), DCS-7130-48LBAS-F (HwRev:B3), DCS-7130-48LBAS-R (HwRev:B2), DCS-7130-48LBAS-R (HwRev:B3), DCS-7130-48LBS, DCS-7130-48LBS-F (HwRev:B2), DCS-7130-48LBS-F (HwRev:B3), DCS-7130-48LBS-R (HwRev:B2), DCS-7130-48LBS-R (HwRev:B3), DCS-7130-48LS, DCS-7130-48LS-F (HwRev:B2), DCS-7130-48LS-F (HwRev:B3), DCS-7130-48LS-R (HwRev:B2), DCS-7130-48LS-R (HwRev:B3), DCS-7130-96, DCS-7130-96-F (HwRev:A7), DCS-7130-96-F (HwRev:A8), DCS-7130-96-R (HwRev:A7), DCS-7130-96-R (HwRev:A8), DCS-7130-96E, DCS-7130-96E-F (HwRev:A7), DCS-7130-96E-R (HwRev:A7), DCS-7130-96L, DCS-7130-96L-F (HwRev:A2), DCS-7130-96L-F (HwRev:A3), DCS-7130-96L-R (HwRev:A2), DCS-7130-96L-R (HwRev:A3), DCS-7130-96LA, DCS-7130-96LA-F (HwRev:A2), DCS-7130-96LA-F (HwRev:A3), DCS-7130-96LA-R (HwRev:A2), DCS-7130-96LA-R (HwRev:A3), DCS-7130-96LAS, DCS-7130-96LAS-F (HwRev:B1), DCS-7130-96LAS-F (HwRev:B2), DCS-7130-96LAS-R (HwRev:B1), DCS-7130-96LAS-R (HwRev:B2), DCS-7130-96LB, DCS-7130-96LB-F (HwRev:A2), DCS-7130-96LB-F (HwRev:A3), DCS-7130-96LB-R (HwRev:A2), DCS-7130-96LB-R (HwRev:A3), DCS-7130-96LBA, DCS-7130-96LBA-F (HwRev:A2), DCS-7130-96LBA-F (HwRev:A3), DCS-7130-96LBA-R (HwRev:A2), DCS-7130-96LBA-R (HwRev:A3), DCS-7130-96LBAS, DCS-7130-96LBAS-F (HwRev:B1), DCS-7130-96LBAS-R (HwRev:B1), DCS-7130-96LBS, DCS-7130-96LBS-F (HwRev:B1), DCS-7130-96LBS-F (HwRev:B2), DCS-7130-96LBS-R (HwRev:B1), DCS-7130-96LBS-R (HwRev:B2), DCS-7130-96LS, DCS-7130-96LS-F (HwRev:B1), DCS-7130-96LS-F (HwRev:B2), DCS-7130-96LS-R (HwRev:B1), DCS-7130-96LS-R (HwRev:B2), DCS-7130-96S, DCS-7130-96S-F (HwRev:B1) and DCS-7130-96S-R (HwRev:B1)

- The signal on pps-outX port of device is not specified. (954629)
Series Affected: DCS-7130 Series
- MuxIn Interfaces from non-EB profile, do not get enabled when sourcing directly from different MetaMux instance.
Atleast one another MuxIn Interface in that non-EB profile, must be sourced from either front panel interface or, switch interface from J2C, to make it enabled (982810) (1)
SKUs Affected: DCS-7130LBR-48S6QD, DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R
- An Arista FDK application with Tscore cannot be configured alongside MetaWatch, even if the configurations are disabled.

To use an Arista FDK application with Tscore, all MetaWatch instance

configurations must be removed (not just disabled).

To use MetaWatch, any Arista FDK applications with Tscore must be disabled.
(1089489) (1)

(1) This item is in two or more sections on this document

7358, 7368, 7388, CCS-750, DCS-7300X, DCS-7500R, DCS-7800 and DCS-7800R3 Series

- When a system shuts down due to a power overload, EOS does not log a 'power overload' reload cause (371160)
- Mixed supervisor types are not supported on dual-supervisor modular systems. (450053)
- With scaled VRF, SVI and Arp entries, typically beyond claimed support, the Arp agent may experience a SIGQUIT during SSO switchover, on the new active supervisor. There is no workaround for this issue and switchover is ultimately successful. (495528)
- IP reachability between the ManagementActive interface and physical management interfaces may be impacted if both sets of interfaces are configured in separate VRFs. (885866)
- If Management0 interface redundancy is enabled with ipv6 neighbor monitoring, the backup interface can become active on supervisor switchover instead of the primary interface (893164)

CCS-710P, CCS-720DP, CCS-720XP, CCS-722XPM and CCS-750 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)

7500R3, DCS-7020, DCS-7130, DCS-7280R, DCS-7280R2, DCS-7280R3, DCS-7280R3A, DCS-7500R, DCS-7500R2, DCS-7500R3, DCS-7800, DCS-7800R3 and DCS-7800R3A Series

- The L3 MTU on interfaces is not enforced on SVIs. (11659)
Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SE, 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)

- On DCS-7500 series, stateful switchover is not hitless. (54305)
- 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)
- 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 +/-3 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)
- ACL counters may increment spuriously during modification (117509)
- 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)
- On the DCS-7280E and DCS-7500E series, configuring QinQ or SelectiveQinQ along with configurable TPID on the same interface is not supported (131757)

- VXLAN encapsulated packets received on MLAG peer-link will not be bridged after decapsulation to any MLAG interfaces. The workaround is to redesign the network such that VXLAN encapsulated packets are not received on the MLAG peer-link. (133841)
- IPv6 Egress ACL deny logging is not supported on subinterfaces. (146487)
- 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)
- Congestion on an egress interface used by a GRE tunnel that is a mirroring destination can cause drops on forwarding ASICs where the mirrored traffic ingresses the switch. (158001)
- On the DCS-7500E and DCS-7280 series, traffic outage is expected during transition from a single member LAG to two member LAG and vice versa. Use "platform sand lag hardware-only" to avoid this outage. (160439)
- When MLAG peer mac routing feature is enabled on both MLAG peers, the IPv6 neighbor may not be resolved when source ipv6 address in solicited neighbor is link local address. (165067)
- 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)
- L2-ECMP over VXLAN is not supported for multi-homed hosts under EVPN. (193475)
- Ingress ACLs are not supported on decap groups. (195568)
- 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

- Native VLAN setting on Tap interfaces does not take effect. (227770)
- While in Tap Aggregation mode, ingress VLAN membership filtering does not work when packets are destined to a tool interface configured with egress packet truncation. (227841)
- While in Tap Aggregation mode, interfaces configured for egress truncation will continue to consume hardware resources even when they are down. (228384)
- LLDP functionality is not supported on interfaces with ingress truncation enabled. (229983)
- Sflow agent CPU utilization is around 10% even when hardware acceleration is enabled and the agent is not processing any flow samples. (251674)
- With MTU increased beyond 9100B, there can be some multicast packets of size greater than 9100B that are dropped as a port gets close to linerate. (261446)
- Changing PMF profile on one linecard may cause a brief disruption of PMF-based features on other linecards which should not be affected by the same change. (270849)
- A QoS policy, with a class map which has either set-tc or set-dscp on the ingress port and which has DSCP rewrite map on the egress port, results in indeterministic behavior with respect to DSCP remarking of the outgoing packet. (273320)
- If a QOS and PDP Policy both having police action are applied on an interface and the traffic hits both the policies then QOS Policer Counter is messed up. The "show policy-map type qos pmap-name" will give the PDP Policer Counter, not the QOS Policer Counter. (275994)
- Rate values for subinterfaces shown by "show interface counter rates" may deviate up to 15% from corresponding rate on parent interface. (278700)

- With TCP MSS Clamping feature configured, TCP SYN packets might get dropped under high CPU usage conditions if the incoming TCP SYN packet rate is high. (279370)
- Packets larger than 10196 bytes will be discarded by the forwarding ASICs. (280459)
- PFC on DSCP trusted ports is not supported. (281214)
- Routing is not supported on the parent interface of a L2 sub-interface. (287016)
- 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)
- Stale telemetry flow-spec statistics may be seen until the next "show flow-spec" command is issued. (345195)
- The sFlow sample output interface of IPv6 packet subject to an egress ACL, is set to unknown output interface. (354746)

- IPv4 packets with options in the IP header will not be subjected to Unicast Reverse Path Forwarding (uRPF) checks, and hence will not be dropped even if the source IP is unknown. (355542)
- unicast RPF is not enforced on IP packets when the destination IP needs ARP resolution (356347)
- Software forwarding is not supported for GRE or IP-in-IP tunnel terminated packets using decap-groups. (358369)
- The egress DSCP value in the Outer IP header for packets that undergo headend replication after routing is derived from the Traffic Class (and not copied from the ingress DSCP value). (366189)
- While in Tap Aggregation mode, if the identity tag configuration of a Tap port is changed, traffic received on this port may be sent to an invalid multicast destination momentarily. This will cause the DchUnreachables counter to increment. (372757)
- TX mirroring shows VXLAN encapsulated packets that are actually dropped by split horizon. (375795)
- If more than maximum supported Isec tunnels are configured, it is random as to which ones are assigned flow entries within on-chip encrypt/decrypt engine. The set of Isec 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)
- Syslog message informing user about enabling a counter feature that was previously disabled due to hardware resource conflict may not be emitted. (389299)
- 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)
- 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)
- When the VXLAN routes are programmed using 2 hierarchical FEC levels in hardware, and the outgoing interface of a VXLAN underlay route changes from one forwarding chip to another, then it can negatively impact VXLAN overlay route convergence. (440233)
- Egress MAC ACLs are matched on the incoming packet's ethernet header and not on the rewritten ethernet headers. (458724)
- Dynamic port-channels coconfigured using LACP are not supported as destinations for monitor sessions (463186)
- SflowAccel does not send flow samples to collectors routed via BGP routes with additional backup paths. (476553)
- 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-7280SE, DCS-7280SR, DCS-7280SR2, DCS-7280TR, DCS-7500R and DCS-7500R2 Series

- When GUE is enabled with outer IP hashing, inner IP fields are not included in the load balance key of MplsOverGre transit packets in tc-counters TCAM profile. (491706)
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- When "ip hardware fib optimize" is enabled for a non-nibble aligned prefix-length, local-remote MPLS Pseudowires must have Control Word enabled for MPLS decap forwarding to work correctly. If Control Word is not enabled, then ETHoMPLSoETH traffic that ingress with the first nibble of the inner DMAC starting with "0x4" will be speculative parsed as IPv4oMPLSoETH and incorrectly dropped. (496624)
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- For Packets with GRE tunnel encap header, MTU enforced on the hardware is 3 bytes more than configured value in Cli. (498470)
Series Affected: DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- All asymmetric IRB traffic flows forwarded over EVPN VXLAN toward to a single multi-homed destination are not load balanced across the VTEPs that form the L2 ECMP group. However, note that the traffic destined to different hosts are statically load balanced across the VTEPs. (501343)
- VRF label termination into non-default VRF is not supported for multi-label MPLS packets. (504763)
- When BGP neighbors advertise Flowspec rules with police actions where the number of rules exceed what can be programmed in TCAM, policers will be allocated in hardware for all requested rules even though best-effort programming will only program the subset of rules that fit in TCAM. As a

result, hardware policer resources will be unnecessarily consumed for rules which are not installed in hardware. (506688)

- SVI egress counters are not incremented for IPV4 or IPV6 multicast packets. (508471)
- Egress counter features configured via "hardware counter feature ..." do not include packet or byte counts for packets sent from the local CPU. (525380)
- Counter information flows from the forwarding chips to the tables that support TerminAttr in a very bursty manner. When averaged over a several minute window the rates will be accurate, but for shorter intervals calculated rates may oscillate significantly above and below the nominal rate.

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

- Hitless restart is not supported with Macsec proxy (540419)
- `show forwarding destination` may give incorrect results when the traffic would hit an egress IPv4 ACL using bridged IPv4 traffic or hit an egress IPv4 ACL using one of the following qualifiers: TCP established, fragment, and L4 port range. Note that TCP and UDP egress IPv4 ACL rules use the fragment qualifier. (554480)

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

- Hitless restart of Layer 3 forwarding agent is not supported unless "ip load-sharing prefer local" configuration is removed. Note that this configuration is added by default on this system. (555772)
SKUs Affected: DCS-7280QR-C36
- Route churn under hardware FIB exhaustion may cause BFD sessions to flap. (556089)
- On PTP enabled switches, CRC errors may occur in internal time synchronization. This can cause PTP time to be unstable. PTP time stability will recover on its own after a few seconds. (557370)
SKUs Affected: DCS-7280CR2-60-F

- For products with paired QSFP100 (odd numbered) and QSFP+ (even numbered) ports, the even numbered port may experience link flaps when a transceiver with a different media type from the previous installed transceiver is inserted in the odd numbered port. (565969)
SKUs Affected: DCS-7280QR-C72-F, DCS-7280QR-C72-M-F, DCS-7280QR-C72-M-R, DCS-7280QR-C72-R, DCS-7280QRA-C36S-F and DCS-7280QRA-C36S-R

- VRRP IPv6 using VRRP IPv4 MAC prefix may have unexpected packet forwarding behavior when the VRRP IPv6 and IPv4 have different states and one of the states is Master. (567504)
- Routes resolved over LAG sub-interfaces that have shape rate configured may forward traffic to CPU after SandL3Unicast agent restarts. The workaround is to flap the parent LAG by using the shut and no shut command. (570311)
- Different streams being rate-limited by the same group policer might exhibit bias towards some streams and fair policing across all the streams is not guaranteed. (578041)

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

- When all members of a port channel have "forwarding transmit disabled", traffic will not be forwarded to other VLAN members. (587648)
- While in Tap Aggregation mode, traffic steering service policy rules that match on fields from the MPLS header may result in false matches on packets without an MPLS header.

Remove the MPLS key-fields from feature 'tapagg mac' in the active TCAM profile. (592184)

- While in Tap Aggregation mode, traffic steering service policy rules that match on TCP flags may result in false matches when using the tap-aggregation-user-tcp-flags TCAM profile, or another profile based on this one.

Remove L4 key-fields from the 'tapagg mac' feature in the active TCAM profile. (592191)

- While in Tap Aggregation mode, traffic steering service policies with VLAN matching may match untagged traffic erroneously.

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

- VLAN matching in ACLs attached to monitor sessions may result in false matches on untagged traffic. (598186)
- Packets greater in size than the egress interface MTU (fragmented packets) are not counted in "show interfaces counters ip".

Packet size should be limited to MTU to avoid fragmentation and loss of packet and byte counts. (601808)

- For IPsec tunnel mode on Jericho2, the output of show ip security connection detail incorrectly shows the replay window as 16384, the correct replay window is

128 and
fixed to that value. (609180)

- The SandAcl agent may restart repeatedly when more mirroring ACLs are applied than can be configured on any given FAP in the system.

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

- Security counters are not updated correctly when using Vxlansec with SecTag2 encapsulation format. (630807)
- This platform only supports 64-bit versions of EOS. (646592)
SKUs Affected: 7816R3-FM, DCS-7816-CH, DCS-7816-SUP and DCS-7816-SUP1S

- When 'show forwarding destination' is used in an L2 forwarding scenario, the source MAC address of the packet can be learnt in the MAC address table. This can result in MAC moves if the MAC address was already learnt.

A workaround is to clear the MAC address after running the 'show forwarding destination' command. (650968)

- For configured load-intervals of less than 1 minute, counter rates may fluctuate significantly above and below the actual rate.

The rate will become increasingly accurate as the configured load-interval approaches 5 minutes. (684152)

- Packets originating from or forwarding via the CPU through an SVI are not counted by the egress SVI counter feature ("hardware counter feature vlan-interface out"). (688993)
- Counter sampling for ingress sFlow on sub interfaces is not supported. (718833)
- In an EVPN VXLAN active/active multihoming setup, configuring L2 subinterfaces as non-DF ports is not supported and may lead to unexpected SandMcast agent restarts. (739986)
- IPv4 and IPv6 strict mode Unicast Reverse Path Forwarding (uRPF) does not work if the route matching source IP address resolves via an ECMP nexthop. Such packets will be dropped. (757793)
- With scale QoS scale policy-map config with ACL rules, configure-replace might CLI time-out (766811)
- Configure replace with a scaled QoS policy-map config with multiple ACL rules may timeout. (766984)

- For tunnel destinations, the decision to fragment is based on unencapsulated packet size, while MTU enforcement is based on the encapsulated packet size. This may cause unencapsulated packets near the MTU packet size to be dropped instead of fragmented. (768410)
- At high combined scale of VLANs and pseudowires in VPLS or L2 EVPN, exhaustion of "managedTt inLif" resources may occur. This may prevent hardware programming of any configured patch panel patches, resulting in "Unprogrammed" status for one or more patches. It may also result in VPLS packets drops in the decapsulation direction. This situation can be detected by "show hardware capability".

The workaround is to reduce the scale of these or other features which use the managedTt inLif resource and then perform a shut/no shut on the affected patch panel or subinterface. (781993)

- "ip hardware fib next-hop sharing" is not supported with VXLAN. (795334)
- Traffic-policy 'redirect next-hop' and 'redirect next-hop group' actions on L2 interfaces are not supported. (821920)
- If multiple traffic policies are applied such that the TCAM usage for traffic-policy is close to scale (~100%), then the traffic-policy actions may not work if the Sand agent is restarted.

The workaround is to remove these policies from applied interfaces, and re-apply each of them. (821984)

- IP nexthop group entries are categorized as MPLS tunnels in the 'show hardware capacity' and 'show platform sand l3 summary' CLI commands. (823863)
- The VLAN interface counter feature is not compatible with dynamic VLANs used in EVPN VXLAN prefix routing setups. When the VLAN interface counter feature is enabled, a "Counters unavailable" message will be shown for each of EVPN VXLAN dynamic VLAN in the "show interface counters" CLI command output. (840994)
- The virtual-cable feature is unsupported, but the configuration CLI is not blocked. Configuring this unsupported feature will result in the Sand agent repeatedly restarting.

The workaround is to remove the virtual-cable feature from the running config. This will allow the Sand agent to recover. (893360)

- If 'switchport vlan tag validation' is configured, packets with more than 2 VLAN tags will be dropped. (898006)

- The command "show interface status errdisable" may throw an error when Tap Aggregation is enabled and the switch is linked to another switch using Multi-Chassis Link Aggregation (MLAG). As a workaround, disable MLAG on the switch before running the command. (913208)
- Interface traffic-policies cannot match on VXLAN decapsulated bridge packets. (913467) (1)
- If a Tap Aggregation policy map is applied to an interface while running in the tap-aggregation-default profile, or any user profile that does not have the count action in the Tap Aggregation features, then the message "ACE-6-ACTION_NOT_AVAILABLE" will be logged for StatId and StatProfile. These messages mean that the policy ACL counters will not work for these policies. Other configured actions are not affected. (939008)
- Arista Networks Release Notes for EOS Release 4.31.1, 4.31.2 and 4.32.0 incorrectly documented the 7816B-CH in the New Platforms and Hardware and Supported Hardware sections. 7816B-CH is not introduced in those releases. 7816L-CH is introduced/supported in 4.31.1, 4.31.2, 4.32.0 and newer releases. Newer versions of EOS release notes after April 20, 2024 correctly document the support for 7816L-CH. (949378)
- Replacing one PBR rule with another rule (by changing the match criteria and/or actions) at the same logical priority in the PBR policy is considered both an addition of a rule, and a removal of the rule previously configured. Any policy update containing this replacement operation cannot be programmed using in-place update of TCAM. (955171)
- Configuring L2 EVPN MPLS on a subinterface with a number greater than 4095 causes the SandTunnel agent to restart repeatedly. The workaround is to use subinterface number less than 4095. (981028)
- A TCAM bank that is allocated for CBF does not get released until all CBF rules are removed or the feature is disabled. A workaround is to disable and then re-enable CBF by removing all color to DSCP mappings and then adding them back. (990098)
- Source VTEP pruning with VTEP-to-VTEP bridging is not supported if the underlay L3 interface is an SVI and the nexthop MAC is unknown. Workaround is to configure a static MAC address for the nexthop. (993682)
- Counter for packets matching ACLs configured in monitor sessions cannot be reset. A workaround is to remove the monitor session with these ACLs and reconfigure them. (1015985)
- The EAPI output of `show interfaces counters` for subinterfaces categorizes packets as "inUcastPkts", but it actually counts all packets (unicast, broadcast, multicast). (1046158)

- User-Defined Fields, or UDFs, are defined as part of an access-list filter and are comprised of an offset, length and pattern match. The behaviour of a UDF where the offset value exceeds the packet's size is undefined and may result in false matches. (1058051)
- On modular switches, enabling disabled ports on linecards not in Tap Aggregation mode using the "no-errdisable" command is permitted erroneously. The correct operation of such ports can not be guaranteed. (1081839)

DCS-7280R and DCS-7500R Series

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

Series Affected: DCS-7020, DCS-7280CR2, DCS-7280R, DCS-7280R2, DCS-7280SE, 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)
- SandTcam agent may unexpectedly restart while changing the system TCAM profile to a profile that has feature 'tcp-mss-ceiling ip' on switches that has some linecards using Algomatch as the access-list mechanism and some other linecards using TCAM as access-list mechanism.

Workaround : Configure 'hardware access-list mechanism tcam' (474595)

- 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

- On TCAM profile change, sub-interfaces that have a L2 Protocol Forwarding profile applied or those sub-interfaces for which a sub-interface on the same physical interface has a L2 Protocol Forwarding profile applied, will briefly forward all traffic that is otherwise supposed to be trapped to CPU by default. (525564)
- If the CPU is included in monitor sessions with multiple destinations, it will only work if the device mirror0 is selected. If multiple mirroring to CPU sessions with multiple destinations are configured, all packets will be sent to mirror0. (525919)
- With RFC5549 ACLs configurations, the Egress IPv6 ACL deny logging on port channel subinterfaces stops working on performing agent SandL3Unicast restart. (549181)
- MPLS explicit NULL termination isn't supported with flex-route configuration containing non-nibble aligned prefix length. (576801)
- If MPLS tunnels are allocated in uncountable egress banks due to transiently high scale, they can be left in an uncountable bank which will prevent traffic counters from being aggregated for these tunnels. As a workaround, a hitfull restart of the forwarding plane agent will reallocate the tunnels in countable egress banks 1 & 2. (635152)

DCS-7280R2 and DCS-7500R2 Series

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

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

- During normal operation, the internal time synchronization in a switch may log the message 'TIMESYNC_ERROR ... (slave signal CRC error)'. The internal time synchronization will correct this error without disrupting the normal operation of the switch. (818143)

SKUs Affected: DCS-7280CR2-60

- With "no platform sand ipv6 host-route exact-match" config, outer IPv6 IP-IP/GRE/GUE encapsulated transit packets will not have PBR and other similar policy-map features applied.

One workaround is to re-enable the "platform sand ipv6 host-route exact-match" config. (885787)

7500R3, 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)
- 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)
- IP forwarded traffic that is mirrored to a GRE tunnel destination will have the PCP bits of their VLAN tag set to 0. (400909)
- Application of low rate shaping on multiple interfaces with line-rate traffic into the box will create a backlog in DRAM. Due to this majority of traffic will be served to egress from slower DRAM instead of SRAM. In the congestion state, even after the removal of shaping with line-rate traffic ingress from multiple ports, switch won't be able to deliver full throughput. The recovery from DRAM to SRAM will happen over a period of time once the congestion is removed. (414481)
- Set dcsp will not apply to mpls routed packets that egress with mpls labels (418791)
- Enabling Tx mirroring may cause the counter VOQ_SRAM_ENQ_RJCT_PKTS to increment, with reason QnumNotValid, as shown in the output from "show

platform fap counters pipeline". There is no impact to mirrored or forwarded traffic. (459862)

- Unidirectional links are not supported at 40G speed on QSFP100 interfaces using the CRT50216 PHY. (467672)
SKUs Affected: 7368-16C, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3-48CQM-LC, DCS-7280CR3-32D4, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-96, DCS-7280CR3-96-F, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-96 and DCS-7280CR3K-96-F
- SFlow samples of packets which are dropped by an ACL will have incorrect output interface indication. (470650)
- 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)
SKUs Affected: 7500R3-24D-LC, 7500R3-24P-LC, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-36D-LC, 7800R3-36P-LC, 7800R3-48CQ-LC, 7800R3-48CQM-LC, 7800R3K-36DM-LC, 7800R3K-48CQ-LC, 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-32D4A-F, DCS-7280CR3K-32D4A-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-32P4A-F, DCS-7280CR3K-32P4A-R, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F and DCS-7280PR3K-24-F
- The sFlow flow samples generated for VXLAN decapsulated packets, have the first 6 header bytes of the sampled packet set to zeros. (502957)
- 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
- Encapsulating PTP packets over VXLAN tunnel is not supported. (526461)

- When Hardware Accelerated Sflow is configured the uptime in sFlow datagram is updated every 1/4 of milli sec, rather than every milli sec as prescribed in sFlow standard. (531769)
- Sequence number in sFlow datagram doesn't reset to one when hardware accelerated sFlow is disabled.
The sequence number in sFlow datagram gets reset to 1 for those sFlow sub-agent when the line-card is reset/ restarted on which the sFlow sub-agent is residing. (531983)
- With loose mode uRPF configuration, the switch doesn't drop packets with an uRPF lookup (using source IP) matching a drop route. (531998)
- When SandCounter agent is down, datagram counters are not accounted. Hence in show sflow output, the "Number of Datagrams" entry displays lesser value than expected. (536213)
- If all front-panel ports are populated, and traffic load is high, an improperly programmed current limit could be exceeded generating a spurious power-off. (536331)
SKUs Affected: DCS-7280CR3-96 and DCS-7280CR3K-96
- Symmetric hashing is always on for IPv6 addresses, when both source and destination addresses are included in the applied load-balance profile. (557714)
- In transparent two-step mode, PTP packets egressing the same port with the same sequenceId may cause transient jumps in correctionField.

The workaround is to use one-step transparent clock mode (564682)

- Switching over in SSO mode fails when initiated on a system with no linecards inserted. (604537)
- In Tap Aggregation mode, when both header removal and truncation are configured on a tap port, a packet will not egress the switch if it has a size less than 64 or has an invalid ethertype, after header stripping and truncation actions are performed. (733144)
- MAC access-lists are not applied on L2 protocol packets even when L2 Protocol Forwarding profile is applied. (735328)
- When using the Uniform-4x100G-2-R L1 module profile, the "/3" and "/4" interfaces will not support single lane speed modes (e.g. 50G-1). (786990)
SKUs Affected: 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3A-36P-LC, 7800R3A-36PM-LC, 7800R3AK-36DM-LC, 7800R3AK-36DM2-LC, 7800R3AK-36PM-LC, DCS-7280DR3A-54-F, DCS-7280DR3AK-54-F and DCS-7280DR3AM-54-F

- Ethernet frames ingressing on a routed interface with multicast MAC as destination MAC address may be unnecessarily copied to CPU. (910627)

7500R3, DCS-7130, DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series

- When operating under fabric egress replication mode (``platform sand multicast replication default fabric-egress``), multicast traffic egressing (or dropped in the egress pipeline of) Ethernet ports is incorrectly counted as unicast traffic in the "Egress Queue Counters" section of the ``show interfaces counters queue`` command. In other words, all traffic will be counted as unicast. (488383)

DCS-7280R3A and DCS-7800R3A Series

- Packets mirrored from a TX source to a GRE tunnel destination will have incorrect metadata in the "key" field of the GRE header. (763403)

DCS-7020 Series

- The interface bin counter for both in and out packets with sizes 1523-Max may be inaccurate. Packets with sizes between 9217-9236 are not counted by the switch for interfaces Ethernet 1-48. (188121)
- Mirroring and sFlow can not be supported on the same source interface at the same time due to a hardware limitation. (209060)
- AES128/SHA-1 and AES256/SHA-256 are the only supported encryption/authentication algorithm pairs. (342005)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Using an Ipsec tunnel interface for policy-based routing nexthop is not supported. (378861)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Ipsec tunnel interfaces are not supported as part of nexthop-group. (378890)
SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- Dynamic NAT, AKA Port Address Translation, is not supported for IPSec tunnels. (384624)
- Phy loopback configuration is shared among groups 4 ports. E.g. when one of the ports in the range of et1-4, or et5-8 etc., is put into loopback mode, all ports in this group will be in loopback mode. Traffic is affected on all these ports. (990020)
SKUs Affected: DCS-7020TR-48-F, DCS-7020TR-48-R, DCS-7020TRA-48-F and DCS-7020TRA-48-R

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

- L4 ops and fragments rule cannot be programmed together. (94197)
- Egress IPv4/IPv6 RACL does not work for VXLAN encap-route case. (154551)
- 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)

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

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

- Vxlan decapped packets will not match rules that match on 'tcp established' in egress acls (226582)
- VXLAN and GRE tunnels cannot be configured at the same time. (248239)
- If a PTP packet is mirrored to a GRE destination which has time stamping enabled, the timestamps on such packets are not reliable. (280362)
- If a packet is mirrored to a GRE tunnel, but the forward copy is trapped to the CPU, then the ingress VLAN ID carried in the GRE key will be set to 0 when the mirrored copy is encapsulated in the GRE envelop. (434871)
- In the event that two rate three color (trTCM) policers are configured for an incoming broadcast or multicast traffic, drop precedence (dp) classification may not happen correctly for those packet that are subject to special Copp class policers. A portion of the ingress traffic is marked with as yellow. (481688)

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

- Egress MAC ACLs do not work on interfaces that are not front-panel interfaces, even though they can be configured. (483583)
- VxLAN encapsulation will not work if mirroring to GRE destinations are configured. (548429)
- MPLS label popped packets will not match rules that match on 'tcp established' in egress IP acls (565015)

- PTP v1 packet forwarding is incompatible with both Tap Aggregation BR/VN tags stripping and GUE packet decapsulation.

The workaround is to disable these features before enabling `ptp forward-v1` configuration. (572447)

- BGP Flowspec rules will not work on SVIs if the TCAM profile does not include a "port qualifier size N", N > 1 configuration for the flow-spec features. (588565)
- L4 destination port matching in egress IP ACLs does not work on untagged bridged traffic. (609461)
- Hitless traffic-policy update is not supported on Sand Gen3 systems. (622066)
- When the non-head-fragment field qualifier is not present for the traffic-policy feature on Sand Gen3 chips, L4 qualifiers will not match on any fragment packets, including head fragments.

As a workaround, add an explicit rule for handling fragmented packets. (628773)

- L2 EVPN MPLS won't work in presence of TCAM feature "VLAN Editing Extended-64" in its TCAM profile. (648996)
- In Tap Aggregation mode, egress IP ACLs are not supported on tool interfaces that have egress truncation configured. (668716)
- The following two configurations are incompatible:
 - IPv6 route optimization with "ipv6 hardware fib optimise prefix-length 128 []" for /128 prefix-length.
 - Using "no platform sand ipv6 host-route exact-match" to disable IPv6 host routes in exact match.

The workaround is to remove either of the two commands from the configuration. (671966)

- Enabling IPoGUE outer UDP hashing may break IP-in-IP hashing. (676230)
- On adding a member to a port-channel sub-interface that has a named QoS-map configured, the ingress packets get classified based on the parent interface configuration for a short duration of time. (680689)
- Egress Port ACL does not work for traffic that is bridged after VXLAN decapsulation. (688880)
- TCAM feature 'acl port ipv6 egress' does not add support for egress IPv6 PACLs when added to a profile. (699451)

- If the active load-balance profile contains "packet-type mpls-over-gre outer-ip" configuration and the active TCAM profile contains "key field propagated-ttl" configuration under "feature mpls", then incoming MPLS-over-GRE traffic with MPLS TTL equal to zero or one may be incorrectly forwarded instead of dropped when terminating the GRE tunnel and forwarding based on the MPLS payload.

Workaround: Remove "packet-type mpls-over-gre outer-ip" from the load-balance profile or remove "key field propagated-ttl" from the TCAM profile. (701931)

- In an EVPN IRB setup, enabling both VXLAN and MPLS is not supported. Traffic loss is expected in this configuration. (703701)
- BFD/BGP/OSPF/IS-IS protocol packets received over VXLAN take generic CoPP queues to CPU instead of the protocol specific ones. (719542)
- L2 EVPN MPLS configuration will not work when the same trunk port is used to carry traffic towards both the Provider Edge (PE) and the Customer Edge (CE). (727163)
- When IP PIM sparse mode is enabled on the VLAN interface, egress port ACLs applied to switch ports that are a member of that VLAN match on random fields for bridged multicast traffic. (728267)
- Egress IPv4 ACLs do not apply to unknown multicast traffic (728269)
- QoS policy-maps are not supported on dot1ad, unmatched or untagged subinterfaces. (732743)
- QoS policy-maps are not supported on dot1ad, unmatched or untagged subinterfaces. (733077)
- IS-IS over VXLAN is not supported when the CLI command "vxlan decap floodset unknown-multicast unknown-unicast exclude cpu" is configured. (749279)
- The "Uniform-4x10G-1" L1 module profile is incompatible with the "boot port numbering octal interfaces 4" CLI configuration.

Users should remove the "boot port numbering octal interfaces 4" command from their configuration and apply the Uniform-4x100G-2-R" L1 module profile on card's where 4 interface octal port numbering schemes are desired. (779375)

SKUs Affected: 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3A-36P-LC and 7800R3A-36PM-LC

- Changing the profile for hardware accelerated sflow via "sflow hardware acceleration profile" will not take effect until the switch is restarted. (851412)
- The IPv6 decap-groups feature is unsupported when IPv6 host routes are programmed in exact match tables. There is a CLI warning when configuring IPv6 decap-groups if this feature conflict happens, but a system log message is not generated.

The workaround is to disable the IPv6 host routes in exact match table feature, which is enabled by default, with "no platform sand ipv6 host-route exact-match". (860223)

- Enabling a GRE destination on a monitor session while traffic is running on the monitor source may cause the DCH unreachable discard counter to increment by a small amount, due to some superfluous traffic being mirrored into the fabric without a valid destination. (882484)
- QoS features configured using tx-queue mode of interface are not supported on ports having 2 priority level.
"show platform fap mapping interface <interfaceName>" shows number of priorities supported by the port under QPairs column

List of unsupported QoS features on the tx-queue of the 2 level priority port include:

- 1) Shaping/Scheduling
- 2) WRED
- 3) ECN
- 4) PFC
- 5) Latency based drops. (893100)

- Private VLAN feature is not supported on switches configured with a tcam profile having "vlan editing extended-64" tcam feature. In such cases, private VLAN packet forwarding behaviour is undefined.
A workaround is to apply a tcam profile not having "vlan editing extended-64" tcam feature. (957675)

7500R3, DCS-7130, DCS-7280R3, DCS-7280R3A, DCS-7500R3, DCS-7800R3 and DCS-7800R3A Series

- Priority tagged packets (802.1Q tagging with the VLAN set to 0) being reflected by an Ethernet reflector interface configured in the egress direction with MAC address swap action will have the priority tag stripped from the packet before they are transmitted. (354190)
- Only port-channel interfaces are supported as peer links. Single member port-channel interfaces are supported. (409958)

- If a packet that is decapsulated with a decap group is mirrored to a GRE tunnel, the packet's ingress VLAN will be set to 0 in the upper 2 bytes of the GRE key. (430986)
- Egress ACLs does not match packets with UDP ports 4754 and 4755. (455922)
- Traffic matching PBR rules will be dropped if it violates the MTU of the PBR destination port (471548)
- A learnt MAC address won't age out in case there is only continuously running unicast RPF drop traffic from the source MAC. The workaround is to clear the dynamically learnt MAC from the MAC address table. Once the MAC address is cleared, the MAC address won't be learnt if all traffic with the source MAC address is dropped by unicast RPF. (536215)
- If there are multiple mirror to GRE monitor sessions sharing the same GRE tunnel parameters (i.e. source and destination IP address, DSCP values, TTL, protocol, all matching), and if any of these sessions has rate limiting enabled, then all these monitor sessions may either have rate limiting functionality disabled, or share the rate limit configured in one of these sessions, following a switch reload.

A workaround is to remove these monitor sessions and reconfigure them again. (548973)

- During EVPN MPLS encapsulation, the VLAN priority bits of payload Dot1Q header will be marked using the traffic-class to CoS map configuration in the global QoS map. (550987)
- QoS policy-map is not supported on Pseudowire decapsulated traffic flows. (574600)
- On switches with B52 external PHY devices, the interface command "traffic-loopback source system device phy" will result in packets being both transmitted and looped back. (582409)
- When there is congestion in Tc6/Tc7 across multiple queues. Tc6/Tc7 Queues can hog up entire DRAM. This can lead to DRAM rejects of lower TC traffic of unrelated stream thus causing some discards. In the current case we had 2 to 3% lower instead of 100% throughput of data traffic (591523)
- Renamed TCAM Qualifiers l4ops_ip_pacl, l4ops_ip_qos, l4ops_ip6_pacl, and l4ops_ip_racl into l4ops_24b, l4ops_7b, l4ops_3b and l4ops_18b. This can cause customer scripts using `show platform fap pmf` command to fail if they are using any of the above renamed TCAM qualifiers. (595868)
- On Gen4 Faps , SSO for TCAM based ACL rules is hit full. There is a brief window where ACLs are not applied to the flows. (606756)
- On Gen4 Faps , SSO is hit full for TCAM based ACL rules. (606764)

- When outer IP hashing is enabled, GUE decap followed by MPLS NULL label termination is not supported. (607977)
- Configuration of "encapsulation dot1q vlan" is only supported on sub-interfaces. (608383)
- Software forwarding of MPLS-over-UDP terminated packets is not supported. (613102)
- Added support for L4 port match on VXLAN IPv6 underlay transit packets (forwarded on either outer L2 or L3 header). (614691)
- If a packet matches both a Tap Aggregation steering rule, as well as either a security ACL or traffic-policy rule to drop the packet, the steering policy counter will still increment. (620429)
- SandTm restart is hitful for named TC to DSCP QoS maps. (632481)
- SandOam may restart with 200+ DM sessions running at 10 milliseconds. The workaround is to use higher tx-interval like 100 milliseconds, 1 second, 10 seconds, 1 minute, 10 minutes. (642106)
- When the transmit interval of CCMs is increased, some of the remote MEPS are shown as unreachable even though CCMs are being received from the remote MEPS. As a side-effect, this also causes the RDI bit in the generated CCMs to be set.
There are a couple of workarounds:
 1. Wait for 5-10 seconds before triggering a CCM interval change on the peer device. This can potentially avoid this issue completely.
 2. If the issue is still seen, restart the SandOam agent. (643283)
- SandL3Unicast on the peer DUT causes a few RMEPS to be marked as unreachable on the local DUT. A SandOam restart on the local DUT recovers from the problematic state. (646125)
- IP UDFs in ACLs may not work correctly on IPv6 packets with a routing extension header when matching on bits beyond the IPv6 header. (650286)
- VXLAN is not supported with greater than 256-way ECMP in the underlay. (652266)
- In case of Delay Measurement, the device does not work as responder to DMM transmitted by any other device which configures the TLVs in the DMM. The device is incapable of copying the received DMM's TLVs to the DMR and if the other device expects the configured TLVs in the DMR then the device can reject the such DMRs. (654226)
- GUE tunnel decapsulation does not support UDP destination port 2152. (675837)

- In RFC2544 throughput test, oversubscribing the ingress interface causes hardware test engine lockup and test traffic won't be generated. Restart the SandFapNi agent. (684101)
- A pseudowire with Cos-To-Tc or Dscp-To-Tc QoS map applied will classify all of the decap traffic to default traffic-class of 1 in the presence of entropy or the null label. (694370)
- While RFC2544 test is in progress, received and sent packet counts may not match. (694683)
- If egress MAC/IPv4 ACL is applied to the front panel port and egress IPv4/IPv6 ACL is applied to SVI with ACL sharing enabled for the SVIs, only the ACL applied to the SVI will match the routed traffic. (695784)
- When using IPv4 route optimization using "ip hardware fib optimize .." command, the number of VRFs supported on the system is reduced to 30 non-default VRFs. (720042)
 SKUs Affected: 7800R3-36D-LC, 7800R3-36P-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3-48CQM-LC, 7800R3-48CQM2-LC, 7800R3-48CQMS-LC, 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3AK-36DM-LC, 7800R3AK-36DM2-LC, 7800R3K-36DM-LC, 7800R3K-48CQ-LC, 7800R3K-72Y-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-36S, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4A, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4A, DCS-7280CR3K-96, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4S, DCS-7280DR3-24, DCS-7280DR3K-24, DCS-7280PR3-24, DCS-7280PR3K-24, DCS-7280SR3-48YC8, DCS-7280SR3K-48YC8, DCS-7280SR3K-48YC8A and DCS-7280SR3M-48YC8
- The MPLS no-php mode is not supported with L3VPN and VPWS services. (720213)
- The packet counting logic for CFM Loss Measurement adheres to Version 1 of the protocol although the LMM is sent with version 0. This can result in an incorrect loss being reported. (720859)
- OAM doesn't generate SLR's in response to SLM's if Test ID in the SLM is more than 1048575 (0xFFFFF), instead it generates LMR's. (726371)
- The TLVs present in the Loss Measurement Message (LMM) are not copied to the Loss Measurement Reply (LMR) that is generated in response. This can cause the LMM initiator device to invalidate the LMR. (727007)
- The "Flag" field in the Loss Measurement Reply that is generated in response to an ITU-T.Y1731 version 1 Loss Measurement Message is set to 0, indicating an on-demand session. (727056)
- CFM Loss Measurement will report incorrect loss if the loss between consecutive LMM/LMRs is greater than 32K packets. (729327)

- When using IPv4 route optimization using "ip hardware fib optimize .." command, the number of VRFs supported on the system is reduced to 14 non-default VRFs. (731721)
SKUs Affected: DCS-7280SR3-40YC6 and DCS-7280SR3E-40YC6
- Traffic throughput on an internal recirculation interface is capped at 100Gbps. (738474)
- Three level hierarchical IP nexthop groups are not supported. (751053)
- When policy configuration changes are made with segment-security, some counts may be missed or misclassified under "show segment-security counters" for up to 30 seconds. (769949)
- Post SandOam restart RFC2544 generator test is initiated for subinterfaces that are operationally down. (772320)
- VXLAN bridging of packets that are received on a 802.1Q tunnel port with more than one 802.1Q tag does not work (782097)
- Traffic sent towards control plane for unresolved ARP or next-hop which is marked with software drop-precedence 1 gets dropped. (792663)
- CFM PDUs received on the Up-MEP interfaces that are to trapped/dropped get counted in the following platform counters: "Tx Total Err On Sop Pkt Counter", "Tx Total Pkt With Err Counter", "Total Byte With Err Counter", "Tx Total Err On Sop Byte Counter", "TxTotalPktWithErrOnEopCounter" and "TxBytesWithErrOnEopCounter". (796338)
- When ETH/MPLS/IP/GTP packets are MPLS terminated, hashing on GTP TEID will not occur.

A workaround to hash on GTP TEID is to use the custom UDP payload hashing feature. (813466)

- Hardware resources are exhausted when hardware accelerated sFlow is enabled with scaled subinterfaces, predominantly with port-channel subinterfaces. Possible workarounds are disabling sFlow hardware acceleration or disabling subinterface ifindex stamping by running "no sflow sample input subinterface". (813602)

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

- Powering off a Linecard which is transmitting high bandwidth on traffic class 6 and 7 may lead to flap of protocol session on other Linecards. Workaround is to shutting down all interfaces of the Linecard before its removal. (832211)
- After deleting the operational TCAM profile and then reloading the switch, no TCAM banks would be granted.

The workaround is to re-configure the TCAM profile. (836367)

- When the device is configured with 'vxlan qos dscp ecn propagation decapsulation disabled', the inner DSCP for VXLAN decapsulated packets is retained even if the ingress port is in cos or untrusted mode and global DSCP rewrite is enabled. (836505)
- Mirror copies, or sFlow samples, of decapsulated VXLAN traffic will have bytes missing from the underlay headers. (843636)
- In Tap Aggregation mode, if an incoming packet has 3 or more VLAN tags and 2 tags are selected to be removed via the Tap Aggregation DOT1Q tag removal feature, the packet may be malformed at egress with an additional VLAN tag.

Using the DOT1Q remove feature to remove a single tag is supported. Removal of additional tags can be implemented by recycling the packet via a tap-tool and traffic loopback feature. (850175)

- The switch does not learn the MAC address of routed packets ingressing on VLAN interfaces when using IPv4 route optimization involving two prefix lengths.

The workaround is to set the ARP aging timer to be smaller than the MAC address aging timer, so that ARP packets refresh MAC addresses. (868827)

- The following two configurations do not work together:
 1. IPv4 optimize configuration of internet profile or any other compression configuration - "ip hardware fib optimize prefixes profile internet" or "ip hardware fib optimize prefix-length 21 compress <>".
 2. PBR policy matching on nexthop-group.
policy-map type pbr PMAP1
10 match ip any any nexthop-group <>

The work around is to do either of the below:

1. Remove the optimize configuration
 2. Use an optimize configuration without compression
 3. Use urpf-internet profile. (893189)
- If routes are learnt via BGP, RFC5549 traffic egressing from L3 sub interfaces and SVIs is not subject to egress IPv4/IPv6 ACLs.

Egress IPv4/IPv6 ACLS on RF5549 traffic will be functional after configuring "forwarding-type" key field in the "acl vlan ipv6 egress" TCAM feature in the operational TCAM profile. Some unused fields from "acl vlan ipv6 egress" feature of the operational TCAM profile needs to be removed to make space for this "forwarding-type" key. (903500)

- CFM protocols do not work if egress sFlow is active on the interface on which CFM MEP with direction Up is defined.
Disable sFlow on CFM UP MEP interface. (905833)
- RFC2544 Initiator benchmarking tests with 64 byte packet size is only able to test max throughput of 256Gbps because of hardware limitation.
400Gbps RFC2544 Initiator throughput testing is supported for packet sizes greater than 100bytes. (915701)
- A traffic-policy match which includes 'encapsulation gtpv1 ipv4', but no inner match criteria will match on both IPv4 and IPv6 traffic. This same behavior is also seen for 'encapsulation gtpv1 ipv6'.
Workaround: add qualifiers inner-src-ip-label, and inner-dst-ip-label, to feature "traffic-policy port ipv4/6" in the TCAM profile (917902)
- When hardware accelerated sFlow is enabled, the UDP checksum field of sFlow datagrams is set to 0. (949490)
- SandTunnel agent restart or SSO may lead to traffic loss for traffic matching VRF selection policies. (954545)
- When using the mirror-on-drop packet buffer drop reporting feature, if for a given forwarding chip and core there are several high frequency packet buffer drops with different drop reasons occurring, or if there are high frequency packet buffer drops occurring across multiple interfaces connected to that forwarding chip and core, then lower frequency packet buffer drops may not be mirrored. (957516)
- Untagged traffic entering VPWS port-based connectors are always classified as traffic-class 1, and ignores any QoS mapping configurations. (967855)
- Packets entering a VLAN via an L2 subinterface will not be matched by segment security. (974329)
- For hardware accelerated sFlow, flow samples ending in the pattern 0xFFFFFFFF may be truncated by 4B. (989032)
- The scale of VTEP decap counters is limited to 1365 VTEPs. (1015265)
- On multi-chip systems, some MAC entries may prematurely age out resulting in unexpected flooding until these entries are relearned. (1030082)
- Due to hardware limitations, it is not possible to mirror the entire VXLAN packet when VXLAN is enabled on the switch. (1044064)
- RFC2544 test fails due to VOQ latency based packet drops for packet generated from SAT.
Workaround : disable VOQ latency based drop using CLI knob "platform sand VOQ tail-drop latency disabled" (1048704)

- At 200G or 400G speeds, the FEC degraded SER indication may be transmitted to the peer. The link otherwise operates normally. (1101261)
 SKUs Affected: 7500R3-24D-LC, 7500R3-24P-LC, 7500R3-48Y4D-LC, 7500R3K-48Y4D-LC, 7800R3-36D-LC, 7800R3-36P-LC, 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-7280CR3-36S-F, DCS-7280CR3-36S-R, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32D4A-F, DCS-7280CR3K-32D4A-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-32P4A-F, DCS-7280CR3K-32P4A-R, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F and DCS-7280PR3K-24-F
- When L4 source and/or destination ports are matched as part of an ingress ACL, VXLAN bridged packets that are decapsulated on this device, may not match the expected ACE and may end up matching an unexpected ACE. (1111896)
- With outer IP hasing configuration, non-/32 prefixes are not supported in IP decap-groups. MPLS-over-GUEv4 packets are misforwarded with such a configuration..

Workaround is to expand the non-/32 decap IP prefixes and configure them as individual /32 IPs within the decap group. (1133677)

- When using Hardware Accelerated Sflow, the uptime in sFlow datagrams start with a random value. (1150365)

DCS-7280R and DCS-7280R2 Series

- CBF does not work in AlgoMatch mode on some custom TCAM profiles (603206)

7300X3, 7320X, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, 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)
- 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, Link Local Multicast traffic will not be intercepted by the ACL attached to the VLAN interface. (105504)
- There may be a momentary traffic disruption every time the nexthop specified in a PBR policy resolves to a new destination: affected packets may be forwarded by normal L3 lookup, or be dropped.

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

- When any PBR policy is applied, packets that violate the egress MTU and are destined to flood the VLAN will end up being flooded instead of being sent to the CPU. (105680)
- Subinterfaces cannot be used to send or receive VXLAN encapsulated packets. (105767)
- When "hardware access-list resource sharing vlan in" is configured, changing the ACL on a VLAN interface:

1. From an ACL that is not shared with other VLAN interfaces to one that is shared with other VLAN interfaces can cause deny traffic to get permitted until the new ACL takes effect.

2. To another ACL that doesn't fit in the TCAM can cause deny traffic to get permitted until the system reverts back to the previously applied ACL. (106646)

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

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

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

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

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

- 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)
- 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-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)
- 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 and DCS-7300X Series

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

Series Affected: 7300X3, CCS-720XP, CCS-750, DCS-7050CX3, DCS-7050SX3, DCS-7050TX3, DCS-7050X3, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7260CX3 Series

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

Series Affected: DCS-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 and DCS-7060X Series

- A configured mirroring ACL may not be applied in hardware. This can be confirmed by checking the output of 'show platform trident tcam mirror'. (236908)
- When a VXLAN encapsulated packet is received on a VXLAN VLAN where PBR policy is applied, the switch will not PBR forward the decapsulated packet. (244093)
- Sflow output interface determination will not work when the output interface is a subinterface. (250424)

- Cut-through mode is not supported on SFP+ interfaces. (252731)

Series Affected: DCS-7050X3, DCS-7050X4, DCS-7060X4, DCS-7060X5 and DCS-7260X3 Series

- A switch may not be able to bridge or route VXLAN traffic if it hits one or all of following conditions:
 - Nexthop resources are exhausted as indicated by syslog `VXLAN_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)
- 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)
Series Affected: 7358, 7368, 7388, DCS-7050X4, DCS-7060PX4, DCS-7060X4 and DCS-7060X5 Series
- Ports with ingress MAC ACL still learn MAC addresses although traffic may be dropped (400211)
- Packets originated from the CPU are not included in egress SVI/subinterface counters. The CPU originated packet/octet count can be obtained from the switch bash shell using the linux ifconfig command. (400885)
- 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)
- A 10G or 25G SFP transceiver plugged into a QSFP port using a QSFP-SFP Adapter (QSA) reserves 4 logical ports. (415439)
SKUs Affected: DCS-7050SX3-48C8, DCS-7050SX3-48YC8, DCS-7050TX3-48C8, DCS-7260CX3-64 and DCS-7260CX3-64E
- 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

- When a VXLAN VTEP is reached through a resilient ECMP route, (VXLAN bridging or routing) traffic to that VTEP may get dropped in HW. Workaround is to remove the Eesilient ECMP config. (421216)
- 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-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X4, 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
- BUM traffic is not sampled by egress sFlow. (459902)
- 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. (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 store-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)

- Mirror on drop does not mirror dynamic NAT traffic drops. (506047)
- Egress mirroring of multicast packets that are going to be encapsulated on the switch, independent of mirroring, is not supported. (512880)
- Mirror on drop does not mirror VXLAN IPv6 underlay traffic drops. (516212)
- Issuing "reboot -f" from the shell can cause continuous machine check errors, requiring a power cycle.

Instead, use the "reload" CLI command, or "reboot" in bash. (525314)

SKUs Affected: DCS-7010TX-48, DCS-7010TX-48-DC, DCS-7010TX-48-DC-F, DCS-7010TX-48-DC-R, DCS-7010TX-48-F and DCS-7010TX-48-R

- The mirroring logic can only capture and mirror packets seen by the ingress pipeline. Packets generated by the memory-management unit do not traverse the ingress pipeline and hence cannot be mirrored. This applies to all types of mirroring. (527668)
- When policer is attached to the class-default of an egress policy-map consisting of both IPV4 and IPV6 ACLs, IPV4 and IPV6 traffic hitting class-default would be policed individually with rate configured on default class-map (541843)
- L3 MTU configuration is not honored on a NAT enabled interface for multicast traffic. (546193)
SKUs Affected: CCS-720XP-24Y6, CCS-720XP-24ZY4, CCS-720XP-48Y6, CCS-720XP-48ZC2, DCS-7050CX3-32S, DCS-7050CX3M-32S, DCS-7050SX3-48C8, DCS-7050SX3-48YC12, DCS-7050SX3-48YC8 and DCS-7050TX3-48C8
- The Strata agent restarts after running "platform trident diag psr <port>" command. (546588)
- Routing between VXLAN VLANs with VRRP is not supported. (548160)
Series Affected: DCS-7050QX, DCS-7050SX, DCS-7050TX, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series
- The traffic load across ports comprising a port channel(s) may be uneven for certain traffic patterns and/or port channel configuration and state.

Workaround:

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

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

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

Workaround is to check the entry and packet hits for the entry using output

from "show platform trident tcam detail". This command shows all ACL entries in TCAM and packet hits per entry. (563228)

- Enabling L2 protocol forwarding of LLDP on a given port will implicitly enable forwarding of EAPOL and MKA packets on the same. (569478)
- Deny action is not supported on ACLs configured for filtering based packet sampling used for mirroring packets to the collector. (582014)
- After performing SSU from a PTP unsupported release to a PTP supported release, PTP will not be functional.
The workaround is to reset the forwarding switch by issuing the CLI command "platform trident reset" after the SSU and prior to enabling PTP. This will cause all interfaces to flap and a brief traffic outage to occur. (582970)
Series Affected: DCS-7060X4 Series
- BGP packets are ignored by IP counters if hardware IP counters are not enabled. (599275)
- When Port Security configuration is removed from an interface while StrataL2 agent is initializing , we can end up in a state where some MAC addresses on the interface do not get aged out and can also not be cleared via configuration.
A workaround is to flap the interface to remove the MAC addresses. (635439)
- Packets egress mirrored to a Greenspan destination with subinterface nexthops may have the internal VLAN in the inner VLAN portion of the packet. (642351)
- The 32-bit image of EOS does not support the 7050X4, 7358X4, 7060X5, and 7388X5 series switches. Use the 64-bit version instead (646587)
- When MACsec encryption traffic spans across multiple queues encrypted packets may be sent out of order. (669251)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- The outPktCtrl MACsec counters will not increment. (671159)
Series Affected: CCS-750 Series
- When the packets matching FBPS rule are also marked for drop due to other ACL rules, the packets are not mirrored and dropped in hardware. Workaround is to remove the configuration that results in dropping these packets. (673800)
SKUs Affected: CCS-720XP-48ZC2-F
- Mirroring packets to CPU can cause congestion on CPU queues and disrupt control plane traffic. It is not recommended to leave this feature on for extended periods of time. (675773)

- When attempting to VLAN mirror packets ingressing a VXLAN edge VLAN, the VXLAN encapsulated packets that ingress the core interface and are destined to be decapsulated to the edge VLAN will not be mirrored. (679353)
Series Affected: DCS-7050QX, DCS-7050QX2, DCS-7050SX, DCS-7050SX2, DCS-7050TX, DCS-7050TX2, DCS-7260CX3 and DCS-7260QX Series
- The maximum amount of VLAN sources allowed in a VLAN mirror session depends on the configuration's TCAM usage and the number of VLAN mirror sessions configured. A maximum of up to 3 mirror sessions are supported. Using a configuration that utilizes more VLAN sources than supported will result in VLAN mirroring sessions not being programmed. (680940)
Series Affected: CCS-710P, CCS-720XP, DCS-7010TX, DCS-7050CX3, DCS-7050SX3 and DCS-7050TX3 Series
- DirectFlow redirecting to interface with 'switchport source-interface tx' configured is unsupported. (681960)
- The outPktsUntagged counter will increment for unencrypted traffic egressing an interface even when the interface is not MACsec enabled. (682111)
Series Affected: CCS-750 Series
- The inPktsBadTag MACsec counter will not increment. Instead the drops will be accounted under the inPktsNotValid MACsec counter. (695748)
Series Affected: CCS-750 Series
- A chassis may only contain all 7300X3 or all 7300X linecards. Combining 7300X3 or 7300X linecards may result in repeated restarts of the forwarding agents, and links may stay down. Power cycle the system after removing the incompatible cards to recover. (699880)
Series Affected: DCS-7300X and DCS-7300X3 Series
- If "dst-ip" key-field is removed from "acl vlan ip ingress" feature in TCAM profile then, implicit link-local-multicast TCAM entry will not be programmed for ACL applied to SVI in unshared mode.

Workaround:

To overcome this issue add "dst-ip" key-field to "acl vlan ip ingress" feature in TCAM profile. (704466)

Series Affected: 7300X3, 7368, CCS-710P, CCS-720XP, CCS-750, DCS-7050CX3, DCS-7050SX3, DCS-7060PX4 and DCS-7260CX3 Series

- 7300-SUP/7300-SUP-D and 7300-SUP2-D supervisors are incompatible with each other and cannot be mixed in the same chassis. (728201)
SKUs Affected: 7300-SUP, 7300-SUP-D and 7300-SUP2-D
- The CLI "show hardware resource IpostLagLagMember" displays current usage statistics of hardware resources.
That hardware table may appear partially used even after some/all LAG(s)

that were active before have been deleted. This is expected behaviour with no other adverse effects. (729756)

- In PTP boundary or generalized PTP mode, configuring the `ptp forward-unicast` feature with interfaces using IPv6 transport is not supported. A workaround would be to use the IPv4 or L2 transport. (730817)
- If any currently unsupported mirroring configuration options such as ACLs, truncation, multiple destinations, or rate limits are entered in such a way that bypasses the CLI guards (such as config replace or editing the startup configuration), these options may be displayed as configured and in-use but will not be in effect.

Workaround: Remove any unsupported mirroring options from the startup configuration and/or un-configure them via the CLI. (731708)

- In cases of failed autonegotiated link up, link training may report as "off" on the primary lane. This is cosmetic. (731948)
- Configuring PTP MACsec bypass using the per-profile "ptp bypass" command is not supported.

Enable the "ptp bypass" command in MAC security configuration mode to have PTP packets bypass MACsec processing. (755578)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8, CCS-722XPM-48ZY8-F, CCS-755-CH and CCS-758-CH

- Only 100G copper downlink and AOC uplinks are supported on this Reverse Airflow platform. (758569)
SKUs Affected: DCS-7060DX5-64S-R
- When the switch is configured in ALPM mode and when the default route is forwarded over a VXLAN tunnel, L3 agent forwarding restart may result in some routes not getting programmed in the hardware. There is no workaround for this issue (764586)
- Ports with the speed set to 10Mbps do not recirculate packets when they are used in Recirc Channels. (766783)
- Per interface L3 MTU CLI configuration is not supported.

Workaround is to use global L3 MTU configuration in "interface defaults" mode. (777699)

Series Affected: CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-722XPM, CCS-750 and DCS-7010TX Series

- If the same ACL is applied on control plane and data plane, the byte counter for the ACL will be incremented for only data plane packet hits while the packet counter will be incremented for both data plane and

control plane packet hits.

A workaround is to create different ACLs to apply on control plane and data plane. (788296)

- Mirroring to CPU dataplane traffic using a non-default traffic-class will use the CPU queue of that traffic class, instead of the "Other" CPU queue. Workaround is to mirror to port instead of CPU destination. (791599)
Series Affected: 7368, DCS-7060X, DCS-7060X2 and DCS-7260X Series
- %ENVMON-3-TEMPSENSORINVALID and %ENVMON-3-TEMPSENSOROK could be reported after overheating of Linecard of Fabric card. The temperature reported by temperature sensors from disabled cards may also report low temperature in "show system environment temperature" CLI command output. (793777)
SKUs Affected: 7320X-32C-LC
- A feature using TCAM resources will not be able to use the last entry in the last slice allocated to it. For example, if a TCAM slice has 512 entries free, and if there are 2 unused slices, a feature using these slices, will be able to program only $(512 \times 2) - 1 = 1023$ entries as the last entry in last slice can not be used. (815178)
Series Affected: 7388 Series
SKUs Affected: 7358X4-SC, 7388X5-SC, DCS-7050DX4-32S, DCS-7050DX4-32S-F, DCS-7050PX4-32S, DCS-7050PX4-32S-F, DCS-7060DX5-64, DCS-7060DX5-64-F, DCS-7060DX5-64S, DCS-7060DX5-64S-F, DCS-7060DX5-64S-R, DCS-7060PX5-64E and DCS-7060PX5-64E-F
- This supervisor module is only supported with 7300X3 linecards. (826371)
SKUs Affected: 7300-SUP2-D
- The 64-bit image of EOS does not support the 7010T series of switches. Use the 32-bit version instead (846011)
- MTU configured on the egress interface is not enforced for the routed IP multicast traffic. This affects both regular IP multicast traffic as well as multicast traffic routed through tunnels. (938808)
- Packets destined to an L3 interface in a VRF that does not have routing enabled, are copied to the CPU on the `other` queue instead of the `self-ip` queue. This can potentially lead to packet loss for sessions on that L3 interface.
Workaround is to enable IP routing on that VRF. (938852)
- When switch is configured with multi VTEP MLAG, all traffic that are head end replicated will use the MLAG shared IP address. (989109)
- With DLB enabled and Strata-FixedSystem agent shutdown, link flaps from peer will cause complete traffic drops due to DLB being non functional (1000399)
- Ingress Filtered Mirroring with Inner VLAN ID is not supported (1003362)

- In case of CPU queue exhaustion, some features in startup config may not work due to CPU queue unavailability in a deterministic way. (1043368)
- MSS-G and common VARP MAC configuration in the network may result in traffic drops on L2 VLANs.

Workaround: Configure different VARP MAC on affected devices with L2 VLANs. (1045376)

- Port-channel negotiating LACP goes down when egress MAC ACL is applied on the port-channel interface.
To workaround this issue either the ACL can be removed from the interface or, a permit LACP rule can be added to the end of the ACL. (1051026)

CCS-750 Series

- Mirroring PTP packets when PTP is enabled is not supported. (738080)
- PTP two-step E2E and P2P transparent clock mode is not supported. (738097)
- Dropping PTP management message using the per-interface `ptp management drop` configuration is only supported in boundary clock mode. (752342)
- When PTP is configured, PTPv1 packet forwarding is not supported. (752367)
- Switches acting as either a Boundary Clock or Transparent Clock, downstream devices may see higher jitter in PTP time of up to +/- 1000ns (1025195)
SKUs Affected: CCS-755-CH and CCS-758-CH

7300X3, 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)
- Mirror To GRE destinations with multiple ECMP nexthops or LAG is not supported. A single link of the LAG or ECMP nexthops will receive all mirror traffic. (639980)

Series Affected: 7368 and 7388 Series

SKUs Affected: DCS-7060DX4-32 and DCS-7060PX4-32

- The 64-bit image of EOS does not support 7260QX series switches. Use the 32-bit version instead. (723646)
Series Affected: DCS-7260QX Series

DCS-7260X3 Series

- Configuring 100G, 50G, 40G, 25G and 10G simultaneously might result in unexpected packet forwarding behavior. (235075)
- Load Balance Number based hashing is not supported with L3 EVPN. (837409)

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)
- 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)
- Per vlan routing and sub interface are not supported together. (486439)
- 3m QSFP-DD passive copper cables are not supported in ports 1-4 and 29-32 (639613)
SKUs Affected: DCS-7060DX4-32
- Double Dot1Q tagged packets with IP payload may not be classified as IP packets in hardware. (813027)
- When a device behind phone/hub/switch connected to an Arista switch is authenticated via Dot1x MAC Based Authentication, it cannot be moved to different port if the other port already has a different supplicant authenticated on it.

A workaround is to configure 'dot1x timeout idle-host' (1066149)

CCS-710P, CCS-720DP and CCS-720DT Series

- Triple-wide TCAMs are not supported. (603246)
Series Affected: CCS-710P Series
- Generation of PFC frames is not supported. (699868) (1)

CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series

- XPN cipher suites are not supported. (590170)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- Non-unicast traffic egressing on Port-Channel interfaces that are MACsec enabled do not use hashing to load balance. A single Port-Channel member will be designated to transmit non-unicast traffic for a given replication group (VLAN or multicast group). (604539)
Series Affected: CCS-722XPM Series
- "show platform trident counters" output is misleading for certain counters with MACsec encryption enabled. The workaround is to use the "show interfaces counters discards" and "show mac security counters interface" commands where possible. (623564)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- Generation of PFC frames is not supported. (699868) (1)
- VLAN tagged MACsec is not supported (852797)
Series Affected: CCS-722XPM Series
- DirectFlow is not supported on MACsec enabled interfaces. (907033)
Series Affected: CCS-722XPM Series
- If "traffic unprotected allow" is configured, and MACsec falls back to unprotected traffic, any encrypted packets will be dropped. (1007567)
Series Affected: CCS-722XPM Series

CCS-720DF and CCS-720XP Series

- When flow group ACL rules are changed, flow group counters displayed in "show platform trident flow counters" will continue to have counters incremented for active flows corresponding to older flow groups till the flows are aged out. (552450)
- On interfaces with speeds less than 10G, PFC cannot be transmitted. (689087)

CCS-750 Series

- The switch is not capable of generating pause flowcontrol frames. (480688)
Series Affected: CCS-750 Series
- Per interface L3 MTU configuration is not supported.
The workaround is to use the global L3 MTU configuration command instead. (520825)

- The supervisor uplink SCD does not support hitless FPGA upgrade. (528417)
SKUs Affected: CCS-750-SUP100 and CCS-750-SUP25
- The switch is not capable of generating Priority Flow Control (PFC) frames. (555294)
Series Affected: CCS-750 Series
- Shaper below 1 Mbps is not supported. (605407)
- MACsec encrypted flow control packets and priority-flow-control packets that are received will not be visible in "show flowcontrol" or "show priority-flow-control counters". (670539)
- When a mirror session has a LAG member as the source, packets sent to the LAG will not be mirrored. (712186)
- TCAM profile does not effect any saving of TCAM slices / banks for Egress MAC ACL and IP ACL applied to Ethernet, LAG ports (723449)
SKUs Affected: CCS-755-X3-SC and CCS-758-X3-SC
- Ports on this linecard share a total sustainable bandwidth.

Though this linecard can absorb concurrent bursts received on all ports, the total sustained bandwidth is limited to a total of 148 Gbps for an average frame size of at least 276 Bytes.

To optimize sustained bandwidth availability, the recommendation is to distribute the bandwidth equally across the groups of eight ports. The groups are: Ethernet1-8, Ethernet9-16, Ethernet17-24, Ethernet25-32, Ethernet33-40, Ethernet41-48

For example, if there are 24x5G and 24x2.5G devices to connect, it is recommended that 4x5G and 4x2.5G be connected to each of the 6 groups of front-panel ports. (788129)

SKUs Affected: CCS-750X-48ZXP-LC

- Stateful Switchover triggered by removing the active supervisor will cause a long data-plane outage for uplink traffic destined to the removed uplinks. Workaround is to unplug the uplink interfaces, wait for the link to go down and then remove the active supervisor. (839923)
Series Affected: CCS-750 Series
- If a device is behind a switch or hub and is authenticated via Dot1x MAC Based Authentication it will not be able to authenticate via Dot1x MBA on a different interface.

A workaround is to configure 'dot1x timeout idle-host' (839960)

- When Greenspan session is configured with trunk port or sub-interface as next hop destination, the mirrored packets will have missing dot1q and VLAN tag. (979999)

7300X3, AS7326, AS7726 and DCS-7050X3 Series

- The egress queue and per hop latency of the INT terminator node are always shown as zero. (562000)
Series Affected: DCS-7050CX3 and DCS-7050SX3 Series

7358 and DCS-7050X4 Series

- Sampled counters are not required. (437955)
- Color based DSCP or CoS assignment is not supported. (679840)
- Packets destined for the control plane may appear to be dropped within the ingress pipeline. This will appear in the output of `show platform trident counters` as a RX drop, and is expected due to more drop events available in the pipeline. (685445)
- On interfaces Ethernet33 and Ethernet34, PFC cannot be transmitted. (737611)
SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4-32S-F and DCS-7050PX4-32S-F
- Egress ACL does not match on policy rewritten DSCP nor remarked ECN values. (741424)
Series Affected: DCS-7050X4 Series
- Maximum replications for a non unicast packet is limited to 1023. (750699)
Series Affected: DCS-7050X4 Series
- A maximum of two QoS features can be configured on Trident4X platforms. (772245)
- Mirror to GRE with forwarding drop is not supported on T4X11.
Workaround is to use Mirror on Drop feature for capturing drop packets. (930244)
Series Affected: DCS-7050X4 Series
SKUs Affected: 7358X4-SC

7388 and DCS-7060X5 Series

- Mirror to GRE with forwarding drop is not supported on Tomahawk4.
Workaround is to use Mirror on Drop feature for capturing drop packets. (930240)
- Metadata fields of LNS, Congestion, QueueID, Queue Depth, Queue Transmit Byte count and MMU Stats are not supported (934700)

Transceiver Series

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

```
* "speed auto" / "speed auto 1gfull" / "speed auto 1000full" : enable autoneg, advertise only 1G
* "speed auto 100full" : enable autoneg, advertise only 100M
* "speed 100full" : disable autoneg, force speed to 100M (159401)
SKUs Affected: 1000BASE-T
```

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

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

- During reboot, a peer may see a link-flap on ports using 1000BASE-T transceivers.
The workaround is to configure the link-debounce period to 30 seconds for the link-up transition on the peer interfaces. (268080)
SKUs Affected: 1000BASE-T
- For SKUs CAB-O-2Q-400G, CAB-O-4Q-400G, CAB-D-2Q-400G and CAB-D-4Q-400G, when operating at 25G or 50G NRZ speeds, forward error-correction should be manually configured (enabled or disabled) on both sides of the link as the default error-correction on both ends may not match. (392124)
- SFP-10G-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 10gfull/"speed auto 10000full" : enable autoneg, advertise only 10G
* "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 (413941)
SKUs Affected: SFP-10G-T
- Revision 20 OSFP-400G-2FR4 transceivers may not correctly report the operational value of the configured Tx Input Equalization. This can be observed in 'show interfaces transceiver tuning detail', in which the Operational Tx Input Equalization may not match the Configured value. If the 'Tx Input Equalization' Configured column shows auto, automatic Tx Input Equalization is enabled, despite the Operational value shown. (423045)
SKUs Affected: OSFP-400G-2FR4
- Using a default speed configuration on any interface associated with a 400G optical transceiver may prevent all interfaces associated with the transceiver from linking up.

The workaround is to explicitly configure speed on all interfaces associated with the 400G optical transceiver. (426232)

- FEC should be configured manually when using OSFP/QSFP-DD to QSFP200/SFP50 breakout cables at 25G or 50G-2 speeds (432776)
 SKUs Affected: CAB-D-2Q-400G, CAB-D-4Q-400G, CAB-D-8S-400G, CAB-O-2Q-400G, CAB-O-4Q-400G and CAB-O-8S-400G

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

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

- Interfaces with BCM82391 PHY do not support 1000BASE-T transceivers. (535121)
- QSFP200 cables and optics encoded using the CMIS specification are not recognized in QSFP100 ports. (536318)
- OSFP & QSFP-DD port LEDs are not supported in hardware LED CLI config mode (650615)
- Interfaces with the SFP-10G-MRA-T 10GBASE-T module may link up during system reload if the peer is enabled. During the reload the port will not pass traffic. This will occur on enabled or disabled ports. (651386)
- SFP 10GBASE-T transceivers which do not support Soft RX_LOS or Soft TX_DISABLE are not supported on DSFP ports. (702025)
 SKUs Affected: SFP-10G-T
- The rate adapting transceivers SFP-10G-MRA-T, SFP-10G-RA-1G-LX, and SFP-10G-RA-1G-SX will not allow PAUSE or PFC frames to pass to and from the peer device. (881411)
 SKUs Affected: SFP-10G-MRA-T
- 400GBASE-ZR transceivers do not support Zero Touch Provisioning. (882209)
 SKUs Affected: OSFP-400G-ZR, OSFP-400G-ZRP, QDD-400G-ZR and QDD-400G-ZRP
- Alert state and count of temperature sensor in "show system environment temperature transceiver detail" CLI output is never updated (886972)
- The SFP-1G-T pluggable 1000BASE-T SFP is not capable of advertising PAUSE or ASM PAUSE during auto negotiation. This does not impact the ability to send or receive PAUSE per the specific switch's abilities and flow control configuration. (904116)
 SKUs Affected: 1000BASE-T

- After using "reload fast-boot" to upgrade to a new EOS release, certain transceivers may have new default tuning values that are not applied until the next system reboot or after the transceiver is reinserted. (1032377)

7300X3, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX 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)
- Directflow configuration with "set vlan", but no "action set output interface" configured is not supported. (345852)
- MAC Egress ACL configurations matching IP packets may not work when the IPv4/v6 per-interface counter feature is also configured. (411243)
- CLI output of "show platform trident l3 shadow mpls-vc-and-swap-label" will show push as MplsLabelAction even for swap operation. (416134)
- Packets dropped by the MACsec engine may not be reported in the CLI (420085)
SKUs Affected: DCS-7050CX3M-32S
- 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
- When IPV6 VXLAN underlay core interface is an access port on a SVI, VXLAN encapsulated packets may be sent out with VLAN tag in the outer ethernet header. The work around is to configure the port as a trunk port. (533929)
- Packets mirrored to GRE with a size within 16 bytes of the MTU limit may be dropped on egress due to exceeding the MTU size after the mirrored packet is encapsulated. (592957)
- The UDP header checksum is not set correctly for PTP packets after NAT (601036)

- "Dont fragment" bit is not set in the outer IP header of the VXLAN encapsulated packets when the payload is non-IP. (605003)
- When Postcard Telemetry is enabled, EVPN Multicast will not be supported. (614072)
Series Affected: 7300X3, CCS-720XP, DCS-7050SX3, DCS-7050TX3 and DCS-7050X3 Series

- ACL resource sharing feature is not supported for NAT translations of tunneled packets like VXLAN-tunneled packets.

Workaround is to turn off ACL resource sharing: "no hardware nat static access-list resource sharing" (620222)

- If the destination MAC address of a NAT flow is not learnt on an interface, then the packet will be dropped. Flooding the translated packet on all interfaces belonging to the VLAN is not supported.

Workaround is to adjust the ARP and MAC address timeouts, so that MAC addresses is learnt again. (642094)

- DirectFlow actions on interfaces configured as dot1q-tunnel are not supported. (686339)
- If an egress ACL is filtering on the NAT flow's source IP, then the ACL rules have to permit both original source IP and translated IP to avoid NAT flows getting dropped. (688883)
- With EVPN Multicast with underlay as HER configured, known multicast traffic might not be load balanced across all ECMP links. (709053)
- Ingress mirroring double tagged packets to CPU will have the outer VLAN tag stripped.
Workaround is to use a port mirroring destination instead. (722891)
- Packets with Broadcast source MAC may be copied to CPU instead of being dropped. (745001)
- For features which use egress TCAM the match is on outer header for a tunnelled encapsulated packet. There is no workaround for this bug because this is a hardware limitation. (801113)
- 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. (822757)

- When ingress sFlow is enabled along with VXLAN data plane learning, VXLAN remote VTEP Learning could fail resulting in remote MACs not being learnt for the sampled packets. (825322)
- Under most scenarios, traffic that enters a VTEP VXLAN-encapsulated is not allowed to also leave the VTEP VXLAN-encapsulated (whether destined to the same or a different VTEP). This split-horizon filtering is broken on t3x-based devices when using IPv6 underlay if the target MAC address is known unicast on the receiving VTEP (that is, if the target MAC address is learned as a remote MAC behind a VXLAN tunnel). Split-horizon filtering for BUM traffic is working as expected. This issue does not affect IPv4 VXLAN underlay. (831822)
- Layer 2 Multicast forwarding will experience higher traffic outage during SSU. Outage will continue till the router re-learns the multicast receivers post SSU operation. There is no workaround for this problem (833449)
- With EVPN multicast configured with IPv4 underlay tunnels, overlay IPv6 unknown multicast traffic may not get flooded in broadcast domain. Work-around is to use VLANs without IGMP-proxy for unknown IPv6 multicast traffic. (869000)
- Packets that match on a dynamic NAT ACL will be sent to CPU if the packets are egressing out of any other interface that also has NAT rules configured. The workaround is to make the NAT ACL less greedy by either reducing the scope of matching IPs or by adding deny rules for specific destination IPs that should not be NATed. (900348)
- A warning will not be emitted when 'dot1x mac based access-list' is missing, but 'dot1x aaa unresponsive action traffic allow access-list' or 'dot1x authentication failure action traffic allow access-list' is configured. (918137)
- When one or more L2 MAC addresses cannot be programmed in the hardware due to events such as hash collisions, user will not be able to see them using 'show platform trident mac-address-table missing' command. There is no workaround for this CLI output issue. (968260)

7358 and DCS-7050X4 Series

- If 300K routes with 128-way ECMP are programmed, system can encounter out of memory conditions. Agents may restart as the scale exceeds platform capabilities. There is no workaround for this issue. (785832)
- Greenspan packets where the size of the mirror payload plus the Greenspan header exceeds the configured MTU size of the ingress interface will be fragmented on a second pass of the pipeline.

Note that high rates of fragmentation can result in high CPU utilization.
(791223)

- By default IPv6 packets are set to not fragment. If an added IPv6 Greenspan header causes a packet to exceed the MTU size of the destination next hop, then it will be dropped. Workaround is to increase MTU size of the nexthop port. (818665)
- Filtered mirrored Greenspan packets cannot be egress sampled by sFlow. (863261)
- Trident4 does not have configurable debug counters, so IP Counters are only possible with Flex Counters. (892598)
- Early shipments of system bundles with this switch card were loaded with the incorrect EOS-4.29.0.2F binary (32-bit) instead of the EOS64-4.29.0.2F binary (64-bit). Unexpected traffic drops may be seen if the 32-bit binary is used. Upgrade to a 64-bit version of EOS to avoid such traffic drops. (909191)

SKUs Affected: 7358X4-SC

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.