

Arista Networks EOS 4.23.16M Release Notes

Version: 0.1

03 February 2025

Table of Contents

EOS 4.23.16M Release Highlights

New Platforms and Hardware

SNMP MIBs

SNMP Traps

Supported Hardware

Reference to MLAG ISSU Upgrade/Downgrade Table

Resolved Caveats in 4.23.16M

General

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

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

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

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

Known Software Caveats in 4.23.16M

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

Containerized EOS

vEOS Router / CloudEOS

CVX

General

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MPLS

Multicast

vEOS Router / CloudEOS

DCS-7150 Series

DCS-7170 Series

DCS-7170 Series

CCS-720XP Series

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

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

DCS-7500E Series

DCS-7280R and DCS-7500R Series

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

DCS-7020 Series

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

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

DCS-7280R and DCS-7280R2 Series

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

7368 and DCS-7060X4 Series

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

Transceiver Series

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

DCS-7160 Series

Limitations and Restrictions in 4.23.16M

BGP EVPN L2/L3 VXLAN/MPLS

CVX

General

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

MPLS

Multicast

vEOS Router / CloudEOS

DCS-7170 Series

CCS-720XP Series

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

DCS-7280R and DCS-7500R Series

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

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

DCS-7280R and DCS-7280R2 Series

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

7368 and DCS-7060X4 Series

Transceiver Series

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

DCS-7160 Series

Conditions and Impacts

EOS 4.23.16M Release Highlights

New Platforms and Hardware

- None

SNMP MIBs

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

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

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

SNMP Traps

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

Supported Hardware

Fixed System

* DCS-7150S-24-CL

* DCS-7280QR-C36-M-R

* DCS-7050QX-32S-SSD-F

* DCS-7150S-24	* DCS-7280QR-C36-R	* DCS-7050QX-32S-SSD-R
* DCS-7150S-24-CL-F	* DCS-7280SR-48C6	* DCS-7050SX-64
* DCS-7150S-24-CL-R	* DCS-7280SR-48C6-F	* DCS-7050SX-64-F
* DCS-7150S-24-CL-SSD-F	* DCS-7280SR-48C6-M-F	* DCS-7050SX-64-R
* DCS-7150S-24-F	* DCS-7280SR-48C6-M-R	* DCS-7050SX-64-SSD-F
* DCS-7150S-24-R	* DCS-7280SR-48C6-R	* DCS-7050SX-64-SSD-R
* DCS-7150S-24-CL-SSD-R	* DCS-7280TR-48C6-R	* DCS-7050TX-128-F
* DCS-7150S-52-CL	* DCS-7280TR-48C6-F	* DCS-7050TX-128-R
* DCS-7150S-52-CL-F	* DCS-7280TR-48C6-M-F	* DCS-7050TX-48
* DCS-7150S-52-CL-R	* DCS-7280TR-48C6-M-R	* DCS-7050TX-48-F
* DCS-7150S-52-CL-SSD-F	* DCS-7280QR-C72	* DCS-7050TX-48-R
* DCS-7150S-52-CL-SSD-R	* DCS-7280QR-C72-F	* DCS-7050TX-48-SSD-F
* DCS-7150S-64-CL	* DCS-7280QR-C72-M-F	* DCS-7050TX-48-SSD-R
* DCS-7150S-64-CL-F	* DCS-7280QR-C72-M-R	* DCS-7050TX-64
* DCS-7150S-64-CL-R	* DCS-7280QR-C72-R	* DCS-7050TX-64-F
* DCS-7150S-64-CL-SSD-F	* DCS-7280SR2-48YC6	* DCS-7050TX-64-R
* DCS-7150S-64-CL-SSD-R	* DCS-7280SR2-48YC6-F	* DCS-7050TX-64-SSD-F
* DCS-7150SC-24-CLD	* DCS-7280SR2-48YC6-M-F	* DCS-7050TX-64-SSD-R
* DCS-7150SC-24-CLD-F	* DCS-7280SR2-48YC6-M-R	* DCS-7050TX-72-SSD-F
* DCS-7150SC-24-CLD-R	* DCS-7280SR2-48YC6-R	* DCS-7050TX-72-SSD-R
* DCS-7150SC-64-CLD	* DCS-7020TR-48	* DCS-7050TX-96-SSD-F
* DCS-7150SC-64-CLD-F	* DCS-7020TR-48-F	* DCS-7050TX-96-SSD-R
* DCS-7150SC-64-CLD-R	* DCS-7020TR-48-R	* DCS-7010T-48
* DCS-7170-64C	* DCS-7280QRA-C36S-F	* DCS-7010T-48-F
* DCS-7170-64C-F	* DCS-7280QRA-C36S-R	* DCS-7010T-48-R
* DCS-7170-64C-R	* DCS-7280SR2A-48YC6	* DCS-7050SX-72
* DCS-7170-32C	* DCS-7280SR2A-48YC6-F	* DCS-7050SX-72-F
* DCS-7170-32C-F	* DCS-7280SR2A-48YC6-M-F	* DCS-7050SX-72-R
* DCS-7170-32C-M-F	* DCS-7280SR2A-48YC6-M-R	* DCS-7050SX-72-SSD-F
* DCS-7170-32C-M-R	* DCS-7280SR2A-48YC6-R	* DCS-7050SX-72-SSD-R
* DCS-7170-32C-R	* DCS-7280SRA-48C6	* DCS-7050SX-96
* DCS-7170-64C-M-F	* DCS-7280SRA-48C6-F	* DCS-7050SX-96-F
* DCS-7170-64C-M-R	* DCS-7280SRA-48C6-M-F	* DCS-7050SX-96-R
* DCS-7170-32CD	* DCS-7280SRA-48C6-M-R	* DCS-7050SX-96-SSD-F
* DCS-7170-32CD-F	* DCS-7280SRA-48C6-R	* DCS-7050SX-96-SSD-R
* DCS-7170-32CD-R	* DCS-7280TRA-48C6	* DCS-7050TX-128-SSD-F
* FAN-7001C	* DCS-7280TRA-48C6-F	* DCS-7050TX-128-SSD-R
* PWR-1100AC	* DCS-7280TRA-48C6-M-F	* DCS-7050TX-72-F
* PWR-1100AC-F	* DCS-7280TRA-48C6-M-R	* DCS-7050TX-72-R
* PWR-1100AC-R	* DCS-7280TRA-48C6-R	* DCS-7050TX-96-F
* PWR-750AC	* DCS-7020TRA-48	* DCS-7050TX-96-R

* PWR-750AC-F	* DCS-7020TRA-48-F	* DCS-7010T-48-DC
* PWR-750AC-R	* DCS-7020TRA-48-R	* DCS-7010T-48-DC-F
* PWR-500-DC	* DCS-7280CR2A-60	* DCS-7010T-48-DC-R
* PWR-500-DC-F	* DCS-7280CR2A-60-F	* DCS-7260QX-64
* PWR-500-DC-R	* DCS-7280CR2K-60	* DCS-7260QX-64-F
* PWR-500AC	* DCS-7280CR2K-60-F	* DCS-7260QX-64-R
* PWR-500AC-F	* DCS-7280SRAM-48C6	* DCS-7260QX-64-SSD-F
* PWR-500AC-R	* DCS-7280SRAM-48C6-F	* DCS-7260QX-64-SSD-R
* PWR-745AC	* DCS-7280SRAM-48C6-R	* DCS-7050SX-72Q
* PWR-745AC-F	* DCS-7280CR2-60	* DCS-7050SX-72Q-F
* PWR-745AC-R	* DCS-7280CR2-60-F	* DCS-7050SX-72Q-R
* FAN-7010	* DCS-7280CR2A-30	* DCS-7050TX-72Q
* PWR-1900-DC	* DCS-7280CR2A-30-F	* DCS-7050TX-72Q-F
* PWR-1900-DC-F	* DCS-7280CR2K-30	* DCS-7050TX-72Q-R
* PWR-1900-DC-R	* DCS-7280CR2K-30-F	* DCS-7050QX2-32S
* PWR-1900AC	* DCS-7280QRA-C36S-M-F	* DCS-7050QX2-32S-F
* PWR-1900AC-F	* DCS-7280QRA-C36S-M-R	* DCS-7050QX2-32S-R
* PWR-1900AC-R	* DCS-7280SR2K-48C6	* DCS-7050SX2-128
* PWR-747AC	* DCS-7280SR2K-48C6-M-F	* DCS-7050SX2-128-F
* PWR-747AC-F	* DCS-7280SR2K-48C6-M-R	* DCS-7050SX2-128-R
* PWR-747AC-R	* DCS-7020SR-24C2	* DCS-7050TX2-128
* FAN-7002H	* DCS-7020SR-24C2-F	* DCS-7050TX2-128-F
* FAN-7002H-R	* DCS-7020SR-24C2-R	* DCS-7050TX2-128-R
* FAN-7000H-R	* DCS-7280CR2M-30	* DCS-7050SX2-72Q
* PWR-1600AC	* DCS-7280CR2M-30-F	* DCS-7050SX2-72Q-F
* PWR-1600AC-F	* DCS-7280SRM-40CX2	* DCS-7050SX2-72Q-R
* FAN-7001D	* DCS-7280SRM-40CX2-F	* DCS-7060CX-32S
* FAN-7001D-F	* DCS-7280SRM-40CX2-R	* DCS-7060CX-32S (Prime)
* PWR-400-DC-BLUE	* DCS-7020SR-32C2-F	* DCS-7060CX-32S-F
* PWR-400-DC-RED	* DCS-7020SR-32C2-R	* DCS-7060CX-32S-F (Prime)
* FAN-7001C-F	* DCS-7020SRG-24C2	* DCS-7060CX-32S-R
* PWR-400AC-BLUE	* DCS-7020SRG-24C2-F	* DCS-7060CX-32S-R (Prime)
* FAN-7011M	* DCS-7020SRG-24C2-R	* DCS-7060CX2-32S
* FAN-7011M-F	* DCS-7280CR3-32P4	* DCS-7060CX2-32S-F
* FAN-7011M-R	* DCS-7280CR3-32P4-F	* DCS-7060CX2-32S-R
* PWR-1021-AC-RED	* DCS-7280CR3-32P4-M-F	* DCS-7260CX-64
* PWR-511-AC-BLUE	* DCS-7280CR3-32P4-M-R	* DCS-7260CX-64-F
* PWR-511-AC-RED	* DCS-7280CR3-32P4-R	* DCS-7260CX-64-R
* PWR-621-AC-RED	* DCS-7280CR3K-32P4	* DCS-7260CX-64-SSD-F
* FAN-7001C-R	* DCS-7280CR3K-32P4-F	* DCS-7260CX-64-SSD-R
* PWR-400AC-RED	* DCS-7280CR3K-32P4-R	* DCS-7060SX2-48YC6

* FAN-7001DH	* DCS-7280CR3-32D4	* DCS-7060SX2-48YC6-F
* FAN-7001DH-F	* DCS-7280CR3-32D4-F	* DCS-7060SX2-48YC6-R
* PWR-1611-AC-RED	* DCS-7280CR3-32D4-M-F	* DCS-7260CX3-64
* PWR-1611-DC-RED	* DCS-7280CR3-32D4-M-R	* DCS-7260CX3-64-F
* FAN-7011H	* DCS-7280CR3-32D4-R	* DCS-7260CX3-64-R
* FAN-7011H-F	* DCS-7280CR3K-32D4	* DCS-7050CX3-32S-F
* FAN-7011H-R	* DCS-7280CR3K-32D4-F	* DCS-7050CX3-32S-R
* PWR-511-DC-BLUE	* DCS-7280CR3K-32D4-R	* DCS-7050SX3-48YC12
* PWR-511-DC-RED	* DCS-7280PR3-24	* DCS-7050SX3-48YC12-F
* PWR-1511-AC-BLUE	* DCS-7280PR3-24-F	* CCS-720XP-24ZY4
* PWR-1511-AC-RED	* DCS-7280PR3-24-M-F	* CCS-720XP-24ZY4-F
* PWR-1511-DC-BLUE	* DCS-7280PR3K-24	* CCS-720XP-48ZC2
* PWR-1511-DC-RED	* DCS-7280PR3K-24-F	* CCS-720XP-48ZC2-F
* PWR-1011-AC-RED	* DCS-7280DR3-24	* DCS-7050SX3-48YC8
* PWR-721-DC-RED	* DCS-7280DR3-24-F	* DCS-7050SX3-48YC8-F
* PWR-621-AC-BLUE	* DCS-7280DR3-24-M-F	* DCS-7050SX3-48YC8-R
* FAN-7012H-RED	* DCS-7280CR3-96	* DCS-7260CX3-64E
* PWR-1011-AC-BLUE	* DCS-7280CR3-96-F	* DCS-7260CX3-64E-F
* PWR-1011-DC-BLUE	* DCS-7280CR3K-96	* DCS-7260CX3-64E-R
* PWR-1011-DC-RED	* DCS-7280CR3K-96-F	* DCS-7060DX4-32
* PWR-2411-AC-RED	* DCS-7280CR3MK-32P4	* DCS-7060DX4-32-F
* PWR-2411-DC-RED	* DCS-7280CR3MK-32P4-F	* DCS-7060PX4-32
* FAN-7000	* DCS-7280CR3MK-32P4-R	* DCS-7060PX4-32-F
* FAN-7000-F	* DCS-7050TX-128	* CCS-720XP-24Y6-F
* FAN-7000-R	* DCS-7050TX-72	* CCS-720XP-48Y6-F
* PWR-460AC	* DCS-7050TX-96	* DCS-7050CX3-32S-SSD-F
* PWR-460AC-F	* DCS-7050CX3-32S	* DCS-7050CX3-32S-SSD-R
* PWR-460AC-R	* CCS-720XP-24Y6	* DCS-7050CX3M-32S
* PWR-460DC	* CCS-720XP-48Y6	* DCS-7050CX3M-32S-F
* PWR-460DC-F	* DCS-7050SX-128	* DCS-7050CX3M-32S-R
* PWR-460DC-R	* DCS-7050SX-128-F	* CCS-720XP-24Y6-R
* DCS-7280TR-48C6	* DCS-7050SX-128-R	* CCS-720XP-48Y6-R
* DCS-7280QRA-C36S	* DCS-7050SX-128-SSD-F	* DCS-7160-32CQ
* DCS-7020SR-32C2	* DCS-7050SX-128-SSD-R	* DCS-7160-32CQ-F
* DCS-7280SE-64	* DCS-7250QX-64	* DCS-7160-32CQ-R
* DCS-7280SE-64-F	* DCS-7250QX-64-F	* DCS-7160-48YC6
* DCS-7280SE-64-R	* DCS-7250QX-64-R	* DCS-7160-48YC6-F
* DCS-7280SE-72	* DCS-7250QX-64-SSD-F	* DCS-7160-48YC6-R
* DCS-7280SE-72-F	* DCS-7250QX-64-SSD-R	* DCS-7160-48TC6
* DCS-7280SE-72-R	* DCS-7050QX-32	* DCS-7160-48TC6-F
* DCS-7280SE-68	* DCS-7050QX-32-F	* DCS-7160-48TC6-R

- * DCS-7280SE-68-F
- * DCS-7280SE-68-R
- * DCS-7280CR-48
- * DCS-7280CR-48-F
- * DCS-7280QR-C36
- * DCS-7280QR-C36-F
- * DCS-7280QR-C36-M-F

- * DCS-7050QX-32-R
- * DCS-7050QX-32-SSD-F
- * DCS-7050QX-32-SSD-R
- * DCS-7050QX-32S
- * DCS-7050QX-32S-F
- * DCS-7050QX-32S-R

- * FAN-7002
- * FAN-7002-F
- * FAN-7002-R
- * FAN-7002H-F
- * FAN-7000H
- * FAN-7000H-F

Modular

* FAN-7002	* DCS-7500-SUP2	* 7800R3K-48CQ-LC
* FAN-7002-F	* DCS-7500-SUP2-D	* 7808R3-FM
* FAN-7002-R	* 7500R-36CQ-LC	* DCS-7800-SUP
* FAN-7002H-F	* 7500R-36Q-LC	* DCS-7800-SUP1A
* FAN-7000H	* 7500R-48S2CQ-LC	* DCS-7508
* FAN-7000H-F	* 7504R-FM	* DCS-7504
* PWR-3K-AC	* 7508R-FM	* 7300-SUP
* PWR-3K-AC-F	* 7512R-FM	* 7300-SUP-D
* PWR-3K-AC-R	* DCS-7512N	* 7300X-32Q-LC
* PWR-2700-DC	* 7500RM-36CQ-LC	* 7300X-64S-LC
* PWR-2700-DC-F	* 7500R-8CFPX-LC	* 7300X-64T-LC
* PWR-2700-DC-R	* 7500R2-18CQ-LC	* 7304X-FM
* PWR-3K-DC-BLUE	* 7500R2-36CQ-LC	* 7308X-FM
* PWR-3K-DC-RED	* 7516R-FM	* DCS-7304
* PWR-3KT-AC-BLUE	* DCS-7516-SUP2	* DCS-7308
* PWR-3KT-AC-RED	* DCS-7516-SUP2-D	* 7316X-FM
* PWR-D1-3041-AC-BLUE	* DCS-7516N	* DCS-7316
* PWR-2900AC	* 7500R2A-36CQ-LC	* 7320X-32C-LC
* 7500E-36Q-LC	* 7500R2AK-36CQ-LC	* 7324X-FM
* 7500E-48S-LC	* 7500R2AK-48YCQ-LC	* 7328X-FM
* 7500E-72S-LC	* 7500R2M-36CQ-LC	* 7300X3-32C-LC
* 7504E-FM	* 7500R2AM-36CQ-LC	* 7304X3-FM
* 7508E-FM	* 7500R3-36CQ-LC	* 7308X3-FM
* DCS-7500E-SUP	* 7504R3-FM	* 7300X3-48YC4-LC
* DCS-7500E-SUP-D	* 7508R3-FM	* 7368
* 7500E-12CM-LC	* 7500R3K-36CQ-LC	* 7368-16C
* 7500E-6C2-LC	* 7500R3-24P-LC	* 7368-SUP
* DCS-7504N	* DCS-7808-CH	* 7368-SUP-D
* 7500E-48T-LC	* 7800R3-36P-LC	* 7368X4-SC
* 7500E-12CQ-LC	* 7800R3-48CQ-LC	* 7304X3-FM2
* DCS-7508N	* 7800R3-48CQM-LC	* 7308X3-FM2
* 7500E-6CFPX-LC		

Transceiver

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

Reference to MLAG ISSU Upgrade/Downgrade Table

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

Resolved Caveats in 4.23.16M

General

- A PTP Management message with a certain TLV size close to 65535 will cause the PTP agent to restart unexpectedly. Agent will recover after the restart.
Workaround is to block such PTP Management messages from entering the PTP network. (638107)
- Certain ill-crafted certificates might cause infinite loop during verification. (674519)

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

- If a QoS policy with rules of different ACL types fails to program in hardware for one of the ACL types, the subset of rules that was programmed successfully may not match incoming traffic correctly if other QoS policies are configured and applied in the system. (571599)

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

- When PBR/BGP Flowspec/QoS/CPU policies are applied on SVIs, and the linecard of one of the SVI members gets removed and re-inserted, the SandAcl agent may restart unexpectedly.

A workaround is to remove the SVI member before its linecard gets removed and re-inserted. (532575)

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

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

The workaround is to configure forced speed. (522773)

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

- MAC addresses might not age out on port-channels spanning across multiple chips (616314)

SKUs Affected: DCS-7250QX-64, DCS-7250QX-64-F, DCS-7250QX-64-R, DCS-7250QX-64-SSD-F, DCS-7250QX-64-SSD-R, DCS-7260CX-64, DCS-7260CX-64-F, DCS-7260CX-64-R, DCS-7260CX-64-SSD-F, DCS-7260CX-64-SSD-R, DCS-7304 and DCS-7308

Known Software Caveats in 4.23.16M

Accelerated System Update

- Performing SSU when `/mnt/flash/persist/persistentRestartLog` doesn't exist may result in extended traffic outage due to a watchdog reset that can occur. (158117)
- SSU with Splunk enabled may result in extended boot times, which could lead to protocol timeouts or switch cold rebooting. (343629)

BGP EVPN L2/L3 VXLAN/MPLS

- In an EVPN deployment, ARP entries learned from MAC/IP advertisement routes may remain stale after removing the associate MAC-VRF or changing the BGP AS number.

The workaround is to restart the BGP agent. (578375)

- In a Multi-VTEP EVPN VXLAN MLAG IRB deployment, reloading both peers may result in the lost of dynamic VLANs on one or both peers required by VRF to VNI configurations.

A workaround is to restart the VXLAN agent. (586826)

Containerized EOS

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

vEOS Router / CloudEOS

- A bug in Google Cloud KVM implementation causes Live Migration performed during a maintenance event to reload the vEOS router.

Until a new KVM instance is deployed the only workaround is to modify `/mnt/flash/kernel-params` to disable hyperthreading using the `nosmt` parameter. Please contact Arista for detailed steps to apply the workaround. (579590)

- In certain CloudEOS deployments the VM may be oversubscribed and thus subject to processing delays while sending packets to the CPU. This primarily affects BFD control packets and result in the BFD detecting a link

flap.

One workaround is to increase the interval, min-rx and multiplier to larger values for each affected interface. For example, use the command: "bfd interval 500 min-rx 500 multiplier 5" to increase the timeout to 2.5 seconds in the interface sub-mode. (581475)

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)
- When a group on NSX-T containing more than 1000 hosts, policy rules for the group on CVX will only generate ACLs for the first 1000 hosts in that group. (459644)
- The Policy Control Service (PCS) fails to apply any security policy configured on NSX-T on a switch interface after the switch configuration is modified to remove all tags from an interface followed by re-adding tags, using the configuration replace feature. (463958)
- When 'client state preserve' feature is configured on CVX, state of the primary switch of an MLAG VTEP will not be preserved upon disconnection. If the primary switch did not reboot and disconnected from CVX due to network faults, it may cause disruption of VXLAN control plane leading to VXLAN traffic loss.

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

- In a Macro-Segmentation Service (MSS) deployment with Panorama, uncommitted changes to the list of managed devices may be used by MSS to program rules on the switches. (536412)

General

- OpenFlow Flows may not be programmed in the right priority order after changing OpenFlow Table Profile config to "auto" mode. (118767)
SKUs Affected: DCS-7260CX3-64, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F and DCS-7260CX3-64E-R
- The OpenFlow agent fails to handle an OpenFlow group modify (OFPGC_MODIFY) message, resulting in an agent restart. Post agent restart, the OpenFlow controller is able to program the switch correctly. (142337)
- 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)
```

- Support for VRRPs in OpenConfig not fully implemented (268736)
- When a MLAG peer (configured as a service VTEP for MSS) reloads, MSS re-installs intercept flows which may cause traffic to drop for a short duration. (349254)
- If multiple front-panel ethernet interfaces are configured as reflector interfaces for the same flow of traffic to be reflected, loops can be formed and hosts can flap from one interface to another when the reflector configuration is subsequently changed. The workaround is to only configure one reflector interface to handle traffic reflection. (360869)
Series Affected: DCS-7280E, DCS-7280R, DCS-7280R2, DCS-7500E, DCS-7500R and DCS-7500R2 Series
- If a large number of rules are setup on a single 'ip access-list', the OpenConfig agent may use a lot of CPU time and may restart unexpectedly. (407002)
- When an OpenConfig NETCONF client is used to add new classes to policy map, the new classes are added before any previous classes. (451025)
- Any configured QoS mappings for traffic class to COS and traffic class to DSCP bits are not applied on packets that are reflected via an Interface Reflector that is configured in MAC swap and direction out mode. The source and destination MAC addresses are swapped as expected, but the COS and DSCP

bits of the reflected packets are copied unmodified from the original ingressing packet. (476507)

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

- EosSdk agents using a certain combination of Sysdb mount paths may cause the Sysdb agent to restart on connecting to Sysdb. If the agent is configured and stored in startup-config this can cause Sysdb to get stuck in a restart loop. (496371)
- Value for mlag leaf might not be mapped in certain circumstances. (518968)
- The "arp reply relay" feature does not support a VXLAN domain with L2 VTEPs. (540369)
- In certain circumstances, when a process takes large amount of memory, it may trigger a kernel panic. (556742)
- Adding or modifying Route-map via OpenConfig may not result in applying the new Route-map in hardware. (575325)
- When configuring interfaces in stateful DHCPv6 environments, ZeroTouch agent will select the longest matching prefix from the Neighbor Discovery router advertisement rather than assigning a default prefix-length of /64. Defaulting the prefix-length to /64 can lead to reachability problems resulting in the agent unable to communicate with the ZTP bootstrap-server. (645404)

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
- Groups of four adjacent ports share a common hardware resource. For example SFP ports Et1-4 share one, SFP ports Et5-8 share another, QSFP ports Et3/1-4 share a third etc.

During switch hitless reload or hitless restart of the forwarding agent, all of the ports which share a common hardware resource may flap if at least one of the ports has the speed-misconfigured errdisable cause, leading to a

brief traffic outage. The workaround is to remove the invalid speed configuration. (508444)

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

SKUs Affected: DCS-7050SX3-48YC12 and DCS-7050SX3-48YC12-F

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

Series Affected: DCS-7060X4 Series

- If the slice agent is restarted with more than the maximum logical ports configured, port active state will not be guaranteed after the agent restarts. Ports that were previously active may become inactive and vice-versa.

Workaround is to avoid configuring more than the maximum logical ports on a pipe. If the maximum logical ports must be exceeded, unexpected activeness changes can be avoided by preemptively restarting the slice agent using "platform trident reset". (682740)

SKUs Affected: DCS-7050SX3-48YC8

Layer 2

- If Ebra agent restarts on a VTEP that is running EVPN control plane with route type 5, then it can result in VxLAN traffic loss. (276063)
- When running rapid-PVST at high scales of interfaces and vlans, the device may occasionally not transmit a BPDU, which can lead to a BPDU timeout on the neighbor. (442663)

Layer 3

- During reload or configuration changes, BFD session may be stuck in init state on one side. Restart Bfd agent to work around the issue. (195545)
- 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)
- Removing LDP mappings in a large scale LDP configuration may result in LDP neighbor sessions flaps. (283972)

- Auto discovery of multiple IPv6 link-local BGP neighbors over a single interface does not work and will result in only one of the neighbors transitioning to established state. (289875)
- If a decap group with multiple IPv6 addresses are configured in either a config-session or config-replace, forwarding agent may restart unexpectedly, or the decap-group may not get programmed correctly. (421010)
- When layer 3 forwarding agent is restarted, it is possible that the agent does not reach warm state. The workaround is to restart the agent again, the second restart may cause traffic loss. (435193)
- When a series of back-to-back configuration changes are applied to a BFD session, the session may not have the latest configured value. To work around the issue, un-configure and re-configure the BFD session. (458955)
- When there're a lot of NAT translations synchronized through NAT Sync in a very short time period, NAT Sync may hang and stop synchronizing. Restarting Nat agent with "agent Nat terminate" in CLI Enable mode may help. Retrying the command several times may be needed. (467300)
- With 'ipv6 dhcp relay install routes' configured on more than 100 SVIs, DHCP relay agent may exit unexpectedly when a IPv6 DHCP helper address is configured simultaneously on all SVIs. (471338)
- In situations in which a large number of routes are forwarded over the same dynamic tunnel interface, the Arp agent may exit unexpectedly if the forwarding information of these routes are changed. The Arp agent will restart normally afterwards. (472837)
- When adding a VRF during config replace or quickly deleting and re-adding a VRF, the internal state of the ARP agent may go out of sync and ARP entries may be missing from "show arp." Workaround is to delete the VRF and re-add it. (474469)
- Cspf agent may restart when TI-LFA protection is configured and the number of segments requiring protection is over 2K (481703)
- TI-LFA backup paths will be computed slowly when there are large number of node-segments/anycast-segments/adjacency-segments that are protected. (489097)
- An LSR which is the penultimate hop in an RSVP LSP may incorrectly report a "bandwidth unavailable" error if the tunnel destination IP is an interface IP that does not match the last hop in the ERO.

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

- BFD sessions may continuously flap after a speed change on an interface. A workaround is to restart StrataBfd. (499185)
- Protocol messages exceeding 1500 bytes are ignored, in particular BUNDLE and SREFRESH messages, which is the case for jumbo frames. The symptom can be dropped LSPs or long convergence times. (521363)
- An RSVP RESV message is immediately sent as a reply to a NACK object, which can lead to a packet storm if an RSVP neighbor incorrectly uses NACK objects as an error reporting mechanism. Configuring RSVP with `refresh method explicit` will prevent this issue but disable Refresh Overhead Reduction. (522041)
- When any form of dynamic NAT is enabled, the first packet for translation is sent to CPU and forwarded in software and after the translation is programmed in the forwarding ASIC, the subsequent packets are forwarded by the ASIC. If the NAT destination is reachable via ECMP paths, the path chosen by the first packet by software forwarding may not be the same as subsequent packets forwarded in hardware. This can potentially cause issues if the NAT destination is to load-balancer or an anycast address that maintains NAT flow state is maintained per receive port.

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

Series Affected: DCS-7170 Series

- RSVP tunnels might not get established or may use non-best paths when IS-IS system ID or OSPF router ID is changed for any node in the IGP network and the best path goes via that node. One workaround is to change the TE router ID also for the same node. (546904)
- On scale setup during certain network events, ARP agent's internal state may go out of sync, leading to ARP agent repeatedly deleting and adding the same entry to the kernel. This may cause the CPU usage to spike to 100% usage.

Workaround is to restart ARP agent. (559855)

- SrTePolicy agent may restart if the top label of the Segment List is resolved over LFIB entry with outgoing label as Explicit Null Label. (561061)

- When TE router ID of one of the two routers which have the same TE router ID is changed to be same as that of a third router, Cspf agent may sometimes restart (565678)
- An IP configured on an interface in a non-default VRF may cause an RSVP LSP to not come up if the same IP appears in that LSP's path for a different router. (591762)

Multi-agent

- When a large number of updates are backed up in the transmit socket buffers, it may sometimes lead to the BGP neighbors session getting stuck in Idle(PendingNotification) state when the session reset is attempted and BGP is unable to send a notification. Not being able to send out updates may lead to traffic misforwarding or drops. (466575)
- Sometimes BGP (multi-agent) would advertise an incorrect originator-ID when sending updates to iBGP clients when configured as route-reflector. The problem only manifests when different paths (could be paths belonging to the same prefix or different prefixes) have the exact same path attributes, including NextHop. The problem corrects itself if the paths are resent to the route-reflector from the source whose paths had the incorrect originatorId. This problem can manifest with add-path configuration as well when the additional paths have the same path attributes. (472025)

Rib

- Arista BGP implementation does not advertise a best-path route back to its originator. When the best-path of a route changes, a BGP speaker might need to send out an "echo" withdrawal towards the new best path originator.

When a BGP speaker queues a large number of echo withdrawals for an update group, it might run into a job starvation situation that stops route advertisement while driving up CPU usage. The workaround is to clear the BGP session of the peer or peers that have echo withdrawals queued. (450701)

MPLS

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

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

- MPLS LFIB routes and tunnels may be slow to reconverge under heavy system load after a change in the network topology. This is a result of removing the route and tunnel when a downstream LSR is transiently lost. (678467)

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)
- Fast reload fails when multicast is configured on a non-default VRF that is not PIM enabled. The default VRF is always PIM enabled, but a non-default VRF is PIM enabled when it contains at least one PIM interface. Workaround is to enable a PIM interface on each non-default VRF. (398256)
- In an environment with a very large number of MSDP peers, during cleanup the MSDP Agent may unexpectedly restart. (418190)
- When using Igmp Snooping proxy in the network, DCS-7170 devices will drop the proxied IGMP reports which will result in multicast traffic not being bridged. (448375)
Series Affected: DCS-7170 Series
- When using Sfe based forwarding configuration, the BessMgr agent may restart unexpectedly when multicast packet size is greater than 1500 bytes (481192)
- If MLD Snooping is configured, IPv6 multicast packets may not be bridged as expected in a PIM SM or PIM SSM non DR VLAN. (493226)
Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- BessMgr agent could restart unexpectedly when resolving a large number of unresolved packets at once. (495863)

vEOS Router / CloudEOS

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

The workaround is to restart the software forwarding engine (Sfe) agent. (523066)
- Changing the mapping of Vrf to VNI (VXLAN Network Identifier) twice will cause the software forwarding engine (Sfe) agent to restart. There will be a brief moment of traffic loss which will recover eventually. (539064)
- When the VXLAN VNI mapping for a VRF is modified, paths towards one or more DPS peers may get updated incorrectly, resulting in loss of traffic towards such peers.

Removing and re-adding the VNI mapping may fix this problem. (549493)

DCS-7150 Series

- When Lanz Streaming is moved to non-default VRF, in some cases, it may not accept the client connections in new VRF and if Streaming trying is trying to connect continuously, Lanz Agent may hog the CPU and usage might go very high. The workaround is to disable and reenale Lanz streaming. (338825)

DCS-7170 Series

- On the DCS-7170 series MLAG protocol traffic shares the same CPU queue with traffic destined to CPU. A burst of such traffic, destined to CPU, can lead to MLAG flaps. It will recover automatically when the burst is over. (449993)
- On the DCS-7170 series L3 protocol traffic shares the same CPU queue with traffic destined to CPU. A burst of such traffic, destined to CPU, can lead to protocol flaps. It will recover automatically when the burst is over. (464981)
- On DCS-7170 series, a policer configuration with burst unit in mbytes will result in continuous crash of BfnQos agent. Remove the configuration to recover, and the issue can be worked around by specifying the burst in kbytes. (469790)
- IPv4 Multicast routes may be prematurely aged out (472845)
- Hardware errors in the switching ASIC may get silently ignored after an initial reporting. (484277)
- Monitor session with TX source to a port-channel may not be rate limited correctly. Link flap of the port channel may cause the mirrored traffic to stop. Workaround is to remove non-working port from port-channel, or remove egress mirror rate limit. (516828)
- The issue is found if there is no space left to add new daisy chain entries but a change in the nexthop-group causes a need for a new entry. The BfnLoadBalancer agent can terminate if the server belonging to the entry that was supposed to be added to the daisy chain table goes down.

If the issue is not resolved after agent restart it can be solved by toggling:

```
load-balance server
  connection history source switch
```

```
To:
load-balance server
  connection history source server
```

and back:
load-balance server
connection history source switch (527242)

DCS-7170 Series

- The 7170 load balancer solution supports 512 L2 adjacencies. The usual topology usually uses 1 or 2 ecmp links to the next-hop router. If by mistake 512 L2 adjacencies are programmed and then the user tries to remove the incorrect configuration the adjacencies are not being removed. The workaround is to restart the BfnL3 agent. (515679)
- When a monitor destination is continually congested due to fan-in/high rates of traffic, and LAG is used in the system, the slice agent may restart when there is LAG membership change or port status change. (516260)

CCS-720XP Series

- Groups of four adjacent 2.5GBASE-T ports share a common hardware resource; for example, Et1-Et4 and Et5-Et8 and so on. Due to an issue with the shared hardware resource, all 4 ports may blackhole traffic even though they report linkup.

Workaround is to run `shutdown` followed by `no shutdown` on the affected 4 ports. (381587)

SKUs Affected: CCS-720XP-48ZC2

- When two or more hardware flow trackers are configured the forwarding agent may restart unexpectedly. In this case even if the extra flow tracker is removed the forwarding agent may continue to restart.

Workaround is as follows:

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

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

(config)# no flow tracking hardware
and reconfigure flow tracking hardware and add only one flow tracker. (426694)

- 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. (436978)

- When a device behind phone/hub/switch connected to Arista switch is authenticated via Dot1x MAC Based Authentication, it cannot be moved to different port if dynamic VLAN is not received or dynamic VLAN is same as native VLAN. (477280)

Series Affected: CCS-720XP Series

- Authentication of any data device behind a tagged phone will fail after the phone is authenticated and authorised via EAP in phone VLAN on that port. Also, if a data device authenticates before tagged phone doing EAP, the authorisation of phone in phone VLAN will fail. (490365)
- NAS sends two EAPOL identity requests in response to EAPOL start from supplicants. This can cause authentication timeouts for some Windows supplicants. (498610)
- Executing ASU with phone trunk ports configuration can lead to Dot1x agent being killed. (503859)

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

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

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

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

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

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

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

- Groups of four adjacent 2.5GBASE-T ports share a common hardware resource; for example, Et1-Et4 and Et5-Et8 and so on. Due to an issue with the shared hardware resource, all 4 ports may blackhole traffic even though they report linkup.

Workaround is to run `shutdown` followed by `no shutdown` on the affected 4 ports. (545508)

SKUs Affected: CCS-720XP-24ZY4

- While doing ASU, Dot1x Agent may restart unexpectedly if there were 802.1X supplicants present before ASU. (558476)
- Front panel BASE-T ports might stop forwarding traffic even though they are reported as linkup.

Workaround is to run `shutdown` followed by `no shutdown` on the affected interfaces. (572686)

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

- Some link partners might receive bad traffic at 100M rate on the BASE-T ports. (631123)

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

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

- When an acl configuration is applied to a large number of SVIs in a configuration session, the acl configuration fails. (367513)
- Under certain conditions where small packet are bursting across many interfaces, chip can get in a state of dropping half packets coming into the chip. (488904)

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

- ACL/PBR/QOS will not take effect on ports with VLAN Translation configured. (246982)
- When a bidirectional PIM multicast receiver is connected to only one MLAG peer switch, the IGMP reports from the receiver may get dropped resulting in multicast traffic drops towards the receiver. (259308)
- 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)
- After applying or removing shaping configuration on an L2 subinterface, the L2 subinterface must be flapped for the configuration to take effect. (278730)
- Power cycling a 7500E linecard in a chassis with 7500E fabric modules can cause packet loss on 7500R linecards. (295825)

- In the presence of virtual MAC with mask configuration, packets that undergo recirculation may not be properly forwarded. The workaround is to eliminate the configuration that causes the recirculation to happen. (387535)
- Traffic subject to CPU traffic policy `_policer_` action will not be `_forwarded_`. This includes multicast control traffic which is replicated to both the CPU but also forwarded (e.g. PIM, OSPF HELLOs, etc.). Unroutable unicast traffic is also flooded, and so will not be forwarded if matched by a CPU traffic policy policer action. An example policy of concern is

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

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

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

- Insertion of a fabric module may cause brief traffic loss on the device (402508)
SKUs Affected: 7504R3-FM and 7508R3-FM
- With scaled Acl configurations, SandAcl agent restart may result in additional agent restarts. (408260)
- Configuring 4094 different Egress IPv4 RACL on 4094 SVI will cause repeated SandAcl Agent crash.
No workaround is available except for scaling down Egress IPv4 RACL. (415577)
- When the outgoing interface of a VXLAN underlay route changes from one forwarding chip to another, it can negatively impact VXLAN overlay route convergence. (440233)
- Performing many updates to a PBR policy in a short amount of time may cause the SandAcl agent to restart unexpectedly. (441912)

- With VXLAN configuration, if the route to a remote VTEP churns while the tunnel hardware resource usage is closer to the hardware limits, then the SandL3Unicast agent might restart unexpectedly. (445834)
- V6 unknown multicast packets are not headend replicated to remote VTEPs. (448950)
- 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)
- On DCS-When a port-based connector is used for a pseudowire or local cross-connect patch, dot1q-tagged packets will have the outer tag incorrectly removed on ingress/encap if any one of several other features are configured on the switch (on any interface) including: L2 subinterfaces, Dot1q Tunnel, Vlan Mapping, and Vxlan. The only workaround is to unconfigure these other features. (474093)

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

- If the number of BGP Flowspec matches exceed the hardware TCAM banks and entries available, and then Flowspec is enabled on additional interfaces while in the overflow state, the matches may fail to be programmed on the additional interfaces, even if the published Flowspec matches later fall back under the hardware limits. (477150)
- The SandAcl agent may crash if PBR is applied on SVIs during a linecard hotswap (486402)
- When BGP neighbors advertise Flowspec rules containing the police action and the system's policer resources are exhausted, any further change to Flowspec rules may cause the SandAcl agent to restart unexpectedly. Upon a SandAcl agent restart, existing Flowspec rules are removed from hardware and the system will attempt to program the new set of advertised rules, which will succeed if there are enough hardware resources.

As a workaround, restart the SandAcl agent after freeing enough policer resources, e.g., remove rules from features such as QOS policies that use policer resources. (486688)

- In a 6PE configuration, packet ingressing on an L3 subinterface with explicit null label and destined to this switch is not supported. (504582)
- Removing a field-set configuration (i.e., via "no field-set ipv4 prefix <field-set-name>") does not cause traffic-policies that match on the field-set to be updated in the switch. Packets will still match on the IP prefixes or L4-

ports of the removed field-set configuration. The removal of the field-set from hardware will take effect in a subsequent change to the policy configuration. (512073)

- TCAM entries may become corrupted when two CLI commands which read TCAM entries from hardware are executed simultaneously. The commands which cause a TCAM hardware read are 'show platform fap acl tcam diff', and 'show platform fap acl tcam hw hits'. (521356)
- SandFap agent restart or a "config replace" operation may cause some subinterfaces to transition to "notconnect" state.

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

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

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

- While in Tap Aggregation mode, changing the composition of a tool group may trigger a cascade of events which could lead to the group not able to replicate and transmit packets.

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

- A config-session that applies primary policy to an interface and updates ACLs/classes used by the primary policy may lead to the policy to not be programmed correctly in hardware if the interface(s) where this policy is applied is/are also covered by a fallback PBR policy. (578146)

DCS-7500E Series

- SAND_POSSIBLE_STUCK_PORT may be logged on an Ethernet interface after switch initialization, which can cause an interface to not transmit packets despite being up.

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

DCS-7280R and DCS-7500R Series

- The SandFap-Linecard agent may restart after logging "TX (or RX) DMA unhealthy". (138960)

- 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

- BFD may briefly drop packets when the standby supervisor comes back online. (389057)

SKUs Affected: DCS-7504, DCS-7504N, DCS-7508, DCS-7508N, DCS-7512N and DCS-7516N

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

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

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

- Configuring more than 6 MPLS labels for the pseudowire path (including VC label and flow label), will cause the SandTunnel Agent to crash and not recover. The only workaround is to reduce the number of transport labels and/or to remove flow label such that the total is 6 or fewer. (443140)
- Optimized routes may not be programmed properly if TCAM profile is changed from <profile-1> to <profile-2> with both profiles having 'feature flex-route' in it.

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

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

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

- A statically configured MAC entry may be replaced by a dynamically learned MAC entry leading to transient mis forwarding of packets destined to that MAC address. This transient condition gets corrected in the next poller cycle. (428580)
- SandAegis agent may crash when the last IPv6 rule is deleted from a traffic-policy which has both IPv4 and IPv6 rules (435733)
- VXLAN IPv6 overlay configuration along with VARP and Virtual VTEP configuration is not supported. (442914)
- A. When Priority Flow Control (PFC) is enabled, incoming congestion on first port of a FAP with a particular priority may cause control traffic of that priority to be dropped on that FAP.
B. Conversely, if there is a lot of control traffic egressing on a FAP, it may cause PFC frames to be transmitted out of the first port on that FAP. The workaround would be to shutdown the first port on the FAP which is experiencing this issue. (443075)
- On Switch with Jericho2 FAP(s), if Pimsm spt threshold is configured as infinity, Pimsm multicast traffic forwarding may not work properly. (447366)
- The SandFapNi agent may restart continuously preventing ports coming up. (462680)
SKUs Affected: DCS-7280CR3-96, DCS-7280CR3K-96, DCS-7280DR3-24, DCS-7280PR3-24 and DCS-7280PR3K-24
- When tc-to-txQueue map is changed, the traffic enqueued into traffic class that had ECN config applied earlier might still be subject to ECN marking. The workaround is to remove and reapply the previous ECN config on the tx-queue. (472217)
- Forwarding of IPv4 and IPv6 multicast packets over MPLS Pseudowires is not supported. (493142)
- Configuring 513 different Egress IPv6 SubInterface ACLs will cause repeated SandAcl Agent restarts. No workaround is available except for scaling down Egress IPv6 SubInterface ACLs. (502185)
- In interface traffic-policies, when configuring a destination port after configuring source port field-set, the ports are treated as field-set names (506962)
- When two or more of the following features: Port ACL, PBR, QoS, Flowspec, Aegis, and CPU Policy are applied on the same interface with a non-default TCAM profile, one or more of the features may not work (526913)
- If the "acl subintf ip/ipv6/mac" feature is added to a TCAM profile without the corresponding "acl port ip/ipv6/mac" feature being added to the same profile, subinterface ACLs will not work. (532393)

- MTU violating VXLAN routed packets with either IPv6 overlay or an Ethernet tag on the underlay ethernet header can cause the egress port to lock up. When in this state, no packets will egress the port and traffic will get deleted on the ingress VOQ. Shutting and unshutting the port will restore the traffic. (537323)
- A VXLAN packet that would be bridged, after decapsulation, to the same interface as it ingressed will be dropped instead of being bridged. (549785)
- MPLS tunnel terminated GRE-over-IPv6-over-MPLS packets are forwarded incorrectly. (550842)
- When VXLAN is enabled, L4 port match in ACLs on VXLAN decap packets does not work. (567246)
- Upon recovery from a PFC pause storm, Normal transmission of PFC frames on the interface and traffic class where we received the pause storm will not resume.

Workaround: To restore PFC functionality, disable and re-enable PFC on that interface. (571168)

DCS-7020 Series

- When traffic is flowing from a 10G interface to a 1G interface and PFC is enabled on the ingress 10G interface, if the administrator changes the COS to TC mapping all the incoming traffic on the 10G interface is dropped. The only way to recover is by restarting SandFap. (400653)

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

- Configuring multiple DirectFlow flows of different priorities at once may restart the SandAcl agent continuously. A workaround would be to configure the DirectFlow flows one after the other. (543622)
- The flexroute feature may override the forwarding decisions from PBR policies and BGP Flowspec matches. (544872)
- SandAcl agent may restart when programming flows at many different priorities. Workaround is to limit the number of different priorities to 1024. (588420)
- In an EVPN MPLS multihomed configuration, packets egressing towards non designated forwarder ports may not be dropped. (592102)

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

- Removing dot1q-tunnel configuration from a port will allow traffic from any VLAN to be received on that port and to be forwarded in the default VLAN

context .

Workaround: After removing dot1q-tunnel configuration, configure the port to be a port connector in a local cross-connect patch, then remove the port from the patch. This will restore to proper ingress VLAN membership filtering behavior. (355471)

- Following a SandFapNi restart or a switch reload, it has been noticed that traffic-policies may not be installed on interfaces. This can happen especially if the traffic-policies are big or if there are too many traffic-policies to be installed.
The workaround is to uninstall all the traffic-policies, restart the SandAegis agent and reinstall the traffic-policies on the interfaces. (607607)

DCS-7280R and DCS-7280R2 Series

- With AlgoMatch enabled, the tc-counters TCAM profile applied, an IPv6 security ACL applied on an interface, and an IPv6 PBR policy applied on the same interface as the ACL, actions for that ACL (deny, count, etc) will not apply to incoming traffic. (399725)
- Traffic egressing out of a LAG interface may not be subject to egress ACL on reload (473227)

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

- In multi-chip systems with L2 MTU greater than 9236, packets that are larger than 9236 bytes that flow between the switch ASICs will be dropped. (102837)
Series Affected: 7300X3, DCS-7250X, DCS-7260CX and DCS-7300X Series
- On "reload hitless", some port channels may flap unexpectedly. (151197)
- When the utilization of the MAC address table grows close to the full hardware capacity, the StrataL2 agent may encounter an out of memory (OOM) condition in the system, causing the agent to get restarted. (238115)
Series Affected: 7320X, DCS-7010, DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7260X, DCS-7260X3 and DCS-7300X Series
- 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

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

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

- IS-IS configured over LAGs may see neighbor adjacency flapping during a hitless restart. (276232)
- Applying a large configuration containing many interface changes may result in Strata forwarding agent restart, which could cause traffic loss during the reconfiguration. (330840)
- If the slice agent is restarted with more than the maximum logical ports configured, port active state will not be guaranteed after the agent restarts. Ports that were previously active may become inactive and vice-versa.

Workaround is to avoid configuring more than the maximum logical ports on a pipe. If the maximum logical ports must be exceeded, unexpected activeness changes can be avoided by preemptively restarting the slice agent using "platform trident reset". (356578)

Series Affected: DCS-7050SX3 and DCS-7260CX3 Series

- When peer mac routing is enabled, ARP replies destined to MLAG peers may be dropped or the Strata agent may restart when peer link is a physical interface.

The work around is to use a port channel instead of physical interface for the MLAG peer link. (365533)

Series Affected: 7300X3 Series

- 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)
- DHCP unicast packets which are not destined to switch are copied to CPU, which may forward the packet to destination along with one copy forwarded in hardware. This results in duplication of unicast DHCP packets. (393520)

- 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)

- On ECMP shrinkage due to link down events, VXLAN tunnels reachable via ECMP may experience extended outage until routes reconverge. (419753)
- 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)

- During Leaf-SSU or SSO, interfaces with 100GBASE-SRBD transceivers may experience a single link flap. (440451)
- A speed change on an interface may result in interface wrongly pointing to a stale nexthop causing traffic loss on that interface. The workaround is to restart the L3 agent. (444936)

Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX3 and DCS-7260CX3 Series

- Changing Acl configuration rapidly and multiple times may cause the forwarding agent to repeatedly restart unexpectedly. (447584)
- Special filter is installed on MLAG port channel member interfaces to block packets ingressed from MLAG peer via peer link at egress to avoid loop. When the filter is not installed on a member interface before it is added to MLAG port channel, there might be a momentary loop when interfaces belonging to an MLAG port channel link up. (471268)

- Slice agent might restart repeatedly in scale QoS policy-map configuration with policers, which causes hardware policers to exhaust. The workaround is to reduce number of policy-maps with policers attached to them. (489846)

- Programming or updating a Router ACL may result in forwarding plane slice agent restart and can leave the corresponding vlan permanently blocked. (498502)

- During an SSO failover, interfaces with VXLAN configured may result in an increased traffic outage on the associated interface (528749)
Series Affected: DCS-7300X Series
- Groups of four adjacent BASE-T ports share a common hardware resource; for example, Et1-Et4 and Et5-Et8 and so on. Due to an issue with the shared hardware resource, all 4 ports may fail to link up.

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

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

- Attaching a policer to class-default of an egress policy-map consisting of both IPV4 and IPV6 ACLs, is not supported and may lead to forwarding agent restart (534152)
- When an MLAG peer is reloaded with VXLAN dynamic VLAN to VNI configuration, some VNI entries may not get programmed in hardware resulting in traffic outage.

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

- 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

7368 and DCS-7060X4 Series

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

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

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

Series Affected: DCS-7060X4 Series

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

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

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

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)

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

- Egress ACL on SVI for BUM traffic will not match ACL rules for packets being sent across fabric interfaces. (284543)
- When performing SSU, there may be failure in table programming resulting in fatal error cold reboot with extended traffic outage. (399774)
- With ISIS configured on multiple VRFs, traffic loss may occur during SSU. The ISSU sessions on non-default VRFs may flap due to lost hello packets. There is no workaround for this issue. (414464)
- When "platform trident forwarding-table flexible" is configured, L3 unicast forwarding to hosts / L3 multicast forwarding may not work resulting in unexpected behavior. The workaround is to remove the CLI and use "platform trident forwarding-table partition" instead. (421084)
- Performing SSU may result in extended traffic outage, leading to total traffic loss being slightly over 1 second. (477139)
- Egress ACL rules can not be used to match ip protocol 50. (535104)
- VXLAN encapsulated packets with inner destination MAC as the switch MAC may get dropped when received on a L2 only VXLAN VLAN. (582895)

DCS-7160 Series

- In some cases, when PFC (Priority Flow Control) is configured on an interface and if the user performs a "default interface" on a range of interfaces, the XpQos agent is restarted. (245559)

Limitations and Restrictions in 4.23.16M

BGP EVPN L2/L3 VXLAN/MPLS

- On a BGP router, EVPN creates dynamic VLANs for L3 VNIs in receiving EVPN RT2 or RT5 advertisements even when the router does not have any L3 VRF or import route-target configurations. (294328)

CVX

- The Arista OpenStack plugin has been enhanced to create a default route for the routers created in OpenStack, when VRF creation and `vrf_default_route` option is enabled in the `networking-arista` plugin configuration (295274)

General

- 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)
- Mirroring to an MLAG port channel is not supported. (276973)
- EOS support alphanumeric VLAN names, however ODL gives an error when reading the leaf `/interfaces/interface/routed-vlan/config/vlan` if the name contains alphabetic characters.

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

- Enabling next-hop sharing (using "ip hardware fib next-hop sharing" command) could result in slower convergence in the case of link failures. (369305)
- For static tunnel interfaces, duplicate or conflicting GRE tunnels are tunnels that have the same source, destination, tunnel mode (GRE) and key configurations as an existing tunnel. A duplicate tunnel is always in the down state. When there are duplicate static tunnel interfaces and a config-replace is executed, it is non-deterministic which one of the conflicting tunnels will be in the up state. (373135)

- When 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

- A range of 0.01 sec to 255 seconds is allowed for the YANG VRRP advertisement-interval value (range is 1..25500 as advertisement-interval is specified in hundredths of a second), but EOS only accepts values in the range 1 sec to 255 secs. (426391)
- Users are unable to login to the standby supervisor via LDAP if an SSL profile is used (442915)
- Mirror-On-Drop does not monitor IP packets with an invalid IP version. (443839)
- The YANG path '/network-instances/network-instance/mpls/global/config/null-label' is unsupported, but it is incorrectly listed as supported, and accepted by the OpenConfig agent. (514041)
- The following YANG paths are unsupported, but are incorrectly listed as supported and accepted by the OpenConfig agent.

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

- Sampled flow tracking doesn't export proper bridging information (ingress VLAN, egress interface, egress VLAN, etc.) in IPFIX flow records for flows on layer 2 subinterfaces when the bridging VLAN is different than the dot1q VLAN.
The workaround is to use the same dot1q VLAN and bridging VLAN for these layer 2 subinterfaces. (572688)
- Sampled flow tracking doesn't export proper routing information (VRF, next-hop, egress interface, BGP information etc.) in IPFIX flow records for flows on sub-interface with QinQ configuration. (575271)

Layer 1

- 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

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)
- 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)
- The command "aaa accounting dot1x default start-stop group radius multicast", does not yet enable this new feature. (455169)

Layer 3

- 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)
- Tunnel interfaces flap when configuring ttl, tos and path mtu discovery on them. (281464)
- IPv6 addresses are not supported as tunnel source and destination. (283912)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289217)

- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289224)
- The 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 MPLS tunnel egress counters or next-hop group counters are enabled, forwarding plane agent restarts could see traffic loss (435899)
- IPv6 Egress ACLs may not be applied to traffic if 'ip hardware fib next-hop sharing' is also configured (446000)
- PBR policy maps which match nexthop-groups will not be applied on platforms that have hierarchical FEC support, but have disabled it with the "ip hardware fib hierarchical next-hop disabled" or "ip hardware fib hierarchical next-hop max-level 1" commands. (471880)

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)
- Routes which are imported into a target VRF using a "leak routes source-vrf" policy cannot be matched on the basis of the source-protocol from the source-VRF as part of the "rib fib policy" (403739)

Rib

- Match dynamic prefix-list clauses in route maps are only supported for BGP inbound and outbound route maps. (261865)
- Dynamic prefix-list match clauses are not supported in BGP neighbor inbound route maps. (269663)
- The BGP configuration commands, "bgp next-hop-unchanged" and "neighbor next-hop-unchanged", are not supported in ribd mode (single-agent mode).

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

MLAG

- The port-channel configured on each MLAG switch with the same MLAG ID must also have the same channel-group ID. (22890)

MPLS

- BGP LU Tx MPLS routes corresponding to BGP LU advertisements with multiple labels are not supported. (395029)
- MPLS vias in the LFIB are limited to 128 even when the maximum allowed ECMP or UCMP values exceed this number (602879)

Multicast

- 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)
- 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)

vEOS Router / CloudEOS

- 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)
- vEOS supports up a maximum of 8 physical cores. With hyperthreading enabled it is 16 vCPUs (278286)
- In dpdk based vEOS(MODE=sfe in /mnt/flash/veos-config), SR-IOV mode doesn't work with Intel x520/82599 on ESXi. As a workaround use mixed mode instead of the SRIOV only mode in ESXi (296718)
- Hot detaching network interface instances is unsupported. Reboot the instance after detaching a network interface. (306132)

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
- 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
- Full mroute scale may not be achievable under certain scenarios due to hash collision for the multicast route table. (380545)
- VXLAN encapsulated ieee reserved mac addresses are not trapped to CPU and hence
any control protocol that is VXLAN encapsulated will not work. (423963)

CCS-720XP Series

- "show environment power" does not show insertion status or telemetry information for PWR-1021-AC-RED or PWR-621-AC-RED when there is no input power. (369368)
SKUs Affected: PWR-1021-AC-RED and PWR-621-AC-RED
- When hardware flow tracking is enabled, tcpControlBits information element in IPFIX flow records for VXLAN encapsulated TCP flows is incorrect. (416079)

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

- show interfaces counters rates' for port channels may show a slightly inaccurate data rate (including above 100%). (67367)
- 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)
- 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)
- When ingress packet truncation is enabled in Tap Aggregation mode, truncated packets are not counted as packets received on the ingress interface. (215346)

- In Tap Aggregation mode on interfaces with MPLS label stripping enabled, MPLS traffic may be throttled aggressively.

The workaround is to create a custom PDP policy and set its action to "count". (223248)

- 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)
- LLDP functionality is not supported on interfaces with ingress truncation enabled. (229983)
- 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 line rate. (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)
- 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)
- VXLAN routing is not supported on a modular system on which one of the linecards is configured for TAP aggregation. (306180)
- Stale telemetry flow-spec statistics may be seen until the next "show flow-spec" command is issued. (345195)
- unicast RPF is not enforced on IP packets when the destination IP needs ARP resolution (356347)
- 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)
- Drop Precedence based tail drop thresholds can be ignored when multiple queues are congested and buffering resources are low. (395693)
- If a security-ACL is configured on a sub-interface beyond the supported limit of ACLs, SandAcl might restart continuously until the number of configured ACLs is brought back within the supported limits (398815)
- Queue build-up doesn't show up for 64 byte packets in case of congestion via following commands:-
1) platform fap diag diag cosq non_empty_queues

2) show interfaces queue length

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

- 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)
- SflowAccel supports up to 200 VRFs. (423892)
- Header truncation for mirrored packets can not be used when the destination of a monitor session is a GRE tunnel. (428547)
- IPv4/IPv6 counters are not supported on port-channel interfaces. (479658)
- Egress MPLS tunnel counters are not supported for implicit null tunnels. (496375)
- With 'hardware next-hop fast-failover' configuration, the maximum ECMP size supported is reduced to 1023 in R, R2 series and 127 in R3 series switches respectively. (502940)
- "shape rate" configuration is not supported on LAG subinterfaces used for MPLS pseudowires. (555799)

DCS-7280R and DCS-7500R Series

- 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)
- Drop-precedence thresholds are not supported on mirror session ports. (388932)
- We do not support incoming multicast LDP packets with a destination multicast MAC address.

RFC 5332 doesn't mandate the use of sending labeled multicast traffic in a multicast ethernet frame. Multicast MAC destination address could be used in upstream label assignment. (393406)

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

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

- With "match mpls" PBR configuration, MPLS packets matched both by PBR and a LFIB route with a pop action would get forwarded correctly only if the MPLS route contains "access-list bypass" configuration. (370118)
- IP egress ACLs are not supported for MPLS terminated packets. (383263)
- Strict Priority scheduling for 400G interfaces don't work properly for 64 byte packets when traffic rate is above PPS limit of the chip. (388517)
- Cli to change tail-drop threshold for 200/400G ports is not supported by software. (396029)
- Applying a Qos policy on Vlan interfaces may result in SandAcl restart, since this functionality isn't supported. (403280)
- Applying a Qos policy on sub-interfaces may result in SandAcl restart, since this functionality isn't supported. (403281)
- "ip tunnel <ip> nexthop-group" configuration for nexthop-groups with type mpls consisting of direct IP nexthops without label encapsulation is not supported. (458065)

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

- Only port-channel interfaces are supported as peer links. Single member port-channel interfaces are supported. (409958)

DCS-7280R and DCS-7280R2 Series

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

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

- Clause 37 auto-negotiation on SFP28 and QSFP28 ports is not supported. A workaround is to use 1G forced speed configuration on the affected ports. (206778)
Series Affected: 7300X3, CCS-720XP, DCS-7050CX3, DCS-7050SX3, DCS-7050X3, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7260CX3 Series
- 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)
- Flexible port-channel hashing may not hash accurately when the inner L4 header option is configured for IPV6 GRE packets. (267487)
- BGP neighbor sessions may flap when URPF is configured on an interface. (279113)
- Issuing "no shutdown" on a member link of port-channel may send duplicate packets for sub-second on that member link (291514)
- DirectFlow (and features such as the MacroSegmentation Service that use DirectFlow) are unsupported on the 7260, 7300 series of switches. (296715)
- When primary vlan and secondary vlan are part of different MST instances, the hardware programming of lag will not happen. (300119)
- Not all interfaces in the groups Ethernet49-Ethernet52 and Ethernet53-Ethernet56 can be broken out in to 4x25G, 4x10G or 2x50G mode due to MAC resource limitation on the system. Interfaces without MAC resource will be marked inactive and 'STRATA-4-HW_PORT_RESOURCE_FULL' will be logged in syslog.

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

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

- When MLAG peer MAC routing is enabled, OSPF neighborship over VXLAN overlay may never get established. (349242)
- SSU is not supported with Vxlan and MLAG config. (388669)
- Ports with ingress MAC ACL still learn MAC addresses although traffic may be dropped (400211)
- When both "hardware counter feature vni decap" and "hardware counter feature vni encap" are enabled, the egress flexcounter pools will only be released after both features are disabled. (441887)
Series Affected: 7300X3, 7320X, 7368, CCS-720XP, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7250X, DCS-7260X, DCS-7260X3 and DCS-7300X Series
- With VXLAN dataplane learning disabled, locally learned MAC addresses will not age out if VXLAN encapsulated packets from remote VTEPs are received with inner packet's source mac as the locally learned MAC. The locally learned MAC entry needs to be cleared manually by issuing 'clear mac address-table dynamic vlan <VLAN> address <ADDRESS>'. (460607)

7368 and DCS-7060X4 Series

- show running-config incorrectly displays interfaces for non-inserted cards. Interface statuses for missing or powered down cards are displayed as inactive. (306105)
Series Affected: 7368 Series
- Grpc buffer-usage counters are not supported on Tomahawk3 platforms, and can cause forwarding agent crashes if configured (593626)

Transceiver Series

- 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)

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

- Leaf stateful switch upgrade is not supported on this platform. (415340)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- MPLS packet undergoing swap operation does not honor the MTU configured on ingress/egress interface. (422668)
- Only 2K MPLS routes are supported (425691)
Series Affected: CCS-720XP Series

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

DCS-7160 Series

- If there is a port ACL (PACL) or router ACL (RACL) that matches on L4 information for TCP/UDP/ICMP, and if the incoming packet has IP options, then those packets are dropped. (172932)

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.