Arista Networks EOS 4.22.0F Release Notes

Version: 1.0

13 June 2019

Table of Contents

Resolved Caveats in 4.22.0F

Known Software Caveats in 4.22.0F

DCS-7020RA, DCS-7160, DCS-7170, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards Series

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

DCS-7280R and DCS-7500R Series

DCS-7020 Series

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

DCS-7010T Series

DCS-7060X and DCS-7300 Series

DCS-7260X3 Series

7368 Series

DCS-7050X and DCS-7300 Series

DCS-7050X2 and DCS-7300 Series

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

DCS-7160 Series

Transceiver Series

vEOS Router Series

Conditions and Impacts

EOS 4.22.0F Release Highlights

New Platforms and Hardware

- 7368
- 7368-16C
- 7368-SUP
- 7368-SUP-D
- 7368X4-SC
- DCS-7020SR-32C2-R
- DCS-7020SRG-24C2-F
- DCS-7020SRG-24C2-R
- DCS-7150SC-24-CLD-F
- DCS-7150SC-24-CLD-R
- DCS-7150SC-64-CLD-F
- DCS-7150SC-64-CLD-R
- DCS-7260CX3-64E-F
- DCS-7260CX3-64E-R
- PWR-400AC-RED

SNMP MIBs

- SNMPv2, SNMPv3
- RFC 3635 EtherLike-MIB
    - obsoletes RFCs 1650, 2358, 2665
- RFC 3418 SNMPv2-MIB
    - obsoletes RFCs 1450, 1907
- RFC 2863 IF-MIB
    - obsoletes RFCs 1229, 1573, 2233
- RFC 2864 IF-INVERTED-STACK-MIB
- RFC 4292 IP-FORWARD-MIB
    - obsoletes RFC 1354
- ARISTA-SW-IP-FORWARD-MIB
    - IPv4 only
- RFC 4363 Q-BRIDGE-MIB
- RFC 4188 BRIDGE-MIB
- ARISTA-BRIDGE-EXT-MIB
- RFC 4113 UDP-MIB
    - obsoletes RFC 1213
- RFC 4022 TCP-MIB
    - obsoletes RFC 1213
- RFC 4293 IP-MIB
    - obsoletes RFC 1213
- HOST-RESOURCES-MIB
- LLDP-MIB

- LLDP-EXT-DOT1-MIB
- LLDP-EXT-DOT3-MIB
- ENTITY-MIB
- ENTITY-SENSOR-MIB
- ENTITY-STATE-MIB
- RMON-MIB
    - rmonEtherStatsGroup
- RMON2-MIB
    - rmon1EthernetEnhancementGroup
- HC-RMON-MIB
    - etherStatsHighCapacityGroup
- RFC 3636 MAU-MIB
    - ifMauDefaultType and ifMauAutoNegStatus are writeable
- SNMP-TLS-TM-MIB
    - RFC 6353
- SNMP-TSM-MIB
    - RFC 5591
- ARISTA-ACL-MIB
- ARISTA-SNMP-TRANSPORTS-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-SMI-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-PRODUCTS-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-QUEUE-MIB
    - Excluding 7150
- RFC 4273 BGP4-MIB
- RFC 4750 OSPF-MIB
- ARISTA-CONFIG-COPY-MIB
- ARISTA-CONFIG-MAN-MIB
- ARISTA-REDUNDANCY-MIB
    - 7500R, 7500E, 7300X, 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- MSDP-MIB
- PIM-MIB
- IGMP-MIB
- IPMROUTE-STD-MIB
- VRRPV2-MIB
- ARISTA-QOS-MIB
- ARISTA-ENTITY-SENSOR-MIB
- ARISTA-BGP4V2-MIB
- ARISTA-VRF-MIB
- ARISTA-DAEMON-MIB
- ARISTA-IF-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-MAU-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R

- ARISTA-PFC-MIB
    - 7050X2, 7060X, 7260X,7320X, 7060X2, 7160,7020R
- ARISTA-FIB-STATS-MIB
- ARISTA-GENERAL-MIB
- ARISTA-HARDWARE-UTILIZATION-MIB
- ARISTA-NEXTHOP-GROUP-MIB
- ARISTA-SW-IP-FORWARDING-MIB
- ARISTA-TEST-MIB
- ARISTA-XCVR-DWDM-MIB
- ARISTA-XGS-MIB
- IEEE8021-PFC-MIB
- IEEE8023-LAG-MIB
- MPLS-LDP-STD-MIB
- NOTIFICATION-LOG-MIB
- OSPF-TRAP-MIB
- OSPFV3-MIB
- PTPBASE-MIB

SNMP Traps

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

Supported Hardware

Fixed System

- DCS-7010T-48-DC-F
- DCS-7010T-48-DC-R
- DCS-7010T-48-F
- DCS-7010T-48-R
- DCS-7020SR-24C2-F
- DCS-7020SR-24C2-R
- DCS-7020SR-32C2-F
- DCS-7020SR-32C2-R
- DCS-7020SRG-24C2-F
- DCS-7020SRG-24C2-R
- DCS-7020TR-48-F
- DCS-7020TR-48-R
- DCS-7020TRA-48-F
- DCS-7020TRA-48-R
- DCS-7050CX3-32S-F
- DCS-7050CX3-32S-R
- DCS-7050QX-32-F
- DCS-7050QX-32-R
- DCS-7050QX-32-SSD-F
- DCS-7050QX-32-SSD-R
- DCS-7050QX-32S-F
- DCS-7050QX-32S-R
- DCS-7050QX-32S-SSD-F
- DCS-7050QX-32S-SSD-R
- DCS-7050QX2-32S-F
- DCS-7050QX2-32S-R
- DCS-7050SX-128-F
- DCS-7050SX-128-R
- DCS-7050SX-128-SSD-F
- DCS-7050SX-128-SSD-R
- DCS-7050SX-64-F
- DCS-7050SX-64-R
- DCS-7050SX-64-SSD-F
- DCS-7050SX-64-SSD-R
- DCS-7050SX-72-F
- DCS-7050SX-72-R
- DCS-7050SX-72-SSD-F
- DCS-7050SX-72-SSD-R
- DCS-7050SX-72Q-F
- DCS-7050SX-72Q-R
- DCS-7050SX-96-F
- DCS-7050SX-96-R
- DCS-7060CX-32S-F
- DCS-7060CX-32S-R
- DCS-7060CX2-32S-F
- DCS-7060CX2-32S-R
- DCS-7060SX2-48YC6-F
- DCS-7060SX2-48YC6-R
- DCS-7150S-24-CL-F
- DCS-7150S-24-CL-R
- DCS-7150S-24-CL-SSD-F
- DCS-7150S-24-CL-SSD-R
- DCS-7150S-24-F
- DCS-7150S-24-R
- DCS-7150S-52-CL-F
- DCS-7150S-52-CL-R
- DCS-7150S-52-CL-SSD-F
- DCS-7150S-52-CL-SSD-R
- DCS-7150S-64-CL-F
- DCS-7150S-64-CL-R
- DCS-7150S-64-CL-SSD-F
- DCS-7150S-64-CL-SSD-R
- DCS-7150SC-24-CLD-F
- DCS-7150SC-24-CLD-R
- DCS-7150SC-64-CLD-F
- DCS-7150SC-64-CLD-R
- DCS-7160-32CQ-F
- DCS-7160-32CQ-R
- DCS-7160-48TC6-F
- DCS-7160-48TC6-R
- DCS-7160-48YC6-F
- DCS-7160-48YC6-R
- DCS-7170-32C-F
- DCS-7170-32C-M-F
- DCS-7170-32C-M-R
- DCS-7170-32C-R
- DCS-7170-32CD-F
- DCS-7170-32CD-R
- DCS-7170-64C-F
- DCS-7170-64C-M-F
- DCS-7170-64C-M-R
- DCS-7170-64C-R
- DCS-7250QX-64-F
- DCS-7250QX-64-R
- DCS-7280QRA-C36S-M-R
- DCS-7280QRA-C36S-R
- DCS-7280SE-64-F
- DCS-7280SE-64-R
- DCS-7280SE-68-F
- DCS-7280SE-68-R
- DCS-7280SE-72-F
- DCS-7280SE-72-R
- DCS-7280SR-48C6-F
- DCS-7280SR-48C6-M-F
- DCS-7280SR-48C6-M-R
- DCS-7280SR-48C6-R
- DCS-7280SR2-48YC6-F
- DCS-7280SR2-48YC6-M-F
- DCS-7280SR2-48YC6-M-R
- DCS-7280SR2-48YC6-R
- DCS-7280SR2A-48YC6-F
- DCS-7280SR2A-48YC6-M-F
- DCS-7280SR2A-48YC6-M-R
- DCS-7280SR2A-48YC6-R
- DCS-7280SR2K-48C6-M-F
- DCS-7280SR2K-48C6-M-R
- DCS-7280SRA-48C6-F
- DCS-7280SRA-48C6-M-F
- DCS-7280SRA-48C6-M-R
- DCS-7280SRA-48C6-R
- DCS-7280SRAM-48C6-F
- DCS-7280SRAM-48C6-R
- DCS-7280SRM-40CX2-F
- DCS-7280SRM-40CX2-R
- DCS-7280TR-48C6-F
- DCS-7280TR-48C6-M-F
- DCS-7280TR-48C6-M-R
- DCS-7280TR-48C6-R
- DCS-7280TRA-48C6-F
- DCS-7280TRA-48C6-M-F
- DCS-7280TRA-48C6-M-R
- DCS-7280TRA-48C6-R
- FAN-7000-F
- FAN-7000-R
- FAN-7000H-F
- FAN-7000H-R

- DCS-7050SX-96-SSD-F
- DCS-7050SX-96-SSD-R
- DCS-7050SX2-128-F
- DCS-7050SX2-128-R
- DCS-7050SX2-72Q-F
- DCS-7050SX2-72Q-R
- DCS-7050SX3-48YC12-F
- DCS-7050TX-128-F
- DCS-7050TX-128-R
- DCS-7050TX-128-SSD-F
- DCS-7050TX-128-SSD-R
- DCS-7050TX-48-F
- DCS-7050TX-48-R
- DCS-7050TX-48-SSD-F
- DCS-7050TX-48-SSD-R
- DCS-7050TX-64-F
- DCS-7050TX-64-R
- DCS-7050TX-64-SSD-F
- DCS-7050TX-64-SSD-R
- DCS-7050TX-72-F
- DCS-7050TX-72-R
- DCS-7050TX-72-SSD-F
- DCS-7050TX-72-SSD-R
- DCS-7050TX-72Q-F
- DCS-7050TX-72Q-R
- DCS-7050TX-96-F
- DCS-7050TX-96-R
- DCS-7050TX-96-SSD-F
- DCS-7050TX-96-SSD-R
- DCS-7050TX2-128-F
- DCS-7050TX2-128-R
- DCS-7250QX-64-SSD-F
- DCS-7250QX-64-SSD-R
- DCS-7260CX-64-F
- DCS-7260CX-64-R
- DCS-7260CX-64-SSD-F
- DCS-7260CX-64-SSD-R
- DCS-7260CX3-64-F
- DCS-7260CX3-64-R
- DCS-7260CX3-64E-F
- DCS-7260CX3-64E-R
- DCS-7260QX-64-F
- DCS-7260QX-64-R
- DCS-7260QX-64-SSD-F
- DCS-7260QX-64-SSD-R
- DCS-7280CR-48-F
- DCS-7280CR2-60-F
- DCS-7280CR2A-30-F
- DCS-7280CR2A-60-F
- DCS-7280CR2K-30-F
- DCS-7280CR2K-60-F
- DCS-7280CR2M-30-F
- DCS-7280QR-C36-F
- DCS-7280QR-C36-M-F
- DCS-7280QR-C36-M-R
- DCS-7280QR-C36-R
- DCS-7280QR-C72-F
- DCS-7280QR-C72-M-F
- DCS-7280QR-C72-M-R
- DCS-7280QR-C72-R
- DCS-7280QRA-C36S-F
- DCS-7280QRA-C36S-M-F
- FAN-7001C-F
- FAN-7001D-F
- FAN-7002-F
- FAN-7002-R
- FAN-7002H-F
- FAN-7002H-R
- FAN-7010
- PWR-1100AC-F
- PWR-1100AC-R
- PWR-1600AC-F
- PWR-1900-DC-F
- PWR-1900-DC-R
- PWR-1900AC-F
- PWR-1900AC-R
- PWR-400AC-BLUE
- PWR-400AC-RED
- PWR-460AC-F
- PWR-460AC-R
- PWR-460DC-F
- PWR-460DC-R
- PWR-500-DC-F
- PWR-500-DC-R
- PWR-500AC-F
- PWR-500AC-R
- PWR-745AC-F
- PWR-745AC-R
- PWR-747AC-F
- PWR-747AC-R
- PWR-750AC-F
- PWR-750AC-R

Modular

- 7300-SUP
- 7300-SUP-D
- 7300X-32Q-LC
- 7300X-64S-LC
- 7300X-64T-LC
- 7300X3-32C-LC
- 7300X3-48YC4-LC
- 7304X-FM
- 7304X3-FM
- 7308X-FM
- 7308X3-FM
- 7316X-FM
- 7320X-32C-LC
- 7500E-48S-LC
- 7500E-48T-LC
- 7500E-6C2-LC
- 7500E-6CFPX-LC
- 7500E-72S-LC
- 7500E-SUP
- 7500E-SUP-D
- 7500R-36CQ-LC
- 7500R-36Q-LC
- 7500R-48S2CQ-LC
- 7500R-8CFPX-LC
- 7500R2-18CQ-LC
- 7500R2-36CQ-LC
- 7516-SUP2-D
- 7516R-FM
- DCS-7304
- DCS-7308
- DCS-7316
- DCS-7504
- DCS-7504N
- DCS-7508
- DCS-7508N
- DCS-7512N
- DCS-7516N
- FAN-7000H-F
- FAN-7002-F

- 7324X-FM
- 7328X-FM
- 7368
- 7368-16C
- 7368-SUP
- 7368-SUP-D
- 7368X4-SC
- 7500-SUP2
- 7500-SUP2-D
- 7500E-12CM-LC
- 7500E-12CQ-LC
- 7500E-36Q-LC
- 7500R2A-36CQ-LC
- 7500R2AK-36CQ-LC
- 7500R2AK-48YCQ-LC
- 7500R2AM-36CQ-LC
- 7500R2M-36CQ-LC
- 7500RM-36CQ-LC
- 7504E-FM
- 7504R-FM
- 7508E-FM
- 7508R-FM
- 7512R-FM
- 7516-SUP2
- FAN-7002-R
- FAN-7002H-F
- PWR-2700-DC-F
- PWR-2700-DC-R
- PWR-2900AC
- PWR-3K-AC-F
- PWR-3K-AC-R
- PWR-3K-DC-RED
- PWR-3KT-AC-BLUE
- PWR-3KT-AC-RED
- PWR-3KT-DC-BLUE

Transceiver

- 1000BASE-BX10-D
- 1000BASE-BX10-U
- 1000BASE-LX
- 1000BASE-SX
- 1000BASE-T
- 100GBASE-AOC
- 100GBASE-CR4
- 100GBASE-CWDM4
- 100GBASE-DR
- 100GBASE-ERL4
- 100GBASE-LR4
- 100GBASE-LRL4
- 100GBASE-PSM4
- 100GBASE-SR4
- 100GBASE-SRBD
- 100GBASE-SWDM4
- 100GBASE-XSR4
- 100GE-DWDM2
- 10GBASE-AOC
- 10GBASE-CR
- 10GBASE-DWDM
- 10GBASE-DWDM-T
- 10GBASE-DWDM-ZR
- 10GBASE-ER
- 10GBASE-ERBD-D
- 10GBASE-ERBD-U
- 10GBASE-ERLBD-D
- 10GBASE-ERLBD-U
- 10GBASE-LR
- 10GBASE-LRL
- 10GBASE-SR
- 10GBASE-SRL
- 10GBASE-ZR
- 25GBASE-AOC
- 25GBASE-CR
- 25GBASE-LR
- 25GBASE-SR
- 40GBASE-AOC
- 40GBASE-CR4
- 40GBASE-ER4
- 40GBASE-LR4
- 40GBASE-LRL4
- 40GBASE-PLR4
- 40GBASE-PLRL4
- 40GBASE-SR4
- 40GBASE-SRBD
- 40GBASE-UNIV
- 40GBASE-XSR4
- CFP2-100G-DWDM-DCO
- CFP2-100GBASE-ER4
- CFP2-100GBASE-LR4
- CFP2-100GBASE-XSR10
- CFPX-100G-DWDM
- CFPX-200G-DWDM

Picking compatible versions for MLAG ISSU Upgrade/ Downgrade

Upgrading/downgrading using the MLAG ISSU procedure, from release EOS-A.B.C to release EOS-D.E.F can be done in one or multiple phases.

1. If A.B.C is a compatible EOS version for D.E.F (refer to the A.B.X or D.E.X release table) then you can directly do MLAG ISSU using the D.E.F version by following the MLAG ISSU procedure. (A.B.C -> D.E.F).
   Example: The customer wants to upgrade from 4.12.10 to 4.15.4F. In the Release 4.15.X MLAG ISSU Table, the 4.15.4F version has 4.12.10 as a MLAG ISSU compatible version, so customer can directly do an MLAG ISSU upgrade from 4.12.10 to 4.15.4F.
2. If A.B.C doesn't exist in the list of compatible versions for D.E.F (refer to the A.B.X or D.E.X release table), then look for the A.B.Y version in D.E.F's ISSU compatible EOS versions, do MLAG ISSU upgrade/downgrade from A.B.C to A.B.Y and then do MLAG ISSU upgrade/downgrade to D.E.F by following MLAG ISSU procedure (A.B.C -> A.B.Y -> D.E.F). Here, Y can be greater than or less than C.
   Example: The customer wants to upgrade from 4.14.1F to 4.15.4F. In Release 4.15.X MLAG ISSU Table 4.15.4F version has 4.14.9F as MLAG ISSU compatible version and Release 4.14.X MLAG ISSU Table 4.14.9F has 4.14.1F as MLAG ISSU compatible version, so customer can do MLAG ISSU upgrade from 4.14.1F to 4.14.9F and then to 4.15.4F.
3. Look for an D.E.Z version that's MLAG ISSU compatible with both A.B.C and D.E.F. Do MLAG ISSU upgrade/downgrade from A.B.C to D.E.Z and then do MLAG ISSU upgrade/downgrade to D.E.F (A.B.C -> D.E.Z -> D.E.F). Here Z can be greater than or less than F.
   Example: The customer wants to upgrade from 4.14.7F to 4.15.4F. In Release 4.15.X MLAG ISSU Table 4.15.0F version has 4.14.7F as MLAG ISSU compatible version and 4.15.4F has 4.15.0F as MLAG ISSU compatible version, so customer can do MLAG ISSU upgrade from 4.14.7F to 4.15.0F and then to 4.15.4F.
4. Look for an G.H.I version that's MLAG ISSU compatible with both A.B.C and D.E.F. Do an MLAG ISSU upgrade/downgrade from A.B.C to G.H.I and then do Mlag ISSU upgrade/downgrade to D.E.F (A.B.C -> G.H.I -> D.E.F).
5. In a highly unlikely scenario, if any of the above procedures are not acceptable, MLAG ISSU can be achieved by picking multiple intermediate versions.

Release 4.22.X MLAG ISSU Table

EOS Version	Mlag ISSU compatible EOS versions
4.22.0F	4.14.16M, 4.15.10M, 4.16.14M, 4.17.10M, 4.18.10M, 4.19.0F-4.19.3F, 4.19.4M-4.19.12M, 4.19.6.3M, 4.20.1F-4.20.6F, 4.20.7M-4.20.12M,4.21.0F-4.21.5F

Release 4.21.X MLAG ISSU Table

EOS Version	Mlag ISSU compatible EOS versions
4.21.0F	4.14.16M, 4.15.10M, 4.16.14M, 4.17.9M, 4.18.8M, 4.19.0F-4.19.3F, 4.19.4M-4.19.9M, 4.19.6.3M, 4.20.1F-4.20.6F, 4.20.7M
4.21.1F	4.14.16M, 4.15.10M, 4.16.14M, 4.17.10M, 4.18.9M, 4.19.0F-4.19.3F, 4.19.4M-4.19.10M, 4.19.6.3M, 4.20.1F-4.20.6F, 4.20.7M-4.20.9M,4.21.0F
4.21.2F	4.14.16M, 4.15.10M, 4.16.14M, 4.17.10M, 4.18.10M, 4.19.0F-4.19.3F, 4.19.4M-4.19.10M, 4.19.6.3M, 4.20.1F-4.20.6F, 4.20.7M-4.20.10M,4.21.0F-4.21.1F
4.21.3F	4.14.16M, 4.15.10M, 4.16.14M, 4.17.10M, 4.18.10M, 4.19.0F-4.19.3F, 4.19.4M-4.19.11M, 4.19.6.3M, 4.20.1F-4.20.6F, 4.20.7M-4.20.11M,4.21.0F-4.21.2F
4.21.4F	4.14.16M, 4.15.10M, 4.16.14M, 4.17.10M, 4.18.10M, 4.19.0F-4.19.3F, 4.19.4M-4.19.11M, 4.19.6.3M, 4.20.1F-4.20.6F, 4.20.7M-4.20.11M,4.21.0F-4.21.3F
4.21.5F	4.14.16M, 4.15.10M, 4.16.14M, 4.17.10M, 4.18.10M, 4.19.0F-4.19.3F, 4.19.4M-4.19.12M, 4.19.6.3M, 4.20.1F-4.20.6F, 4.20.7M-4.20.11M,4.21.0F-4.21.4F

Release 4.20.X MLAG ISSU Table

EOS Version	Mlag ISSU compatible EOS versions
4.20.1F	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.13M, 4.16.7-FXMLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.7M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M, 4.19.0F-4.19.2F
4.20.2F	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.13M, 4.16.7-FXMLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.7M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M, 4.19.0F-4.19.2F

EOS Version	Mlag ISSU compatible EOS versions
4.20.3F	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.13M, 4.16.7-FXMLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.7M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M, 4.19.0F-4.19.2F
4.20.4F	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.13M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP,4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.9M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M-4.18.6M, 4.19.0F-4.19.3F, 4.19.4M-4.19.5M, 4.20.1F-4.20.3F
4.20.5F	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.13M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.9M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M, 4.19.0F-4.19.4M, 4.20.1F-4.20.4F
4.20.5.1F	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.13M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.9M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M, 4.19.0F-4.19.4M, 4.20.1F-4.20.5F
4.20.6F	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.13M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.9M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M-4.18.8M, 4.19.0F-4.19.7M, 4.20.1F-4.20.5F
4.20.7M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.14M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.9M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M-4.18.8M, 4.19.0F-4.19.8M, 4.19.6.3M, 4.20.1F-4.20.6F
4.20.8M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.14M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.9M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M-4.18.8M, 4.19.0F-4.19.8M, 4.19.6.3M, 4.20.1F-4.20.7M
4.20.9M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.14M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.9M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M-4.18.8M, 4.19.0F-4.19.5M, 4.19.6.3M, 4.19.7M-4.19.9M, 4.20.1F-4.20.8M

EOS Version	Mlag ISSU compatible EOS versions
4.20.10M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.14M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.10M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M-4.18.9M, 4.19.0F-4.19.5M, 4.19.6.3M, 4.19.7M-4.19.10M, 4.20.1F-4.20.9M
4.20.11M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.14M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.10M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M-4.18.10M, 4.19.0F-4.19.5M, 4.19.6.3M, 4.19.7M-4.19.11M, 4.20.1F-4.20.10M
4.20.12M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.10M, 4.15.3FX-7060X.2, 4.15.5FX-7500R, 4.16.6M-4.16.14M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.16.13FX-MLAGISSU-TWO-STEP, 4.16.9FX-7500R, 4.16.10FX-7060X, 4.17.0F-4.17.3F, 4.17.4M-4.17.10M, 4.18.0F, 4.18.1F-4.18.4F, 4.18.2-REV2-FX, 4.18.5M-4.18.10M, 4.19.0F-4.19.5M, 4.19.6.3M, 4.19.7M-4.19.11M, 4.20.1F-4.20.11M

Release 4.19.X MLAG ISSU Table

EOS Version	Mlag ISSU compatible EOS versions
4.19.0F	4.16.6M-4.16.12M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.3F, 4.17.4M-4.17.7M, 4.18.0F, 4.18.1.1F, 4.18.2F, 4.18.2-REV2-FX.1, 4.18.3.1F
4.19.1F	4.16.6M-4.16.12M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.3F, 4.17.4M-4.17.7M, 4.18.0F-4.18.4F
4.19.2F	4.16.6M-4.16.12M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.3F, 4.17.4M-4.17.7M, 4.18.0F, 4.18.1.1F, 4.18.2F, 4.18.2-REV2-FX.1, 4.18.3.1F, 4.19.0F-1F
4.19.3F	4.16.6M-4.16.12M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSUTWO-STEP, 4.17.0F-4.17.3F, 4.17.4M-4.17.7M, 4.18.0F, 4.18.1.1F, 4.18.2F, 4.18.2-REV2-FX.1, 4.18.3.1F, 4.19.0F-2F, 4.19.2.1F
4.19.4F	4.16.6M-4.16.12M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSUTWO-STEP, 4.17.0F-4.17.3F, 4.17.4M-4.17.7M, 4.18.0F, 4.18.1.1F, 4.18.2F, 4.18.2-REV2-FX.1, 4.18.3.1F, 4.19.0F-2F, 4.19.2.1F
4.19.5M	4.16.6M-4.16.12M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSUTWO-STEP, 4.17.0F-4.17.3F, 4.17.4M-4.17.7M, 4.18.0F, 4.18.1.1F, 4.18.2F, 4.18.2-REV2-FX.1, 4.18.3.1F, 4.19.0F-2F, 4.19.2.1F

Release 4.18.X MLAG ISSU Table

EOS Version	Mlag ISSU compatible EOS versions
4.18.0F	4.16.6M-4.16.9M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.3F
4.18.1F	4.16.6M-4.16.10M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.3F, 4.17.4M, 4.18.0F
4.18.1.1F	4.16.6M-4.16.10M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.3F, 4.17.4M, 4.18.0F, 4.18.1F
4.18.2F	4.16.6M-4.16.10M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.3F, 4.17.4M, 4.18.0F, 4.18.1F, 4.18.1.1F
4.18.3F	4.16.6M-4.16.10M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.3F, 4.17.4M, 4.18.0F, 4.18.1F, 4.18.1.1F 4.18.2F

Release 4.17.X MLAG ISSU Table

EOS Version	Mlag ISSU compatible EOS versions
4.17.0F	4.16.6M
4.17.1F	4.16.6M-4.16.7M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.17.0F
4.17.2F	4.16.6M-4.16.8M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.1F
4.17.3F	4.16.6M-4.16.9M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.2F
4.17.4M	4.16.6M-4.16.9M, 4.16.7FX-MLAGISSU-TWO-STEP, 4.16.8FX-MLAGISSU-TWO-STEP, 4.17.0F-4.17.3F

Release 4.16.X MLAG ISSU Table

EOS Version	Mlag ISSU compatible EOS versions
4.16.6M	4.16.7FX-MLAGISSU-TWO-STEP
4.16.7M	4.16.6M, 4.16.7FX-MLAGISSU-TWO-STEP
4.16.8M	4.16.6M-4.16.7M, 4.16.7FX-MLAGISSU-TWO-STEP - 4.16.8FX-MLAGISSU-TWO-STEP
4.16.9M	4.16.6M-4.16.8M, 4.16.7FX-MLAGISSU-TWO-STEP - 4.16.8FX-MLAGISSU-TWO-STEP
4.16.10M	4.16.6M-4.16.9M, 4.16.7FX-MLAGISSU-TWO-STEP - 4.16.8FX-MLAGISSU-TWO-STEP

EOS Version	Mlag ISSU compatible EOS versions
4.16.7FX-MLAGISSU-TWO-STEP	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.14M, 4.15.0F-4.15.4F, 4.15.5M-4.15.7M, 4.16.6M, 4.16.7M, 4.17.0F
4.16.8FX-MLAGISSU-TWO-STEP	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.15M, 4.15.0F-4.15.4F, 4.15.5M-4.15.8M, 4.16.6M-4.16.8M, 4.17.0F-4.17.1F

Release 4.15.X MLAG ISSU Table

EOS Version	Mlag ISSU compatible EOS versions
4.15.0F	4.9.11, 4.10.8.1, 4.11.11, 4.12.10, 4.13.12M, 4.14.7M
4.15.1F	4.9.11, 4.10.8.1, 4.11.11, 4.12.10, 4.13.12M, 4.14.7M, 4.15.0F
4.15.2F	4.9.11, 4.10.8.1, 4.11.11, 4.12.10, 4.13.12M, 4.14.9M, 4.15.0F-4.15.1F
4.15.3F	4.9.11, 4.10.8.1, 4.11.11, 4.12.10, 4.13.12M, 4.14.9M, 4.15.0F-4.15.2F
4.15.4F	4.9.11, 4.10.8.1, 4.11.11, 4.12.10, 4.13.12M, 4.14.9M, 4.15.0F-4.15.3F
4.15.5M	4.9.11, 4.10.8.1, 4.11.11, 4.12.10, 4.13.12M, 4.14.10M, 4.15.0F-4.15.4F
4.15.6M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.15M, 4.14.12M, 4.15.0F-4.15.4F, 4.15.5M
4.15.7M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.13M, 4.15.0F-4.15.4F, 4.15.5M-4.15.6M
4.15.8M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.13M, 4.15.0F-4.15.4F, 4.15.5M-4.15.7M
4.15.9M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.15M, 4.15.0F-4.15.4F,4.15.5M-4.15.8M
4.15.10M	4.9.11, 4.10.8.1, 4.11.12, 4.12.11, 4.13.16M, 4.14.16M, 4.15.0F-4.15.4F,4.15.5M-4.15.9M

Enhancements in 4.22.0F

General

- A new EOS CLI command to check VXLAN configuration. (171387)
- Support for IPv4 and IPv6 ICMP MPLS extensions. (176993)
- Add support for NetSync Monitor: Management messages (IEEE1558 - PTP). (215337)

- Add support for NetSync Monitor: DelayRequest messages (IEEE1558 - PTP) with custom TLVs attached. (215339)
- Add interface-address option to tcpdump command so that it supports capturing traffic on loopback interfaces. (256104)
- Added support for DHCPv6 Relay over a non-default VRF (267203)
- Secureboot is an anti tamper mechanism that works by verifiying the cryptographic signature of SWI images. (286447)
- Add OpenConfig support for VXLAN model (287165)
- Added support for software sampled flow tracking with IPFIX export functionality. (351804)
- Support for these products have been removed from non-2GB EOS images. These products are only supported in 2GB EOS images such as EOS-2GB.swi. (359921)
  SKUs Affected: DCS-7050QX-32-F, DCS-7050QX-32-R, DCS-7050QX-32-SSD-F, DCS-7050QX-32-SSD-R, DCS-7150S-24-CL-F, DCS-7150S-24-CL-R, DCS-7150S-24-CL-SSD-F, DCS-7150S-24-CL-SSD-R, DCS-7150S-24-F, DCS-7150S-24-R, DCS-7150S-52-CL-F, DCS-7150S-52-CL-R, DCS-7150S-52-CL-SSD-F, DCS-7150S-52-CL-SSD-R, DCS-7150S-64-CL-F, DCS-7150S-64-CL-R, DCS-7150S-64-CL-SSD-F and DCS-7150S-64-CL-SSD-R

Layer 1

- Support for the new CLI command "clear phy counters [<interface range>]". This command will clear Forward Error Correction (FEC) counters in the output of "show interfaces phy detail" on supported platforms. (261237)
  SKUs Affected: 7300X3-32C-LC, 7300X3-48YC4-LC, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050SX3-48YC12-F, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-F, DCS-7060SX2-48YC6-R, DCS-7260CX3-64-F and DCS-7260CX3-64-R
- Support for clause 73 auto negotiation over twin-ax copper cables at 25G, 40G and 100G. 25G conforms to IEEE 802.3by. There is no support for Ethernet Consortium 25G or 50G auto negotiation. (281281)
  Series Affected: DCS-7170 Series
- Add support of state-of-polarization fast tracking mode to Acacia CFP2-DCO transceivers, used in DCS-7280SRM-40CX2 switches (344079)
  SKUs Affected: CFP2-100G-DWDM-DCO, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R

Layer 2

- Recover port from backup role once self loop is resolved (226405)

Layer 3

- 'show arp' and 'show ipv6 neighbors' now show correct age information. If an ARP or ND entry is first created, the age might show 'N/A' initially until it is refreshed the first time. (1704)
- Support for "show diff" and "show pending" under IPv4 and IPv6 prefix list configuration mode (150922)
- Support for OSPFv2 Secure Hash Algorithms (SHA) Authentication support (SHA1, SHA256, SHA384 and SHA512) as defined in RFC 5709 as an alternative to the MD5 Authentication. (179611)
- Support for ISIS Administrative Tags. (191761)
- Support for resequencing prefix lists by the new "resequence" command under prefix list configuration mode. (203652)
- Support for IPv6 Control-Plane ACLs in non-default VRFs (209778)
- LDP can be limited to selected interfaces by configuring LDP to have interfaces disabled by default (using "config mpls ldp interface disabled default"), and then explicitly enabling individual interfaces (using "config interface <interface-id> mpls ldp interface"). (249835)
- Add support for configuration of multiple Adj-SIDs for the same IS-IS adjacency. (253511)
- Support for BFD authentication. (256193)
- Add support for configurable TC to EXP map, and EXP to TC map (258287)
  Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Support for IS-IS IPv6 Traffic Engineering extensions (272526)
- Support for removing prefix list entries without providing the sequence number. Example "no permit 1.1.1.1/32". (279238)
- Support for OSPF Traffic Engineering, RFC3630, RFC4203 and RFC7308. (283516)
- Prefix-length support for IPv6 attached-host routes (287319)
- Support for IPv6 Neighbor Discovery Proxy (359709)

Multi-agent

- 4 byte AS number support for route distinguisher (120035)

- support for neighbor prefix-list route filtering "neighbor {<ip-address> | <peer-group>} prefix-list <prefix-list-name> {in|out}" (135885)
- Support for Generalized TTL Security Mechanism (GTSM) to protect both iBGP and eBGP peer sessions (196327)
- Support for Active-Active multi-homing scenario using an EVPN control plane with VXLAN encapsulation (233390)
- ASDOT notation support (239515)
- Adds virtual IPs as receive routes in multi-agent mode (240508)
- The feature implements the route-map setter set as-path prepend last-as <n>where n stands for a number between 1 to 15 representing the number of time the last AS number should be prepended to the AS_PATH attribute. (263816)
- Support for 'match ip address dynamic prefix-list' clause in route-map (279553)
- Support for 'nexthop resolution ribs vrf-unicast-rib' knob with Mpls L3 VPN (280039)
- Support for auto-generation of control plane ACLs from BGP neighbor configuration (284423)
- 4 byte AS number support for route-target (286731)
- SNMP support for OSPF, OSPFv3 and IS-IS (288838)
- Support for VPN export/import in default VRF in Bgp (291164)
- Support for specifying RIBs for next hop resolution of routes in IPv4 and IPv6 unicast address families (294317)
- "local-as <AS> no-prepend replace-as fallback" is now supported for BGP neighbors in mutli-agent mode (294594)
- The tables ptpbaseClockCurrentDSTable and ptpbaseClockDefaultDSTable are now supported from PTP MIB RFC 8173. (355256)
- Add support for RFC7606 excluding AFI/SAFI disable on malformed update. (359895)

Rib

- The RIB Route Control - FIB Policy feature introduces a mechanism for selecting which routes in the routing table (RIB) get programmed into the fowarding table (FIB). (221017)

MLAG

- Added maintenance-mode support for MLAG (113260)

DCS-7150 Series

- Support for configuring more than 255 Twice NAT groups under an interface (250846)

DCS-7300 and DCS-7500 Series

- Support High availability for ManagementActive interface in modular system through standby supervisor when Active supervisor management interface is down. (262997)

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

- Support for configuring ingress Fair Adaptive Dynamic Thresholds (116115)
- Support using unicast buffers for ingress replicated multicast traffic with small (no more than 6) replications (125536)
  Series Affected: DCS-7280R and DCS-7280R2 Series
- Enabling VLAN Aware Bundle Support in EVPN-MPLS network (173942)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Add 'show interface counters queue' support for LAG interfaces. (202290)
- Support interface reflection (210555)
- Support for VLAN mappings increased from 4096 to 8192 per FAP. (222897)
  Series Affected: DCS-7280E and DCS-7500E Series
- SSO support on 7500R-8CFPX-LC, 7500RM-36CQ-LC line cards of 7500 family. (246671)
  SKUs Affected: 7500R-8CFPX-LC and 7500RM-36CQ-LC
- show hardware capacity now shows the utilization separately for divergent and convergent for ISEM table (248305)
- Support time stamping in UTC while in Tap Aggregation mode (260241)
- Support for TCAM databases with 80bit keys for traffic steering while in Tap Aggregation mode (264441)
- Support matching of inner IP/IPv6 header fields in traffic steering policies while in Tap Aggregation mode (267045)

- Support using the same destination port in multiple monitoring sessions simultanously (270248)
- Support for BGP Flowspec Traffic-rate action (271304)
- Support for static GRE tunnel interface counters in Jericho/Jericho+/ QumranMx. (272337)
- Support to match on nexthop-group in Qos policy-map (272898)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Shows ACL counters on a per chip basis when 'detail' option is given to commands such as 'show ip access-lists <aclName> detail' (273545)
- Add support for monitoring VOQ deletes and reporting them in "show platform fap voq delete" command output. (275919)
- VRF support for Macsec proxy sub interface. (281515)
- Add support for 2x50G breakout on the following:
  DCS-7280QRA-C36S: Odd Numbered ports in range 13-23
  DCS-7280QR-C72: Odd numbered ports in ranges 11-25, 47-61.
  Even numbered port in this range are disabled when the prior odd numbered port is configured for 2x50G. (285880)
- Support for QoS policy-map counters on 7280 systems (300399)
  SKUs Affected: DCS-7280SE-64-F, DCS-7280SE-64-R, DCS-7280SE-68-F, DCS-7280SE-68-R, DCS-7280SE-72-F and DCS-7280SE-72-R
- Support for shaping on L2 Subinterfaces (312422)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Support two-way ports on port-channels while in Tap Aggregation mode (339737)
- Support for Stateful Switch Over while in Tap Aggregation mode (342110)
- Support for Overlay VRF in Static GRE Tunnel Interfaces on Jericho, Jericho +, Qumran MX and Qumran AX platform. (346881)
- Support for 'output interface' action in Direct Flow. (347503)
  Series Affected: DCS-7020, DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- Support for CPU traffic policy. (351185)
- When ACL configuration is directly modified in a config session and the acl configuration fails, the config session is rolled back. (351821)
- Allow configuration of explicit congestion notification marking probability (352616)

- Support for DCS-7020SR-32C2-R (363960)
  SKUs Affected: DCS-7020SR-32C2-R

DCS-7280R and DCS-7500R Series

- Support for Egress Acls for bridged IPv4 traffic. (190829)
  Series Affected: DCS-7020, DCS-7280R and DCS-7280R2 Series
- Add support for MPLS label lookup based on two MPLS labels. (271525)
- Support for loose uRPF mode. (283626)
- uRPF support with "ip hardware fib optimize prefix-length" configuration. (285118)
- Support for static VRF label decap with MPLS payload (291269)

DCS-7020 Series

- Provide support for Ipsec tunnels. Up to 128 concurrent Ipsec tunnels are permitted using tunnel mode only, no support for transport mode. Two authentication/encryption algorithm combinations are supported: SHA1 with AES128 and SHA256 with AES256. (229966)
  SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R
- MPLS routing, and Psuedowire support (360254)

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

- Support clearing of PFC (Priority Flow Control) Watchdog counters. (161993)
- Support for GRE encapsulation through Nexthop Groups and GRE decapsulation through Decap Groups. (205153)
- EAPI support for the "show cpu counters queue" command. (251672)
  Series Affected: DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7250X, DCS-7260X and DCS-7300X Series
- Add support for LANZ Mirroring (253307)
  Series Affected: DCS-7300 Series
- Platform interface counters are accessible via SNMP in the aristaXgsIfTable of the aristaXgsMIB. (261048)
- Added support for policy-maps on vxlan interface to rate limit encapped vxlan routed traffic per vlan/VNI and all types of vxlan decap traffic. (268365)
  SKUs Affected: 7300-SUP, 7300-SUP-D, 7300X-32Q-LC, 7300X-64S-LC, 7300X-64T-LC, 7300X3-32C-LC, 7300X3-48YC4-LC, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050QX-32-F, DCS-7050QX-32-R, DCS-7050QX-32-SSD-F,

DCS-7050QX-32-SSD-R, DCS-7050QX-32S-F, DCS-7050QX-32S-R, DCS-7050QX-32S-SSD-F, DCS-7050QX-32S-SSD-R, DCS-7050QX2-32S-F, DCS-7050QX2-32S-R, DCS-7050SX-128-F, DCS-7050SX-128-R, DCS-7050SX-128-SSD-F, DCS-7050SX-128-SSD-R, DCS-7050SX-64-F, DCS-7050SX-64-R, DCS-7050SX-64-SSD-F, DCS-7050SX-64-SSD-R, DCS-7050SX-72-F, DCS-7050SX-72-R, DCS-7050SX-72-SSD-F, DCS-7050SX-72-SSD-R, DCS-7050SX-72Q-F, DCS-7050SX-72Q-R, DCS-7050SX-96-F, DCS-7050SX-96-R, DCS-7050SX-96-SSD-F, DCS-7050SX-96-SSD-R, DCS-7050SX2-128-F, DCS-7050SX2-128-R, DCS-7050SX2-72Q-F, DCS-7050SX2-72Q-R, DCS-7050SX3-48YC12-F, DCS-7050TX-128-F, DCS-7050TX-128-R, DCS-7050TX-128-SSD-F, DCS-7050TX-128-SSD-R, DCS-7050TX-48-F, DCS-7050TX-48-R, DCS-7050TX-48-SSD-F, DCS-7050TX-48-SSD-R, DCS-7050TX-64-F, DCS-7050TX-64-R, DCS-7050TX-64-SSD-F, DCS-7050TX-64-SSD-R, DCS-7050TX-72-F, DCS-7050TX-72-R, DCS-7050TX-72-SSD-F, DCS-7050TX-72-SSD-R, DCS-7050TX-72Q-F, DCS-7050TX-72Q-R, DCS-7050TX-96-F, DCS-7050TX-96-R, DCS-7050TX-96-SSD-F, DCS-7050TX-96-SSD-R, DCS-7050TX2-128-F, DCS-7050TX2-128-R, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-F, DCS-7060SX2-48YC6-R, DCS-7260CX-64-F, DCS-7260CX-64-R, DCS-7260CX-64-SSD-F, DCS-7260CX-64-SSD-R, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F, DCS-7260CX3-64E-R, DCS-7260QX-64-F, DCS-7260QX-64-R, DCS-7260QX-64-SSD-F and DCS-7260QX-64-SSD-R

- Support for queue counter rates for enqueued/dropped packets and bytes via CLI command "show interface counter queue rate". (268800)

- Add support for LANZ Mirroring (289593)
  Series Affected: DCS-7260X3 Series

- IPv6 multicast routing is now supported. Currently, PIM-SSM v6 is the only supported protocol. (295409)
  Series Affected: DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X3, DCS-7250X, DCS-7260X, DCS-7260X3, DCS-7300, DCS-7300X and DCS-7320 Series

- The output of "show cpu counters queue" is now available via terminattr. (345252)

- Support for DCS-7368 (357905)

- Enable MSSL3 support. (359679)
  Series Affected: DCS-7050, DCS-7050T, DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3 and DCS-7250X Series

DCS-7060X and DCS-7300 Series

- Added support for policy-maps on vxlan interface to rate limit bridged encap vxlan traffic per vlan/VNI. (376264)
  SKUs Affected: 7300-SUP, 7300-SUP-D, 7300X-32Q-LC, 7300X-64S-LC, 7300X-64T-LC, 7300X3-32C-LC, 7300X3-48YC4-LC, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-

F, DCS-7060SX2-48YC6-R, DCS-7260CX-64-F, DCS-7260CX-64-R, DCS-7260CX-64-SSD-F, DCS-7260CX-64-SSD-R, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F, DCS-7260CX3-64E-R, DCS-7260QX-64-F, DCS-7260QX-64-R, DCS-7260QX-64-SSD-F and DCS-7260QX-64-SSD-R

- Support for high speed interface counters polling. (391178)

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

- SVI Ingress counters are supported on 7050X3 series switches. SVI Ingress counted can be enabled through the following configuration command switch(config)# "hardware counter feature vlan-interface in". and the counter will be available through the following commands, switch# sh interface vlan <vlanId> counters (255586)
  SKUs Affected: DCS-7050CX3-32S-F, DCS-7050CX3-32S-R and DCS-7050SX3-48YC12-F
- Support for VNI ingress and egress counters (342058)
- Support for high speed interface counters polling. (391179)

DCS-7160 Series

- Support for LACP and hardware programming interlock on 7160 platform (164520)
- This adds FIB-RIB-ACK feature on 7160. This allows 'bgp update wait-install' command to work on 7160 (182928)

Transceiver Series

- Support media type 100BASE-LX10 (252566)

vEOS Router Series

- Support for tunnel interface counters in vEOS DPDK mode (MODE=sfe in /mnt/flash/veos-config). (342087)
- vEOS (dpdk flavor) has support for ingress policing and egress port shaping (359688)
- Add support Dpdk for vEOS on Azure Cloud platform. You can change from kernel to DPDK mode by setting MODE=sfe in /mnt/flash/veos-config and rebooting (359689)
- Adds support for L3 VXLAN in DPDK mode. This is useful for buildling overlay networks on public cloud networks using multipoint tunnels rather than GRE tunnels. (359690)

Resolved Caveats in 4.22.0F

Accelerated System Update

- Interfaces using certain optical transceviers may flap during "reload hitless" and "reload fast-boot" (363205)
  SKUs Affected: DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050QX-32-F, DCS-7050QX-32-R, DCS-7050QX-32-SSD-F, DCS-7050QX-32-SSD-R, DCS-7050QX-32S-F, DCS-7050QX-32S-R, DCS-7050QX-32S-SSD-F, DCS-7050QX-32S-SSD-R, DCS-7050QX2-32S-F, DCS-7050QX2-32S-R, DCS-7050SX-72Q-F, DCS-7050SX-72Q-R, DCS-7050SX2-72Q-F, DCS-7050SX2-72Q-R, DCS-7050SX3-48YC12-F, DCS-7050TX-128-F, DCS-7050TX-128-R, DCS-7050TX-128-SSD-F, DCS-7050TX-128-SSD-R, DCS-7050TX2-128-F, DCS-7050TX2-128-R, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-F, DCS-7060SX2-48YC6-R, DCS-7250QX-64-F, DCS-7250QX-64-R, DCS-7250QX-64-SSD-F, DCS-7250QX-64-SSD-R, DCS-7260CX-64-F, DCS-7260CX-64-R, DCS-7260CX-64-SSD-F and DCS-7260CX-64-SSD-R
- Enable SSU support for multiple MSTP instances. (369596)
- If SSU has been performed from >=4.21.3F and then a subsequent downgrade to a release <4.21.0F occurs, performing SSU to >=4.21.3F may result in a fatal error reboot causing extended traffic outage. A workaround is to remove /mnt/flash/asu-persistent-store.json prior to performing SSU from <4.21.0F. (369618)

CVX

- OpenStack may fail to provision VLAN to VNI mappings on PortChannel interfaces

  This can be worked around by configuring the PortChannel's member interfaces as trunk ports (226933)
- Arista OpenStack integration is incompatible with TACACS+ authentication in Queens and later releases

  As a workaround, a 'roles = <role>' attribute-value pair may be added to the TACACS+ configuration for
  the group of the user configured to authenticate from the OpenStack controller (eapi_username in the
  ml2 config.) (367229)

General

- It is not possible to filter packets based on VLAN tagging when using tcpdump, such as the filters "not vlan", "vlan" or "vlan <vlanid>". (72726)

- In rare circumstances, the kernel can panic with a signature that includes: "IP: [] n_tty_receive_buf_common" (283501)
- RADIUS command authorization does not fall back to local authorization when RADIUS server is unavailable. (287025)
- If there are a large number of routes whose next-hop addresses route to a large number of interfaces (e.g., high ECMP fan-out), the Snmp agent may restart unexpectedly. (297557)
- Due to a CPU erratum, under rare circumstances, agents can unexpectedly restart due to SIGSEGV. (330509)
  SKUs Affected: DCS-7010T-48-DC-F, DCS-7010T-48-DC-R, DCS-7010T-48-F, DCS-7010T-48-R, DCS-7020SR-24C2-F, DCS-7020SR-24C2-R, DCS-7020TR-48-F, DCS-7020TR-48-R, DCS-7020TRA-48-F, DCS-7020TRA-48-R, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050QX-32S-F, DCS-7050QX-32S-R, DCS-7050QX-32S-SSD-F, DCS-7050QX-32S-SSD-R, DCS-7050QX2-32S-F, DCS-7050QX2-32S-R, DCS-7050SX-64-F, DCS-7050SX-64-R, DCS-7050SX-64-SSD-F, DCS-7050SX-64-SSD-R, DCS-7050SX-72-F, DCS-7050SX-72-R, DCS-7050SX-72-SSD-F, DCS-7050SX-72-SSD-R, DCS-7050SX-72Q-F, DCS-7050SX-72Q-R, DCS-7050SX-96-F, DCS-7050SX-96-R, DCS-7050SX-96-SSD-F, DCS-7050SX-96-SSD-R, DCS-7050SX2-128-F, DCS-7050SX2-128-R, DCS-7050SX2-72Q-F, DCS-7050SX2-72Q-R, DCS-7050SX3-48YC12-F, DCS-7050TX-48-F, DCS-7050TX-48-R, DCS-7050TX-48-SSD-F, DCS-7050TX-48-SSD-R, DCS-7050TX-64-F, DCS-7050TX-64-R, DCS-7050TX-64-SSD-F, DCS-7050TX-64-SSD-R, DCS-7050TX-72-F, DCS-7050TX-72-R, DCS-7050TX-72-SSD-F, DCS-7050TX-72-SSD-R, DCS-7050TX-72Q-F, DCS-7050TX-72Q-R, DCS-7050TX-96-F, DCS-7050TX-96-R, DCS-7050TX-96-SSD-F, DCS-7050TX-96-SSD-R, DCS-7050TX2-128-F, DCS-7050TX2-128-R, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-F, DCS-7060SX2-48YC6-R, DCS-7160-32CQ-F, DCS-7160-32CQ-R, DCS-7250QX-64-F, DCS-7250QX-64-R, DCS-7250QX-64-SSD-F, DCS-7250QX-64-SSD-R, DCS-7260QX-64-F, DCS-7260QX-64-R, DCS-7260QX-64-SSD-F, DCS-7260QX-64-SSD-R, DCS-7280QR-C36-F, DCS-7280QR-C36-M-F, DCS-7280QR-C36-M-R, DCS-7280QR-C36-R, DCS-7280QRA-C36S-F, DCS-7280QRA-C36S-M-F, DCS-7280QRA-C36S-M-R, DCS-7280QRA-C36S-R, DCS-7280SE-64-F, DCS-7280SE-64-R, DCS-7280SE-68-F, DCS-7280SE-68-R, DCS-7280SR-48C6-F, DCS-7280SR-48C6-M-F, DCS-7280SR-48C6-M-R, DCS-7280SR-48C6-R, DCS-7280SR2-48YC6-F, DCS-7280SR2-48YC6-M-F, DCS-7280SR2-48YC6-M-R, DCS-7280SR2-48YC6-R, DCS-7280SR2A-48YC6-F, DCS-7280SR2A-48YC6-M-F, DCS-7280SR2A-48YC6-M-R, DCS-7280SR2A-48YC6-R, DCS-7280SR2K-48C6-M-F, DCS-7280SR2K-48C6-M-R, DCS-7280SRA-48C6-F, DCS-7280SRA-48C6-M-F, DCS-7280SRA-48C6-M-R, DCS-7280SRA-48C6-R, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F, DCS-7280SRM-40CX2-R, DCS-7280TR-48C6-F, DCS-7280TR-48C6-M-F, DCS-7280TR-48C6-M-R, DCS-7280TR-48C6-R, DCS-7280TRA-48C6-F, DCS-7280TRA-48C6-M-F, DCS-7280TRA-48C6-M-R and DCS-7280TRA-48C6-R
- Add mapping from EOS route tag value to YANG static route set-tag. Without this a YANG update to change the set-tag leaf is lost. (344827)

- In an environment with the Macro-Segmentation Service (MSS) enabled, and a firewall connected to a MLAG service TOR pair with static port channel configuration, a MLAG peer reboot can result in a loss of traffic (for up to 4 seconds). (344868)
- If the routing and forwarding paths are mounted by OpenConfig, the OpenConfig agent may restart if a second next hop is added to an existing route. (345764)
- An empty YANG container may be created at static-routes/static/next-hops/next-hop/interface-ref when updating next hops on IP routes. (347016)
- The number of active connections to vCenter when using the VmTracer agent increases over time which may cause the VmTracer agent to restart if the number of active connections reaches the maximum allowed active connections on vCenter. (347456)
- Units configured with unit profile which has onboot enabled may enter into maintenance during config-replace.
  The work around is to issue "quiesce" followed by "no quiesce" for the affected units. (353824)
- When a Port-Channel interface has both primary and fallback PBR policy applied, the primary policy is edited and then a port-channel member is removed, this can remove the PBR policy on all members of the Port-channel. (362905)
- In OpenConfig, the management interface is removed from a non-default VRF when any interface configuration is applied to YANG (365893)
- In an EVPN VXLAN IRB setup, when Windows Cluster Failover is used to migrate a Virtual IP (VIP) to a new server, the new Arp binding may not be refreshed on all devices behind remote VTEPs. The EVPN published MAC-IP bindings also may not refresh. (368995)
- Interrupting the command "show ip route vrf all" while it's printing might lead to ConfigAgent crash (370920)

Layer 1

- Ethernet interfaces can be errdisabled due to link flaps even if show interface status state doesn't change as in the case when link debounce is configured.

  Workaround is to set the errdisable max-flaps limit to a value high enough that the flaps that are happening don't cause the port to errdisable:
  (config)#errdisable flap-setting cause link-flap max-flaps 100 time 1

Or to disable link flap monitoring with the below CLI command:
(config-if-EtX)#no monitor link-flap (349264)

- 100G ports may flap or not link up at all in macsec mode (367279)
- Configuring 'speed sfp-1000baset auto 100full' on the platforms that don't support 100M speed, might result in continuous restart of forwarding agent (370191)
- MacSec enabled port will fail to receive if MKA session fails and gets re-established and the port has transmitted more than 4294967296 packets. (371498)
  SKUs Affected: 7500RM-36CQ-LC

Layer 2

- If FIPS restrictions are enabled on only one of the MACsec partners, Security Associations may not get established resulting in traffic being blackholed.

  It is recommended to configure FIPS restrictions on both MACsec partners. (236345)
- Under scaling scenarios, when a large number of arp entries are removed, internal state of Arp agent may be out of sync. This could lead to static arp entries/neighbors not getting successfully added. When this happens, simply restart the Arp agent and Arp agent will recover properly afterwards. There is no other impact on the system after the restart. (358634)
- In the EVPN VXLAN multi-homing scenario, the SandMact agent might restart unexpectedly when the last member leaves the ESI port-channel. (361931)
- When an "interface Port-Channel" configuration is removed and added back while Lag agent is stopped, Lag agent will continuously crash when started again.

  To workaround, stop Lag agent via "agent Lag shutdown", remove the offending port-channel configuration and start the Lag agent before adding the port-channel configuration back. (371441)

Layer 3

- Under scaling scenarios, when a large number of arp entries are removed, internal state of Arp agent may be out of sync. This could lead to static arp entries/neighbors not getting successfully added. When this happens, simply restart the Arp agent and Arp agent will recover properly afterwards. There is no other impact on the system after the restart. (153502)

- When the nexthop of a BGP route is resolved over IGP and the cost of the IGP nexthop changes, the dependent BGP route is not flashed. This results in peers to which this route has been advertised not aware of the new MED or metric and would affect BGP best path calculation on them. (183526)
- Under scaling scenarios, when a large number of arp entries are removed, internal state of Arp agent may be out of sync. This could lead to static arp entries/neighbors not getting successfully added. When this happens, simply restart the Arp agent and Arp agent will recover properly afterwards. There is no other impact on the system after the restart. (228343)
- Under scaling scenarios, when a large number of arp entries are removed, internal
  state of Arp agent may be out of sync. This could lead to static arp entries/neighbors not getting successfully added. When this happens, simply restart
  the Arp agent and Arp agent will recover properly afterwards. There is no other
  impact on the system after the restart. (254692)
- When OSPF receives a NTW LSA that it did not originate and the ID is the same as an ip address that is configured on the router , OSPF increments the sequence number and re-originates the NTW LSA. This action may cause a NTW LSA flood war. (290162)
- Under scaling scenarios, when a large number of arp entries are removed, internal state of Arp agent may be out of sync. This could lead to static arp entries/neighbors not getting successfully added. When this happens, simply restart the Arp agent and Arp agent will recover properly afterwards. There is no other impact on the system after the restart. (339825)
- When RSVP-TE link protection is being used, the RIB agent may exit unexpectedly after the expiry of a non-zero IS-IS LSP fragment, provided this fragment has the TE Router ID TLV and the same TLV is also present in another non-zero IS-IS LSP fragment (341671)
- Under load, the NAT agent might unexpectedly restart. (348199)
- Under scaling scenarios, when a large number of arp entries are removed, internal state of Arp agent may be out of sync. This could lead to static arp entries/neighbors not getting successfully added. When this happens, simply restart the Arp agent and Arp agent will recover properly afterwards. There is no other impact on the system after the restart. (353052)
- The Rib agent may not terminate gracefully if RSVP FRR link protection is disabled during a change in service routing protocol mode (353419)

- Under scaling scenarios, when a large number of arp entries are removed, internal state of Arp agent may be out of sync. This could lead to static arp entries/neighbors not getting successfully added. When this happens, simply restart the Arp agent and Arp agent will recover properly afterwards. There is no other impact on the system after the restart. (356268)
- Under scaling scenarios, when a large number of arp entries are removed, internal state of Arp agent may be out of sync. This could lead to static arp entries/neighbors not getting successfully added. When this happens, simply restart the Arp agent and Arp agent will recover properly afterwards. There is no other impact on the system after the restart. (356463)
- If non-default VRF instance is deleted under router BGP configuration mode, BGP monitoring config for other VRFs (including default) can be lost and BGP monitoring service impacted (359781)
- When RSVP-TE is configured and if the number of routers belonging to a single LAN network is greater than or equal to 30, then the CSPF agent might crash. (368553)
- Some non-static MPLS routes (such as those from ISIS-SR) may become unprogrammed if ip routes exist in both the default and a non-default VRF. These MPLS routes will not reprogram until a change occurs which reprograms the ip routes resolving those MPLS routes, or some of the ip routes (which may not be apparently related to those in the default VRF) in the non-default VRF. (370769)
- Ldp agent might assert if a malformed Ldp PDU was received from peering router, e.g. with 0 bytes in Length field. (371998)
- A race condition in IpRib agent start-up may lead to the installation of drop routes instead of forwarding routes, causing the traffic to be blackholed.

  The workaround is to restart the IpRib agent. (378335)

Multi-agent

- If BGP routes are redistributed into an IGP agent (OSPF/IS-IS), the IGP agent can restart unexpectedly. (338290)
- 'match ip[v6] address dynamic prefix-list <dpl_name>' clause under route-map config mode is not supported (345442)
- On Smart System Upgrade, some BGP neighbors may converge late resulting in traffic loss. (347224)

- In multi-agent mode, when running BGP-related "show" commands, interrupting the output with ctrl-C may occasionally lead to the Bgp agent restarting. (350938)
- If a VTEP withdraws all the EVPN type-5 updates corresponding to a VRF and re-advertises those again immediately, then the IpRib agent on the peer VTEP may restart. (364277)
- Issuing a 'bgp clear counters' will no longer trigger the BGP session flap that could sporadically occur. (365016)
- In an EVPN MPLS IRB setup, when a Virtual IP (VIP) fails from a VM behind a remote PE to a new VM behind the local PE, the duplicate address detection from the new VM using ARP probes may falsely return a proxy ARP reply for the VIP from the local PE. Thereby, the new VM will incorrectly conclude that the VIP has not failed at the old VM, and the failover will not be be successful. (366043)
- In multi-agent mode, Bgp agent may restart when a large number of routes are added or removed at high peer and path scale (for example, hundreds of peers and more than 5M paths). (367521)
- Bgp agent may continuously restart after reboot if 'bgp additional-paths send' is configured under ipv6 address-family in a VRF (369049)
- A static route that specifies an interface and a next-hop, where the route covers the next-hop (for example "ip route 1.1.1.0/24 1.1.1.1 eth1") is not installed in the FIB (369410)
- If a configuration change is made to the outbound policy of a BGP peer (such as a change to the contents of the outbound route-map) shortly after the peer session is established, then existing routes whose attributes and nexthop are not affected by the policy change may fail to be advertised to the peer. (371014)
- An IPv6 MP_REACH_NLRI may contain a 32 byte nexthop, consisting of a 16 byte global ipv6 address followed by a 16 byte link local address. If the second address is not a valid link local address, the entire NLRI will be rejected. (372153)
- If 'additional-paths receive' is enabled for a BGP peer in a VRF, paths received from such peer may continue to be advertised to a VPN peer or leaked to another VRF even if the peer withdraws those paths (372746)
- The Bgp agent may restart after a peer RibOut change triggered by an outbound policy change or the peer entering or exiting maintenance mode. (375458)
- Bgp agent may unexpectedly restart on configuring "maintenance" followed by configuring "quiesce" for a unit in quick succession (375613)

Rib

- When any unused policy config is changed, like an unused route-map is configured, or Selective Route Download route-map is changed and the link corresponding to RIP neighbor is flapped then RIP neighbour will not come up. Workaround is to change any RIP config. (267437)
- Routes can get incorrectly programmed in the kernel when interface resolving the route flaps in quick succession. (348695)
- CLI may crash if a regular expression used in the "ip as-path access-list" command is about 2040 characters long and if it contains an underscore '_' or AS-dot notation '\.' near the end of the regular expression. For each entry in an as-path access-list, the current practical limit on the length of the regular expression is 2000 characters.

  A workaround is, because an as-path access-list can take multiple regular expressions, instead of using one very long regular expression, use multiple regular expressions that are less than 2000 characters long each. (364389)
- The Rib agent may restart if a dynamic peer (configured using 'bgp listen range') has an
  established session without graceful restart capabilities (aside from helper mode)
  and then the same peer opens a new session without closing the previous session and the
  new session advertises graceful restart capabilities in the BGP OPEN message.

  The workaround is to close the existing session when changing the graceful restart capabilities of a dynamic peer before starting the new session. (370128)
- When IS-IS SR Node Segment prefix reachability is ECMP , then the IS-IS SR Tunnel corresponding to the given Node Segment can be installed with next hops even when the corresponding peer attached to those nexthop interfaces does not support Segment Routing. This can result in traffic loss when the BGP prefixes are resolved over corresponding Tunnel. (373068)

MLAG

- If a physical interface that was part of peer-link port-channel is reconfigured as a normal switch port via configure replace, the physical interface state output may be missing from "show spanning-tree detail", and

"show spanning-tree topology status detail" will indicate discarding state for the Mlag peer. (366673)

- Unit System maintenance mode on an MLAG peer can cause MLAG agent to restart unexpectedly and maintenance operation will be stuck. (375497)

Multicast

- Pim might not update the RPF correctly of a S,G route when there is a unicast routing
change. (238557)

- When PIM recovers routes after ASU reload, it might leave the X flag which indicates hardware recovery is in progress, and this might lead to the route not timing out or doing a incorrect RPF lookup when unicast routing changes occur. (341876)

- In an MLAG setup, with Multicast Vxlan underlay feature being used, source route is not injected when the source is connected only to the NON-DR peer. (370077)

DCS-7170 Series

- BfnNat agent can restart under load when fragmentation support is enabled. (367262)

- UDP NAT sessions created just after device boot do not get programmed into hardware and are NAT'd via the slow path instead (372255)

- Under high load the BfnNat agent can crash leaving stale flows in hardware. (374001)

DCS-7300 and DCS-7500 Series

- When 7500R-36CQ-LC with hardware revision less than 15.20 is inserted in a modular system, it causes system fan speeds to be at 65% minimum.
Skus Affected: 7500R-36CQ-LC (355672)
  SKUs Affected: 7500R-36CQ-LC

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

- If a source MAC flaps between multiple ports connected to different ASICs, traffic may be temporarily misdirected to these ports even after the MAC flaps stop. This should be corrected once the next aging cycle kicks in. (96808)

- Egress security ACLs only apply to routed IP traffic. (191753)

- Source port filtering, when disabled, configuration may not be retained across reloads. Configuration needs to be reapplied. (260667)
- On 7500RM-36CQ-LC with hardware revision less than 21.50 and 7500R-36CQ-LC with hardware revision less than 15.20, under some environmental conditions, links connecting to external PHYs or the fabric may degrade and not come up. (262350)
  SKUs Affected: 7500R-36CQ-LC and 7500RM-36CQ-LC
- Double tagged packets ingressing on L2 sub-interfaces configured with a single outer Dot1Q encapsulation, will have both the outer and inner Dot1Q tags terminated. (292226)
- When using hardware accelerated sflow with IPv6 collectors, the UDP checksum is set to 0. This does not apply to sflowaccel-v2 profile. (294373)
- Unidirectional settings are not applied when configured on a port that is down.

  The workaround is to apply the Unidirectional settings when the link is up. (294979)
- Reconfiguring a monitor session destination from GRE-tunnel to non-GRE-tunnel can intermittently cause the SandFap agent to restart. (352042)
- With a large number of vlans and SVIs, wrong ordering of l2 and l3 vlan creation can cause Mlag peer mac/varp mac routing to be not enabled on the l3 vlan interface. (355311)
- When a PBR policy containing mpls match rules are applied on subinterfaces, SandAcl agent can crash if 'pbr-subintf' tcam profile is used. (357556)
- Changing interface speed on ports with sub-interfaces may cause the SandFap agent to restart. (359875)
- If a tool port was configured in TAP aggregation mode, and after leaving TAP aggregation mode the same port is then configured as a trunk port, any traffic egressing from the trunk port may have its VLAN tag stripped. The workaround is to toggle TAP aggregation mode. (360889)
- With hardware accelerated sflow with sflowaccel-v2 profile, sampling will stop when the TCAM profile is changed. (362826)
- When Sflow hardware acceleration is configured with multiple collectors reachable through different routes reachable via the same egress interface, one of the collectors may stop receiving samples. (367236)
- Reduced number of supported IPSec tunnels from 128 to 96. This is to reserve additional flow entries to have the flexibility to implement the design of having two encrypt and two decrypt flow entries per tunnel in the

future. Using two separate encrypt flow entries will avoid simultaneous changes to flow entry tables while traffic is processed on flow entry. (370852)

- While in Tap Aggregation mode with 48-bit header timestamp configured, a switch reload or forwarding plane restart operation may cause the header timestamps to be partially corrupted: the protocol version field in the header will show the encoding for 64-bit timestamp header instead of for 48-bit. However, the timestamp data and other header fields are correct.

  The work around is to flap 'mac timestamp header format 48-bit' in Tap Aggregation configuration mode. (371215)

- SandL3Unicast agent may restart on very high scale VRF configuration. (371806)

DCS-7280E and DCS-7500E Series

- 802.1x single-host mode is not supported. (190684)

DCS-7280R and DCS-7500R Series

- MPLS Ping/Traceroute is not supported for nexthop-groups with 'fec hierarchical' configuration. (261486)

- Programming more than 6k flows in DirectFlow may cause SandAcl to restart. (300324)

- Restarting the SandFap agent with Shaping configuration on a Subinterface which is configured on a Sesto2 Slave port may result in the agent continuously restarting. The workaround is to remove the shaping configuration from the subinterface until the SandFap agent restarts stop and then reapply the shaping configuration. (317160)

DCS-7020 Series

- 1G ports can have a packet latency jitter of up to 2us. (353182)

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

- Control packets sent out through an SVI will have their VLAN priority/CoS value set to 0. (75527)

- The link-debounce command will not correctly prevent the link going down when the debounce time is less than 2.5 seconds. (249612)
  Series Affected: DCS-7010, DCS-7050X, DCS-7050X2, DCS-7250X and DCS-7300X Series

- Irrespective of IGMP snooping being enabled or disabled, IGMP packets will not be flooded to remote VTEPs. (280073)
  Series Affected: DCS-7050, DCS-7050Q, DCS-7050S, DCS-7050T, DCS-7050X, DCS-7060X, DCS-7060X2, DCS-7060X3, DCS-7250X, DCS-7260CX, DCS-7260QX, DCS-7260X, DCS-7260X3, DCS-7300, DCS-7300X and DCS-7320 Series

- If routing is disabled in the default VRF, a restart of the L3 forwarding agent will cause temporary software forwarding, of L3 unicast traffic, in VRFs which have routing enabled. Regular hardware forwarding of the L3 traffic will resume, once the L3 forwarding agent finishes initializing. (284010)

- When DirectFlow flow is configured with Port-channel as output interface, reconfiguring the port channel may result in direct flow not working. The work around is to reconfigure the direct flow configuration. (355323)
  Series Affected: DCS-7050X2 Series

- Resetting an ASIC while it is being initialized may cause a kernel panic. (355340)

- Performing SSU may result in LACP link flap due to malformed LACP SSU configuration file resulting in extended traffic outage. The LACP SSU configuration file can be examined in /mnt/flash/fastpkttx.backup (355892)

- In a MLAG setup, unknown multicast packets entering a Vxlan VLAN where PIM is enabled, may get looped over the peer-link. (356712)
  Series Affected: DCS-7050X, DCS-7060X, DCS-7250X, DCS-7260X, DCS-7300 and DCS-7320 Series

- In a MLAG setup, non routable packets destined to VARP MAC address coming in on a Vxlan Vlan (either on Vxlan edge port or encapped and entering a core port) can get looped over the peer-link. (360305)
  Series Affected: DCS-7250X, DCS-7260X, DCS-7300 and DCS-7320 Series

- Removing or resetting a switch ASIC during bringup may cause a kernel panic. (360631)

- Performing ASU from an image before 4.21.3 can result in Port-Channels flap due to LACP timeouts. This can lead to extended traffic loss. (364508)

- When "platform trident forwarding-table partition 4" is configured to make use of ALPM mode and uRPF is not configured, traffic may get mis-directed under some conditions. (378564)

DCS-7060X and DCS-7300 Series

- SSO switchover may unexpectedly fail, and result in a switch reload. (368229)
  Series Affected: DCS-7300X Series

DCS-7260X3 Series

- Strata forwarding plane agent may restart unexpectedly with a `Shim::SdkImpl::AR_bcm_port_enable_set` signature. (284270)

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

- For TCP packets, the offset for UDF matching will start after the TCP header. (338733)
- The following DirectFlow configurations are not supported: match input interface hardware-loopback
  action output interface <> (more than one)
  action output interface hardware-loopback
  action output flood|all
  action set cos (355927)
- A Single Event Upset (SEU) report followed by a correction may leave ASIC interrupt delivery disabled. This may impact future SEU handling. (358689)
- Unknown IP multicast traffic with TTL 0 or 1 received on a PIM-enabled VxLAN Vlan is not forwarded to either local ports or remote VTEPs. (373797)

DCS-7160 Series

- If XpL3Unicast agents restarts 4 times or more, it can potentially cause host routes to not be
  programmed until XpMact is restarted (367655)
- Multicast data traffic punted to CPU could compete with LinkLocal/ISIS control traffic and control traffic could be dropped by policer. This is with non-PDP mode. (368004)

Transceiver Series

- Ports running at 100G using QSFP28 optics may experience link flaps due to a low rx-output-amplitude value.
  Starting in release EOS-4.20.5F a work around is to override this value with the interface command (this is applied to the first interface of the port, i.e the /1):
    "transceiver electrical lane 1-4 rx-output-amplitude 2"
  The current transceiver electrical tuning can be viewed with the command:
    "show interfaces transceiver tuning" (364794)

vEOS Router Series

- In DPDK mode (MODE=sfe in /mnt/flash/veos-config), vEOS instance is deployed in AWS, a dynamically hot attached network interface cannot be

moved to a VRF. Please reboot the vEOS instance after hot attach of the interface to enable VRF forwarding of the interface. (344837)

- When using a vEOS c5-based instance type in Sfe mode in AWS, traffic attempts over it may fail due to duplicate packets getting sent out. (360632)
- vEOS VM deployed on a Intel hyperthread platfrom can cause performance drop. The workaround is to disable hyperthreading on the server if possible. (364731)
- Sometimes when vEOS is launched on Azure, accelerated networking interface don't show up in vEOS when instance is launced with multiple interfaces. The workaround is try relaunching a new instance. (372192)

Known Software Caveats in 4.22.0F

Accelerated System Update

- When performing ASU2 on an MLAG setup with VXLAN configuration, ASU2 can result in a fatal error resulting in a cold boot being performed resulting in large traffic loss. (245555)
  Series Affected: DCS-7050, DCS-7060X and DCS-7060X2 Series
- Performing reload on MLAG primary causes traffic outage for VXLAN indirect routing traffic. (340320)
- SSU with Splunk enabled may result in extended boot times, which could lead to protocol timeouts or switch cold rebooting. (343629)
- When running 'reload fast-boot' or 'reload hitless', ports with optical transceivers may flap, causing traffic disruption (348172)
  SKUs Affected: DCS-7050SX-128-F, DCS-7050SX-128-R, DCS-7050SX-128-SSD-F, DCS-7050SX-128-SSD-R, DCS-7050SX2-128-F, DCS-7050SX2-128-R, DCS-7050TX-48-F, DCS-7050TX-48-R, DCS-7050TX-64-F, DCS-7050TX-64-R, DCS-7050TX-72Q-F and DCS-7050TX-72Q-R
- SandFap agent may crash during SSU causing SSU to be aborted. The system will undergo cold reboot. There are no workarounds for this issue (374829)
  SKUs Affected: DCS-7280SR-48C6-F, DCS-7280SR-48C6-M-F, DCS-7280SR-48C6-M-R and DCS-7280SR-48C6-R
- Some services may not run properly during SSU reload causing LACP links to flap and ARPs to timeout. There are no workarounds for this issue. (381014)

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)
- MSS revokes redirect rules from a MLAG peer when the last member interface on any of the port channels connected to the service device goes down. This results in redirected traffic traversing that MLAG peer being dropped.

  To work around this failure, use the interface mapping from the MSS dynamic device-set configuration to statically map service device interfaces to their corresponding switch port channels. (230653)
- The MacroSegmentation Service (MSS), working with L2 transparent firewalls, requires the firewall interface be statically identified on CVX via configuration commands (as opposed to using LLDP to detect them dynamically). This prevents issues such as traffic loss in the event that a member port of a aggregated link goes down. (275384)
- In Macro-Segmentation Service (MSS) deployment, DirectFlow rules for the intercepted hosts seen behind an MLAG VTEP might not be installed on a MLAG peer after it is reloaded. This situation will be resolved as soon as the reloaded peer starts receiving traffic from that set of intercepted hosts.

  An alternative workaround is to remove and readd the affected IP addresses from the firewall policy (283229)
- If there are large number of intercept hosts in L2 policies in the order of 10k, then shutting down management cvx feature may not cleanup all DirectFlow rules (285821)
- The Macro-Segmentation Service (MSS) fails to install policy rules with pre-defined applications correctly. To workaround this issue, manually define the application in the firewall policy rule. (362682)
- In Macro-Segmentation Layer 3 deployment, restarting MssClient agent might not redirect routed traffic destined to an intercepted host to an external firewall.

  As a workaround, restarting StrataL3 agent should resolve this issue. (363494)
- The Macro-Segmentation Service (MSS) may not handle a custom session timeout in the L4 service definition. To workaround the issue, use a default session timeout. (367896)

- Restarting the MssPolicyMonitor (MssPM) agent can result in a 2-5 second disruption of traffic when the Macro-Segmentation Service is enabled. (369347)
- The Macro-Segmentation Service (MSS) may fail to re-establish communication with a firewall if the host route on CVX is modified.

  To workaround this, toggle the device-set state via CLI (to 'shutdown', followed by 'active'). (372827)
- The Macro-Segmentation Service (MSS) can fail to remove policies on a failure to communicate with the firewall (when configured to bypass the firewall on failure).

  The workaround is to restart the MssPM agent. (373012)
- The Macro-Segmentation Service (MSS) interfacing with a firewall running in L2 transparent mode does not enforce a firewall policy that requires intercepting traffic on a VNI for the first time while one of the peers of the MLAG Service VTEP is being reloaded. MSS starts intercepting that set of hosts when the peer reload completes. (373587)
- The Arista OpenStack ML2 driver fails to sync with CVX for set of VMs is there is a failure in creating a VM, causing CVX and OpenStack to get out of sync (376978)
- The Arista OpenStack agent incorrectly removes the provisioned network when an instance using the network is shut down in a deployment where hierarchical port binding is in use. (377721)

General

- The SuperServer agent can restart unexpectedly if tcpdump sessions are configured and a link being captured on by one of the sessions flaps. (60515)
- When a policy map of type PBR is attached to an interface that is either in down state or a switchport (or both), and if the policy map is too large to fit into available hardware resources, the CLI will not report the error and roll back the configuration. Later, when the interface becomes an operational routed interface, the policy map will not be programmed into hardware. However, "show policy-map type pbr" shows the policy applied to the interface.

  To fix the discrepancy, remove extra rules from the policy map to make it fit into the hardware. (89931)
- When routing is disabled, ARP control packets share the CPU queue with other CPU bound packets. This could lead to MLAG disassociation and network

outages.

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

- The OpenFlow agent fails to return counters associated with a flow, when queried by the OpenFlow controller using protocol version 1.0 (132812)

- On switches using the tg3 driver for the management interface a kernel panic can happen resulting in a switch reload. (136944)
  Series Affected: 7368 Series
  SKUs Affected: DCS-7010T-48-DC-F, DCS-7010T-48-DC-R, DCS-7010T-48-F, DCS-7010T-48-R, DCS-7020SR-24C2-F, DCS-7020SR-24C2-R, DCS-7020SR-32C2-F, DCS-7020TR-48-F, DCS-7020TR-48-R, DCS-7020TRA-48-F, DCS-7020TRA-48-R, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050QX-32-F, DCS-7050QX-32-R, DCS-7050QX-32-SSD-F, DCS-7050QX-32-SSD-R, DCS-7050QX-32S-F, DCS-7050QX-32S-R, DCS-7050QX-32S-SSD-F, DCS-7050QX-32S-SSD-R, DCS-7050QX2-32S-F, DCS-7050QX2-32S-R, DCS-7050SX-64-F, DCS-7050SX-64-R, DCS-7050SX-64-SSD-F, DCS-7050SX-64-SSD-R, DCS-7050SX-72-F, DCS-7050SX-72-R, DCS-7050SX-72-SSD-F, DCS-7050SX-72-SSD-R, DCS-7050SX-72Q-F, DCS-7050SX-72Q-R, DCS-7050SX-96-F, DCS-7050SX-96-R, DCS-7050SX-96-SSD-F, DCS-7050SX-96-SSD-R, DCS-7050SX2-128-F, DCS-7050SX2-128-R, DCS-7050SX2-72Q-F, DCS-7050SX2-72Q-R, DCS-7050SX3-48YC12-F, DCS-7050TX-48-F, DCS-7050TX-48-R, DCS-7050TX-48-SSD-F, DCS-7050TX-48-SSD-R, DCS-7050TX-64-F, DCS-7050TX-64-R, DCS-7050TX-64-SSD-F, DCS-7050TX-64-SSD-R, DCS-7050TX-72-F, DCS-7050TX-72-R, DCS-7050TX-72-SSD-F, DCS-7050TX-72-SSD-R, DCS-7050TX-72Q-F, DCS-7050TX-72Q-R, DCS-7050TX-96-F, DCS-7050TX-96-R, DCS-7050TX-96-SSD-F, DCS-7050TX-96-SSD-R, DCS-7050TX2-128-F, DCS-7050TX2-128-R, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-F, DCS-7060SX2-48YC6-R, DCS-7150S-52-CL-F, DCS-7150S-52-CL-R, DCS-7150S-52-CL-SSD-F, DCS-7150S-52-CL-SSD-R, DCS-7160-32CQ-F, DCS-7160-32CQ-R, DCS-7170-32C-F, DCS-7170-32C-M-F, DCS-7170-32C-M-R, DCS-7170-32C-R, DCS-7170-32CD-F, DCS-7170-32CD-R, DCS-7170-64C-F, DCS-7170-64C-M-F, DCS-7170-64C-M-R, DCS-7170-64C-R, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7280CR2A-30-F, DCS-7280CR2K-30-F, DCS-7280CR2M-30-F, DCS-7280QR-C36-F, DCS-7280QR-C36-M-F, DCS-7280QR-C36-M-R, DCS-7280QR-C36-R, DCS-7280QRA-C36S-F, DCS-7280QRA-C36S-M-F, DCS-7280QRA-C36S-M-R, DCS-7280QRA-C36S-R, DCS-7280SE-64-F, DCS-7280SE-64-R, DCS-7280SE-68-F, DCS-7280SE-68-R, DCS-7280SR-48C6-F, DCS-7280SR-48C6-M-F, DCS-7280SR-48C6-M-R, DCS-7280SR-48C6-R, DCS-7280SR2-48YC6-F, DCS-7280SR2-48YC6-M-F, DCS-7280SR2-48YC6-M-R, DCS-7280SR2-48YC6-R, DCS-7280SR2A-48YC6-F, DCS-7280SR2A-48YC6-M-F, DCS-7280SR2A-48YC6-M-R, DCS-7280SR2A-48YC6-R, DCS-7280SR2K-48C6-M-F, DCS-7280SR2K-48C6-M-R, DCS-7280SRA-48C6-F, DCS-7280SRA-48C6-M-F, DCS-7280SRA-48C6-M-R, DCS-7280SRA-48C6-R, DCS-7280SRAM-48C6-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F, DCS-7280SRM-40CX2-R, DCS-7280TR-48C6-F, DCS-7280TR-48C6-M-F, DCS-7280TR-48C6-M-R, DCS-7280TR-48C6-R, DCS-7280TRA-48C6-F,

DCS-7280TRA-48C6-M-F, DCS-7280TRA-48C6-M-R, DCS-7280TRA-48C6-R, DCS-7304, DCS-7308, DCS-7504N, DCS-7508N, DCS-7512N and DCS-7516N

- For interfaces undergoing maintenance by virtue of being part of a user configured/builtin unit being put under maintenance, "show maintenance" would not reflect the correct maintenance status of the interface. (141942)

- If a file system is full, Launcher can crash when it tries to start a new agent due to a configuation change. (173902)

- When no on-logging event-handler is configured, the EventMgr agent keeps the syslog file open. As a result, when logrotate eventually attempts to clean up old logs, the space associated with this file is not actually cleaned up, causing space to be wasted under /var/log/. Only one deleted message file is kept open when more logotate operations are performed.

  The workaround is to configure an on-logging event handler and then unconfigure it. Alternatively the EventMgr agent can be manually restarted. (177553)

- The system may power cycle unnecessarily in the event of a correctable error on the system FPGA.

  The workaround is to configure the system with 'hardware fpga error action log'. (201445)
  Series Affected: DCS-7050TX Series
  SKUs Affected: DCS-7050SX-128-F, DCS-7050SX-128-R, DCS-7050SX-128-SSD-F, DCS-7050SX-128-SSD-R, DCS-7050SX-64-F, DCS-7050SX-64-R, DCS-7050SX-72-F, DCS-7050SX-72-R, DCS-7050SX-72-SSD-F, DCS-7050SX-72-SSD-R, DCS-7050SX-96-F, DCS-7050SX-96-R, DCS-7050SX2-128-F, DCS-7050SX2-128-R, DCS-7050SX2-72Q-F, DCS-7050SX2-72Q-R, DCS-7050TX-128-F, DCS-7050TX-128-R, DCS-7050TX-128-SSD-F, DCS-7050TX-128-SSD-R, DCS-7050TX-48-F, DCS-7050TX-48-R, DCS-7050TX-48-SSD-F, DCS-7050TX-48-SSD-R, DCS-7050TX-64-F, DCS-7050TX-64-R, DCS-7050TX-64-SSD-F, DCS-7050TX-64-SSD-R, DCS-7050TX-72-F, DCS-7050TX-72-R, DCS-7050TX-72-SSD-F, DCS-7050TX-72-SSD-R, DCS-7050TX-96-F, DCS-7050TX-96-R, DCS-7050TX2-128-F, DCS-7050TX2-128-R, DCS-7160-32CQ-F, DCS-7160-32CQ-R, DCS-7250QX-64-F, DCS-7250QX-64-R, DCS-7280CR-48-F, DCS-7280QR-C36-F, DCS-7280QR-C36-M-F, DCS-7280QR-C36-M-R, DCS-7280QR-C36-R, DCS-7280SE-64-F, DCS-7280SE-64-R, DCS-7280SE-68-F, DCS-7280SE-68-R, DCS-7280SE-72-F, DCS-7280SE-72-R, DCS-7280SR-48C6-F, DCS-7280SR-48C6-M-F, DCS-7280SR-48C6-M-R, DCS-7280SR-48C6-R, DCS-7280TR-48C6-F, DCS-7280TR-48C6-M-F, DCS-7280TR-48C6-M-R and DCS-7280TR-48C6-R

- Enabling the Macro-segmentation service can result in duplicate packet delivery of broadcast, multicast or unknown unicast traffic sourced from the host whose traffic is intercepted. Additionally, one copy can be delivered to the intercepted host that is the source of the traffic. (212480)

- When configuration of a service-policy of type pdp on an interface fails due to lack of hardware resources, the interface may not be protected by any PDP policy. The workaround is to remove the configuration of the PDP service-policy and to configure the PDP service-policy default-pdp-policy on the interface. (216365)
- A custom PDP policy configured using config-session/CVP, may not be applied to the interfaces. Detaching and reattaching the policy through CLI, should solve the issue. (222185)
- When one or more neighbors is moved to a new update-group as a result of neighbor outbound policy change, we might transiently send withdrawals/ incorrect delta advertisements.

  This is eventually corrected when policy evaluation of new update-group eventually completes. (226201)
- With EVPN IRB routing service enabled, it might result in creating a 0.0.0.0/32 host routes in v4 address family for the transit VRF. (237887)
- The MacroSegmentation Service (MSS) feature doesn't work in conjunction with the VARP VTEP IP configuration required for VXLAN routing where a subset of VTEPs are L2 VTEPs. (263532)
- Support for VRRPs in OpenConfig not fully implemented (268736)
- OpenConfig fails because of an invalid type in ISIS with the log message 'expected attribute type nominal.U8, got uint8' (293256)
- The ZeroTouch agent does not properly handle 301 HTTP return codes when downloading the bootstrap script. This can lead to the device leaving ZTP mode with a default startup-config. (342098)
- Under certain conditions where a firewall policy containing rules tagged for the Macro-Segmentation Service (MSS) are modified by moving the destination IP addresses to the source IP field in the rule, MSS may fail to redirect that traffic to the firewall.

  To workaround the problem, please restart MSS. (346370)
- When a MLAG peer (configured as a service VTEP for MSS) reloads, MSS re-installs intercept flows which may cause traffic to drop for a short duration. (349254)
- If the L3 forwarding agent is restarted when VXLAN MAC addresses are present and the route to remote VTEPs are down, then the L2 forwarding might unexpectedly restart in rare cases. (355110)
- 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

- MSS uses DirectFlow flows to match and redirect/permit/drop traffic. For IP flows, where the ethertype is unspecified, it also matches on ARP traffic.

  The workaround here is to add a high priority flow that matches ARP traffic and is forwarded normally. The configuration is:
  (config-directflow)# flow treat-ARP-normally
  (config-directflow-treat-ARP-normally)# match ethertype ARP
  (config-directflow-treat-ARP-normally)# priority 65534
  (config-directflow-treat-ARP-normally)# exit (365135)

- PTP agent will not send unicast PTP event messages using the G.8275.2 PTP profile if the PTP peer is in a remote network. (365913)

- DirectFlow may be unable to reprogram flows with a nexthop after a Strata restart. (366231)

- CVE-2019-9636: Python 2.7.x through 2.7.16 and 3.x through 3.7.2 is affected by: Improper Handling of Unicode Encoding (with an incorrect netloc) during NFKC normalization. The impact is: Information disclosure (credentials, cookies, etc. that are cached against a given hostname). The components are: urllib.parse.urlsplit, urllib.parse.urlparse. The attack vector is: A specially crafted URL could be incorrectly parsed to locate cookies or authentication data and send that information to a different host than when parsed correctly. (367678)

- When an interface with a many directly connected IPv6 routes configured becomes suddenly inactive, such as due to peer reset, BFD sessions on other interfaces may flap. (367872)

- When using PTP with the G.8275.2 PTP profile, enabling PTP using "ptp enable" on a Port-Channel is unsupported. (372185)

- LDP agent may restart when churning a large number of FECs. (387077)

Layer 1

- Speed change on interfaces with "speed-misconfigured" errdisable state after hitless reload causes forwarding agent restart. This leads to the flap of all connected interfaces. The workaround is to recover all "speed-misconfigured" interfaces by configuring speeds of affected /2, /3 and /4 ports to match the speed of the corresponding /1 ports before hitless reload. (352115)
  Series Affected: DCS-7060X and DCS-7060X2 Series

Layer 2

- Configuring VLANs within a config session that are currently in use as internal VLANs may fail with the message that internal VLANs cannot be created, even if the VLAN would no longer be used as an internal VLAN once the config session is applied. (132550)
- The configured MAC security license may not take effort after a system reboot.

  The workaround is to reconfigure the license after the reboot. (147835)
- MACsec may get disabled on some ports if a "configure replace" operation is done where the MAC security profile applied to an interfaces is switched to another profile.

  The workaround is to remove the MAC security profile from the interface and reapply. (187897)
- If Ebra agent restarts on a VTEP that is running EVPN control plane with route type 5, then it can result in VxLAN traffic loss. (276063)
- MACsec delay protection feature is not working. (282186)
- If "port-channel min-links" is configured on both LACP partners for the port-channel between them, the port-channel might not link up. The workaround is to configure min-links on only one LACP partner. (347483)
- If VXLAN routing is confgured on a MLAG, the VxlanSwFwd agent can crash when there UDP traffic received on port 51023 current used by VxlanSwFwd to synchronize ARP entries. (364633)
- When Macro-Segmentation Service (MSS) for L3 Firewalls is enabled and the intercepted mac address aging time expires, offload flows may not function correctly. The workaround is to disable mac aging with the configuration "mac address-table aging-time 0" (365007)
  Series Affected: DCS-7050X3 Series

Layer 3

- OSPF authentication passwords are displayed in clear text in 'show tech-suuport ribd runnning-config' output. (36313)
- If a loopback interface is only configured with a link-local IPv6 address, OSPFv3 will not advertise the address. (131369)
- When the MTU of routed port is set to a higher value followed immediately by a similar change on the subinterfaces configured on that routed port, the MTU may not get set correctly on the sub-interfaces.

The workaround is to verify the MTU configuration is effective using the "show ip interface" command before configuring the MTU on the subinterface. If the subinterface ever gets into a state where the configured MTU is different from the effective MTU, perform a "shutdown" then a "no shutdown" command on the affected subinterfaces to recover. (141389)

- FHRP state may start flapping between Backup and Master and Fhrp Agent may restart while dealing with large number of Virtual IP addresses. This is seen at a scale of 20K virtual IPs or more. (182399)
- MRIB BGP routes are not installed when a next hop is recursively resolved to the Null0 interface, which happens in iBGP scenarios. (194181)
- BGP LU path next-hop can be potentially resolved over an IGP protocol route which could lead to undesired traffic forwarding behavior if the next-hop IGP router is not in the label switched forwarding path (LSP).

  The issue is resolved by restricting next-hop resolution to a) label switched paths (LSPs) where the BGP next-hop corresponds to a Tunnel end-point in the Tunnel RIB or b) directly connected routes in the RIB (215996)
- On an IS-IS L1-L2 router, when an interface transitions from one level to another, accompanied by the interface going down, the Rib/IS-IS agent may restart. (235375)
- Shutting down all the nexthop interfaces belonging to a default prefix resilient ECMP will cause the L3 Agent to restart. (255064)
- The Ldp agent may restart if configured with a large number of MPLS ECMP routes with many next hops. (273767)
- With scaled VRF config, in the order of 1K VRFs, if the Arp agent restarts, it may fail to come up. (275514)
- In certain scale situations upon a sysdb restart, arp entries can be learned in
  kernel but not show up under 'show arp'. When this happens, the work around is to
  restart the Arp agent. (275749)
- Configuring 'max-metric router-lsa' on an OSPF router may cause summary LSAs to be generated with maximum metric. (276629)
- Removing LDP mappings in a large scale LDP configuration may result in LDP neighbor sessions flaps. (283972)
- Auto discovery of multiple IPv6 link-local BGP neighbors over a single interface does not work and will result in only one of the neighbors transitioning to established state. (289875)

- If a sub-interface creation/deletion overlaps with it's VRF creation/ deletion then the sub-interface may not be deleted properly in kernel. This prevents creation of same sub-interface in kernel again.

  Workaround is to restart Ebra agent. (337707)

- OSPF ABR does not withdraw some translated NSSA LSAs in presence of overlapping NSSA prefixes (350804)

- When OSPF has default-information originate always configured and a static default route is redistributed to OSPF, LSA changes in the network can cause unexpected updates to the external LSA with the default route information. This may cause too many retransmissions and bring down the adjacency. The workarounds are to stop the route changes, or unconfigure default-information originate always. (354145)

- If DirectFlow is shutdown / no shutdown after the Macro-Segmentation Service (MSS) for L3 Firewalls is enabled, the hardware programming for redirection of flows may not get removed resulting in incorrect forwarding. There is no workaround for this issue. (363479)
  Series Affected: DCS-7050, DCS-7050T, DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X3, DCS-7250X, DCS-7260X, DCS-7260X3, DCS-7300 and DCS-7300X Series

- If one or more L3 routed interfaces that were initially assigned to a VRF are operationally DOWN, it may cause the KernelFib agent (process managing the SW routing table) to reprogram the kernel routes in that VRF for destinations that have a special forwarding adjacency (i.e. routes pointing to tunnels, nexthop-groups adjacency or when routing from one IP address-family to another (RFC5549)). This may result in a temporary blackhole of packets forwarded to these destinations in software.

  As a work around for this issue, one can configure a VLAN and then create an SVI in that VRF. The SVI will need to have an ip address assigned and no autostate configured. (364393)

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

  A workaround is to shutdown the MSS service on CVX, enable DirectFlow on all supported switches and re-enable the MSS service. (366022)
  Series Affected: DCS-7050, DCS-7050T, DCS-7050TX, DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X3, DCS-7250X, DCS-7260X, DCS-7260X3, DCS-7300 and DCS-7300X Series

- Directflow match on "input interface Recirc-Channel" is not supported. On Macro-Segmentation Service (MSS) deployment with L3 firewalls, this may result in traffic looping on service VTEP for the traffic returning from

Firewall. (369939)
  Series Affected: DCS-7050, DCS-7050T, DCS-7050TX, DCS-7050X, DCS-7060X, DCS-7060X2 and DCS-7060X3 Series

- As-external OSPFv3 routes may not be installed correctly, if there is more than one path to reach the ASBR router (375929)

Multi-agent

- The forwarding address in OSPF/OSPFv3 external LSAs will not be updated correctly when the route being redistributed is a BGP route whose FEC underwent inplace adjacency replace. (218143)
- If OSPF is configured with max-metric wait-for-bgp and IS-IS is configured with set-overload-bit, then wait-for-bgp might advertise BGP prefixes before BGP convergence.

  To ensure that IGPs wait for BGP convergence before advertising prefixes, configure BGP graceful-restart or wait-for-convergence along with IGP wait-for-bgp configuration. (225972)
- During Smart System Upgrade, there may be unexpected traffic loss when only IGPs are configured. (239982)
- With EVPN IRB routing service enabled, IPv6 hosts may result in creating unwanted routes in IPv4 VRF on VTEPs. (272726)
- Ospf summary-address might not be advertised if all the contributors of summary-address are removed and added in quick succession.

  To workaround, un-configure and reconfigure the Ospf summary address. (346166)
- There may be some traffic loss on ASU or SSO restart at high scale (i.e with a large number of next-hops). (362995)
- The BGP may crash if any VLAN-Aware bundle MAC-VRF is configured under MPLS encapsulation with VLAN/ETID where ETID is set to any value > 4094. To avoid the crash, do not exceed 4094 for ETID configuration. (380138)
- BGP agent may restart unexpectedly when dynamic prefix-list is configured in a VRF which has BGP imported VPN routes in FIB (381120)

Rib

- Rib agent may unexpectedly restart when executing show commands for routing protocols. (259090)
- In scaled BGP deployments, BGP convergence is degraded (379127)

- An eventual Rib crash will occur over time due to an internal resource re-use. Amount of time can vary based on route churn in the system. (379398)

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

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

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

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

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

- Policy cache size may grow unboundedly, leading to unexpected agent restart (388122)

MLAG

- If non-MLAG reload-delay is configured to be lower than the MLAG reload-delay when LACP standby is used, traffic entering the non-MLAG interfaces destined towards hosts on the MLAG interfaces will be lost during the LACP standby period. (83330)

- When one of the MLAG peers goes through a stateful switchover (SSO), the other MLAG peer might lose all the mac addresses synced via the peer-to-peer mount and this could result in flooding until the mac addresses are restored or relearnt. (247396)

- On Mlag setup with dual primary detection enabled, traffic loss might be seen when recovering from dual primary condition. (297307)

- On forwarding agent restart on MLAG system, STP might not process BPDUs received on the MLAG secondary's interfaces. This affects the port roles, states, counters and can disrupt traffic. Workaround is to flap interfaces on MLAG secondary. (382765)

Multicast

- When multicast routing is disabled, changing the state of a port in a non-default VRF will not correctly update the port in L2 multicast floodset. As a result, packets destined to the port's mulicast group will not reach it.

  The workaround is to enable multicast routing. (142942)

- When the route to RP is ECMP, MSDP SA messages are accepted from only one of the possible peers.

  To accept the SA message from a different MSDP peer use "ip msdp default-peer <peer_id> <RP>". (190306)

- During KernelMfib restart, the agent might not program or update a multicast route in the kernel when there are many VRFs. This might cause problems if the restart takes place during SPT switchover. The multicast route might not switch to SPT. (217771)

- In high scale scenarios MSDP agent might crash. (259987)

- An S,G route may not properly cleanup previously sent S,G joins. This would prevent S,Gs from being deleted in the upstream router. A work around is to flap the interface in which the join is sent on. (291756)

- If large number of IGMP reports are received in a very short interval, some of them might get dropped resulting in the multicast route not being created. The routes will be re-learnt on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIgmp", consider increasing the bandwidth for this Copp class. (356675)

- A S,G route, in a multicast vxlan environment using the underlay solution, can loop packets when there is a source and receiver on the same vlan. This can happen during a mlag reload of the DR or when changing DR using the DR priority configuration. (373268)

- With a scale of few thousand multicast route, and multicast traffic stopped and restarted after a while (> 5mins) some routes might get into a fastdrop state causing loss of multicast traffic. Clearing the affected multicast route can be used as a work around. (380214)

DCS-7150 Series

- If a multicast boundary is applied on an SVI which also has IGMP snooping enabled, IGMP snooping will not take effect. (138610)

DCS-7170 Series

- In an MLAG setup, if routing on peer-mac is enabled, then if a switch receives an L2 packet with destination mac address of the frame with peer-mac, then the packet will not get forwarded to the peer.

  The workaround is to disable the feature by running the configuration command "no ip virtual-router mac-address mlag-peer" (352873)
  Series Affected: DCS-7170 Series

- If lag configuration changes continuously such that a port is moved from being member of one lag to the the other lag, the lag from which the member port moved and the the lag to which the member port moved, both would have undefined behavior. (371119)

- After high routing churn, routed traffic may be misforwarded to stale routes. (385787)
  Series Affected: DCS-7170 Series

DCS-7300 and DCS-7500 Series

- On the DCS-7300X series, an SEU on a supervisor SCD will cause a power cycle of the supervisor and possibly a misleading reload cause of "The system rebooted due to a watchdog". The reload cause should be "The system rebooted due to an SEU error on the SCD". (104962)
  SKUs Affected: 7300-SUP and 7300-SUP-D

- When many IPv6 routes are configured and the BFD protocol in use, BFD sessions may flap when a linecard is removed or shut down. (279129)

- ManagementActive Agent crash repeatedly in rare corner case and Management0 interface is not reachable in
  Modular systems. Doing switchover will address the issue. (369565)

DCS-7020RA, DCS-7160, DCS-7170, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards Series

- The AlgoMatch forwarding agent will restart when config replace is used to configure access-lists with large number of rules. (240540)

- Following a physical error on link between Jericho and Algomatch FPGA the SandHalo agent may restart, causing ACLs to be removed and reinstalled. (283917)

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

- A linecard may be detected erroneously as removed and reinserted after an SSO switchover. (34907)
- When in MLAG mode without ip routing enabled, under some control plane oversubscription scenarios the MLAG peer link might come down. The workaround is to either enable IP routing or have a static ARP entry on the MLAG peers for each other. (90247)
- An IPv6oGRE packet terminated using decap-group with a host route destination IPv6 address will be dropped.

  The workaround is to disable exact match host routes using "no platform sand ipv6 host-route exact-match". (95873)
- When a SSO switchover is triggered prior to all the agents on the standby supervisor being ready, this can result in PCIE bus errors and syslogging of the HARDWARE-3-SUPERVISOR_PCIE_SWITCH_BUS_ERROR message for the affected modules. This will result in traffic loss during switchover, and the switchover will be non-stateful. (113038)
  Series Affected: DCS-7500E Series
- On DCS-7280E, DCS-7500E series, DCS-7050X, DCS-7250X, and DCS-7300 series switches, the vlan tag on IGMP messages(report, query, leave) generated by the switch is not translated in the egress direction by the switchport vlan mapping command. (113104)
- Some packet loss may be observed when IP multicast traffic has high fanout on a single physical port (e.g. trunk port with many receivers for the same multicast group on different VLANs), even when the total egress packet rate is well below linerate. The default egress buffer settings may not perform well under this scenario. This issue can be resolved by adding "platform arad buffering egress port multicast max-queues-full 1" to the switch config and rebooting the switch. (115343)
- MAC mirroring ACLs do not match IPv6 packets. (137537)
- When MLAG peer mac routing feature is enabled on both MLAG peers, the IPv6 neighbor may not be resolved when source ipv6 address in solicited neighbor is link local address. (165067)
- When a large number of PBR service policies are configured and the SandAcl agent restarts a number of times, the SandAcl agent may continuously restart.

System restart will remedy the situation. (166081)

- If an ACL applied on a Port-Channel subinterface is removed while the parent Port-Channel is shutdown and then the parent Port-Channel is unshut, the SandAcl agent can restart unexpectedly. (177127)
- Mirrored traffic is not subject to egress security ACLs which are applied to mirroring destination ports. (190637)
- PBR is not supported along with IPv6 in IPv6 packets. (193914)
- If an ACL is attempted to be applied on a subinterface, whose parent port is not a routed port, the CLI will timeout. (229537)
- PBR policies applied on interfaces do not work if the interface has some subinterfaces that also have PBR policies applied. (243286)
- VTEP cannot learn a remote host's MAC address from VXLAN encapsulated IPv6 Neighbor Advertisement packets received from that host. This might be an issue in a pure IPv6 depolyment where the VTEP does not receive any bridged traffic from the host. In such a scenario, the traffic toward the host will be flooded in the VLAN.

  The work around for this is to force Neighbor Solicitation (NS) packets from the host for its gateway or send bridging data traffic from the host. (246676)
- ACL/PBR/QOS will not take effect on ports with VLAN Translation configured. (246982)
- A TCAM profile not compatible with all linecards in a system will not be installed on any linecard in that system. However, hardware tcam profile status may show errornously that the profile is installed on compatible linecards.

  Removing the incompatible linecard(s) or use a TCAM profile that is supported by all linecards in the system. (247842)
- VXLAN bridging traffic sent out of LAG interfaces may be temporarily dropped after an SSO event. (251398)
  Series Affected: DCS-7500E and DCS-7500R Series
- When SflowAccel is configured, removing the AlgoMatch linecard may cause packet drops on flows in the non-AlgoMatch linecards. (251483)
- When linecards are powered on/off, there may be a traffic loss for a short duration on the interfaces where vlan mapping feature is configured. (254476)

- 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 DirectFlow flow is created with an action to drop the outer tag, the packet is still leaving the interface tagged. (264336)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

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

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

- Unexpected DSCP rewrite may occur on a port if the port has an internal VLAN that was previously assigned to a port which had DSCP rewrite map applied.

  The workaround is to apply and remove a dummy map on that interface using "[no]qos rewrite traffic-class to dscp". (274484)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

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

- With TCP MSS Clamping feature configured, TCP SYN packets might get dropped under high CPU usage conditions if the incoming TCP SYN packet rate is high. (279370)

- PIM multicast traffic may experience a brief interruption when one or more linecards on the same modular chassis enters or leaves TAP Aggregation mode. (279972)

- After SSO, LDP Pseudowires can misforward traffic for upwards of 15 minutes. Eventually forwarding will be restored automatically. (281521)

- When the primary and fallback policy are both changed in the same config session, traffic may stop matching on either primary or fallback policy. (287146)

- When primary and fallback policies are both changed in the same config session, some packets will not be subject to either policy for a brief period during the configuration change. (287154)

- On a modular chassis with some linecards in TAP Aggregation mode, restarting the SandFap agents for these linecards could result in a brief

window where traffic entering into Tap or Tool ports may be flooded to all switch ports in the system. (288350)

- When operating in hardware TCAM profiles that support "feature pbr ip/ipv6/mpls", interface speed changes, Linecard removal or insertion, and SandFap agent restarts might cause transient traffic loss for PBR matching flows for Port-Channel sub-interfaces. (294271)
- Power cycling a 7500E linecard in a chassis with 7500E fabric modules can cause packet loss on 7500R linecards. (295825)
- When new flow-spec rules are programmed, existing rules will be briefly lifted and allow matching traffic to pass. (297534)
- When both port-channel load balance for multicast and selective ingress replication is configured, groups that get ingress replicated incorrectly replicate packet to each member of port-channels that are member of the group.
  As a workaround one of the features has to be disabled. (344650)
- On DCS-7500RM-36CQ-LC linecards, rapidly enabling and disabling FEC on interfaces with 100GBASE-PSM4 transceivers can cause the interfaces to not come up, flap, or errdisable.

  Workaround is to wait 20 seconds before applying another FEC configuration. (354921)
- Switching from non-flex route profile to flex-route profile with 'ip hardware fib optimize prefixes profile internet' in configuration, the SandL3Unicast agent can restart unexpectedly twice with transient traffic disruption. (366472)
- When IPv4 or IPv6 flowspec rules exceed the available TCAM capacity, the rules for the other protocol may not be installed even if there is space in the TCAM for it. (369374)
- After SandFap slice agent restarts with PTP operational and using peer-to-peer delay mechanism, all PTP Peer delay messages ingressing PTP enabled ports will no longer reach the PTP agent. The workaround is to toggle PTP mode from disabled back to boundary. (370878)
- TCAM bank reservation may result in extra banks allocated to the feature after a tcam profile change. (372760)
- Traffic coming into subinterfaces are subject to the flowspec rules applied on the parent interface's VRF. (380533)
- Ethernet over MPLS packets being terminated and decapped on a LDP pseudowire will be hashed as if the inner header is IPv4 rather than Ethernet (382763)

DCS-7280E and DCS-7500E Series

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

  The workaround is to always enable autonegotiation on these ports. (70667)
  SKUs Affected: 7500E-72S-LC

- IPv4 egress RACLs will not filter multicast traffic in shared mode. In unshared mode, multicast traffic is filtered only if ingress replication is enabled. (131660)

- An IPv6 Egress ACL applied to a Port Channel is not applied to a newly added member interface, if the member interface already has the same ACL applied to it. The workaround is to remove the ACL from the interface before adding it to the PortChannel. (141736)

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

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

- L3 forwarding agent may continuously crash if EVPN VXLAN multihoming/VXLAN L2 ECMP is configured on unsupported platforms. (381972)
  Series Affected: DCS-7280SE and DCS-7500E Series

DCS-7280R and DCS-7500R Series

- In MPLS L2EVPN configuration, traffic requiring 3 or more LU labels to be pushed, may not be forwarded correctly. (275834)

- On the 7500R-8CFPX-LC linecard modules, rapid removal and reinsertion of the transceiver may result in the Denali agent restarting.

  Inserting transceiver 10 seconds after removal from the same slot is the workaround. (293594)
  SKUs Affected: 7500R-8CFPX-LC

- If a Linecard is powered off due to detection of a hardware issue, any additional OIR on the card slot will not be able to power on the Linecard. (337169)

- 100GBASE-PSM4 modules may not link up after a reboot. The workaround is to remove and reinsert the
  module. (354176)
  SKUs Affected: 7500RM-36CQ-LC

- Restart of the Evora agent could cause dropping of control packets and CRC errors on the received packets. (367523)
  SKUs Affected: 7500R-8CFPX-LC and 7500RM-36CQ-LC

- With CPU traffic policy configured with high VRF scale, the Acl agent may unexpectedly restart. (379101)

DCS-7020 Series

- When multiple Ipsec tunnel interfaces are shutdown, then "no shut" in a short period of time, routes associated with some tunnel interfaces may not be installed after "no shut". The workaround is to do shut/no-shut again on the individual tunnels where routes not installed. (377984)
  SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

- Ipsec tunnel interfaces may not have link up after SandFap agent restarts. The tunnel state is not consistent in this case.

  The workaround to fix this issue is to restart SandIpsec agent. (378536)
  SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

DCS-7020RA, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards Series

- VXLAN packets that are decapped and bridged are incorrectly subject to the IP ACL applied on the ingress port. (221457)

- The HaloFpga agent will continuously restart if SandFap agent is restarted.

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

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

- Traffic streams matching DirectFlow flows with VLAN or MAC rewrite actions could face brief disruption across a StrataL3 agent restart. (111833)
  Series Affected: DCS-7050 and DCS-7050X Series

- The WRED drop decision obtained from the ECN configuration on tx-queue 0 is also used to determine whether packets should be dropped on UC8. This can result in control packets enqueued on UC8 getting dropped if UC0 is experiencing WRED drops for non ECT packets emqueued on it. A workaround is to use a queue other than tx-queue 0 for applying ECN. (125803)
  Series Affected: DCS-7050 Series

- Toggling the priority flow control (PFC) configuration on an interface carrying traffic may cause all incoming traffic to be dropped. The workaround is to restart the Strata agent. (138511)

- Support for ingress and egress subinterface and VLAN counters. (181633)
  Series Affected: DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7250X, DCS-7260X and DCS-7300 Series

- On the DCS-7250X and DCS-7300X series, the packet flow for vxlan decap +routing into ecmp can be dropped if the ecmp members are in different Linecards. (181934)
  Series Affected: DCS-7250X and DCS-7300X Series

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

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

- On an MLAG pair of switches, with fast-MAC redirect enabled, the Macro-Segmentation Service (MSS) can fail to recover from a MLAG port channel member link flap.

  The workaround is to shut/ no-shut the port channel connected to the intercepted host experiencing traffic loss. (213199)
  Series Affected: DCS-7050X and DCS-7060X Series

- The Strata fabric or linecard agent may restart due to fabric link flaps leading to traffic outage. (225302)
  Series Affected: DCS-7300 and DCS-7320 Series

- Inserting and removing the QSFP+ transceiver will result in forwarding plane agent restart. Workaround is to configure a forced speed on the port before inserting/removing a transceiver. (226581)
  Series Affected: DCS-7050X3 and DCS-7260X3 Series

- DirectFlow flow in table type vfp will not match malformed Arp requests. (233922)

- When traffic is redirected using PBR and when the FIB lookup based on destination IP address indicates ARP miss, the PBR redirected traffic may be rate-policed.

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

- When the utilization of the MAC address table grows close to the full hardware capacity, the StrataL2 agent may encounter an out of memory (OOM) condition in the system, causing the agent to get restarted. (238115)
  Series Affected: DCS-7010, DCS-7010T, DCS-7050, DCS-7050T, DCS-7050TX,

DCS-7050X, DCS-7050X2, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X3, DCS-7260X, DCS-7260X3, DCS-7300, DCS-7300X and DCS-7320 Series

- If nexthop resources are exhausted as indicated by syslog VXLAN_HWHER_NEXTHOP_RESOURCE_FULL, forwarding of BUM traffic to that VTEP will be affected. (246391)
  Series Affected: DCS-7050X and DCS-7060X Series

- When Storm Control is configured on a port channel interface, frequent link flaps on the member interfaces can cause the Strata slice agent to restart.

  The workaround is to configure Storm Control on the member interfaces instead of the parent port channel interface. (248663)

- During an SSO failover, Strata Fabric agent may not detect hardware, resulting in a fatal error reboot with increased traffic outage. (250129)
  Series Affected: DCS-7300 Series

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

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

- SSO can fail with a fatal error if switchover occurs before Strata forwarding plane agent has finished initialization. (254603)
  Series Affected: DCS-7300 Series

- 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)

- The Strata agent can restart in the event a mirror session is updated while an OpenFlow flow exists referencing that particular session. (270713)

- If SSO redundancy mode is configured and a linecard is removed before it is fully initialized, RPR switchover may be performed. (283088)
  Series Affected: DCS-7300X Series

- Configuring interfaces with auto-negotiation speeds may cause continuous forwarding agent restarts. (283312)
  Series Affected: DCS-7260X3 Series

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

- Shutting/Unshutting a MLAG downlink might cause the unknown unicast flooding on MLAG peer-link.
  A possible workaround is to flush the MAC entries in the system (288719)
- Show directflow detail may show incorrect status if the flow has been deleted. (297693)
- Under heavy control plane traffic to/from the CPU, an assertion may be printed in the system logs. This log is harmless. (309393)
- 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 Strata linecard forwarding plane agent unexpectedly restarts during SSO, SSO will not complete causing a reload of the switch. After the reload, protocols like ISIS, OSPF or BGP will re-establish their sessions, which will lead to routing flap. (330842)
  Series Affected: DCS-7300 Series
- Configuring more than five different speeds simultaneously on the system will result into continuous forwarding agent restarts. (339513)
  Series Affected: DCS-7050X3, DCS-7260X and DCS-7260X3 Series
- For OSPF sessions running with BFD, if the BFD timeouts are too low there can be BFD flaps during ASU that may cause OSPF sessions to flap leading to traffic outage greater than expected. (341021)
- If ports with default speed other than 100G are configured for default speed, SSU may have extended outage. Workaround is to explicitly configure speed on these ports using 'speed forced <speed>' command. (342254)
  SKUs Affected: DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-F and DCS-7060SX2-48YC6-R
- When Vxlan configuration is present, any unexpected restart of StrataL3 agent may result in continuous restarts of StrataL3 agent. The workaround is to reload the system to recover from this scenario. (344674)
- Stateful switchover may fail causing cold reload of the switch. There is no workaround for this issue (345279)
- Under certain circumstances, some switch chips or fabric chips may no longer be programmed by platform agents until they are manually reset. (350851)
- Some ports (esp. Et73-Et96 and Et103-Et104) might blackhole traffic or fail to linkup. (358361)

SKUs Affected: DCS-7050TX-128-F, DCS-7050TX-128-R, DCS-7050TX-128-SSD-F and DCS-7050TX-128-SSD-R

- RIB agent restart may result in continuous restart of StrataLag agent when the Vxlan tunnels are resolved through underlay ECMP. The workaround is to "shut" and "no shut" on Vxlan interface. (360915)

- Irrespective of IGMP snooping being enabled or disabled, IGMP packets will not be flooded to remote VTEPs. (368106)
  Series Affected: DCS-7050X2 and DCS-7050X3 Series

- Executing "platform trident forwarding-table partition" in quick succession may result in unrecoverable traffic loss. The work around is to do remove and reissue the command "platform trident forwarding-table partition". (368260)

- When removing security ACL by using "no ip / ipv6 / mac access-list <acl>" on a switch and if more than 1 security ACL is applied of same type IPv4, IPv6, MAC then Strata agent may restart due to heartbeat timeout.

  Workaround is to detach the ACL that needs to be removed from all interfaces and then remove the ACL. (378727)

- If a DirectFlow flow is configured with both 'action output nexthop' and 'action output interface', the Strata slice agent can continuously restart until one of the conflicting actions is removed. (380873)
  Series Affected: DCS-7050X2, DCS-7050X3, DCS-7300 and DCS-7320 Series

DCS-7010T Series

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

- PIM packets which are denied by a user ACL and logged will be processed by the PIM agent (216939)

DCS-7060X and DCS-7300 Series

- If RPR or SSO is configured, Strata fabric agent restarts may cause switchover to occur. (368082)
  SKUs Affected: DCS-7304 and DCS-7308

DCS-7260X3 Series

- Changing interface speed on a range of interfaces on the switch might result in unexpected restart of the forwarding agent. (267755)

- Hitless upgrade might fail if ports are left with default or unsupported speed configuration.

Workaround is to explicitly configure supported speeds on these ports. (285124)

- Hitless upgrade will fail if non-flex ports with 40G transceivers are configured for default speed.

  Workaround is to explicitly configure forced speed on all ports with 40G transceivers. (286870)

- Ports with 1g transceivers will still count towards the limit of 5 speed classes even when disabled and may lead to continuous restart of the Strata forwarding agent. It can cause traffic outage on other ports.

  Workaround is to remove the 1g transceiver. (338741)

7368 Series

- An interface may fail to link with a MAC Local Fault indication.

  The workaround is to run interface configuration command "shutdown" followed by "no shutdown". (341316)
  Series Affected: 7368 Series

DCS-7050 and DCS-7300 Series

- 25G optics are not supported. Inserting one would lead the forwarding agent to crash continuously. (199036)

DCS-7050X and DCS-7300 Series

- Configured as MLAG, link flap of one or more LAG peer-link member ports could cause brief flooding of packets over the peer-link and can also lead to double delivery of packets on the MLAG interfaces. (117870)

- (A1 silicon only) MPLS traffic gets incorrectly prioritized, causing other traffic on the same port to get dropped. (160269)

- In an MLAG configuration, if the StrataL2 agent is restarted, traffic to MLAG interfaces may be double delivered by both switches.
  Couple of Workaround options:
  1. Flapping the Peer link interface would recover the state
  2. Reapply the MLAG configurations
  3. Configure the Peer link interface and Mlag Interface in the same chip so that it uses TCAM for the split horizon (164669)

- Altering the priority of transmit queues of a port will result in some traffic loss on that port. (194460)

- When WRR is configured, 'reload hitless' can fail causing a hitful reload of the switch. (219482)
  Series Affected: DCS-7050X Series

DCS-7050X2 and DCS-7300 Series

- When an interface which is transmitting at linerate is disabled and re-enabled, it might end up in a condition where it is unable to transmit any more packets. Disabling the interface again causes the Strata-FixedSystem agent to restart following which the interface continues to transmit properly. (245501)
- If the switch receives a VLAN-tagged VxLAN encapsulated packet destined to the switch's VTEP IP address, but the packet's destination MAC is not the switch's MAC, such a packet is flooded to all ports in the ingress VLAN including the ingress physical port. (261494)

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

- Inserting an SFP-1G-SX or SFP-1G-T transceiver into an SFP28 port will cause continuous forwarding agent restarts until the transceiver is removed.

  The workaround is to configure the first port in the quad as 10G (either leaving the port empty or using a 10G transceiver). The other three ports in the quad must be configured as either 1G or 10G. (225916)
  SKUs Affected: DCS-7050SX3-48YC12-F
- Egress ACL on SVI for BUM traffic will not match ACL rules for packets being sent across fabric interfaces. (284543)
- Cos rewrite does not take effect for packets that have been vxlan terminated. (286841)
- Gratuitous ARP packets sent from CPU with source MAC as VARP Mac address may get encapsulated twice to remote VTEPs with SIP as local VTEP Ip and VARP Vtep Ip resulting in mac address move for the VARP mac address in remote VTEPs. This may result in momentary traffic loss. The workaround is to increase the GARP advertisement interval with command "ip virtual router mac advertisement-interval <interval>" to reduce the frequency of MAC address moves. This issue only affects configurations that contain ip virtual-router address. (289986)
- A single-event upset (SEU) on a TCAM associated data table can cause traffic loss. The workaround to recover from this is to do an ASIC reset. (367805)
- Ingress ACL and Policy Map rules applied to an interface set might take effect on interfaces outside the set as well. (378187)

Series Affected: DCS-7300X Series
SKUs Affected: DCS-7304 and DCS-7308

- Inserting an SFP-1G-SX or SFP-1G-T transceiver into an SFP28 port will cause continuous forwarding agent restarts until the transceiver is removed.

  The workaround is to configure the first port in the quad as 10G (either leaving the port empty or using a 10G transceiver). 1G transceivers can then be used in the other three ports in the quad. (378303)

- If L3 agent undergoes a restart on the switch, routes may be incorrectly programmed in the forwarding chip. This will cause incorrect routing or traffic loss. There is no workaround for this issue. (378632)

- When multicast routing is enabled on T3 systems, route + flood traffic could result in DMA stuck error in Strata agent and that could result in resetting of the Strata agent.

  There is no known workaround. (378787)

- Priority flow control configuration on the master port will not take effect after speed change is performed. The workaround will be to remove and re-apply the Priority flow control configuration on the interface after speed change. (381163)

DCS-7160 Series

- Under certain race conditions between LAG configuration changes and XpCentral agent restart, the system could get into a state where XpCentral agent continuously undergoes a restart.

  The workaround is to remove all LAG interfaces and re-configure them after XpCentral agent has restarted successfully. (161558)

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

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

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

- In the event of the XpAcl agent restarting, the hardware programming of ACL may not be correct.

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

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

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

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

  The workaround is to reboot the dut. (255957)

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

- XPL3Unicast process may crash when the switch is reloaded if PBR policies with nexthop VRF actions are applied on an interface in the non-default VRF. (262713)

- If the capacity of prefixes/routes that can be accomodated is reached, under some conditions,
  new routes are not programmed even if the the number of prefixes/routes are brought down within
  the capacity.

  Workaround is to reduce the amount of learned routes and restart XpL3Unicast agent. (372207)

- Under some scenarios where there is significant about of CPU bound traffic, especially
  with packets from new hosts whose MAC address need to be learnt,, all traffic bound to
  CPU may be dropped. This also requires that the ASIC ports 60-63 (as seen in "show platform xp mapping port") is configured to be in 10G/25G.

  One possible way to recover is is to reset the chip via "platform xp xp0 reset" in 4.22.0
  and beyond. (380517)

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)

vEOS Router Series

- The OpenFlow agent fails to handle packet-in messages (145001)
- A reboot due to a kernel panic can happen during first 10 mintues of vEOS bootup on Microsoft Azure. The kernel panic occurs because of a task blocking without being scheduled for more than 300 seconds in "D" state. (249618)
- Interface may be allowed to be configured with un-supported speed of NIC using "speed <speed>", which will cause PhyEthtool agent crash. Do not use "speed <speed>" command to configure interface speed that NIC does not support. (264395)
- When running vEOS devices in Sfe mode with interfaces having an MTU of 9214, users may encounter a scenario where not all routes are propagated through the network. The workaround is to have interfaces use an MTU of 9202 or less. (289886)
- When we are running vEOS in DPDK mode (MODE=sfe in /mnt/flash/veos-config) and any interface link status changes, this will cause traffic drop to happen on the system. There is no workaround for the issue. (295198)
- On vEOS Router in Sfe/DPDK mode with many GRE/IPSec tunnels, changes to configuration may cause the time to be applied to the datapath. (349552)
- When loading a vEOS instance in an Azure cloud, it is possible for the device to experience a kernel panic just after bootup, causing the device to reboot a second time. (354411)
- In an EVPN IRB deployment on vEOS, changing the VXLAN source-interface interface may result in EVPN ip-prefix routes not getting installed into one or more VRFs for one or more remote VTEPs. (369816)
- If two vEOS routers are configured to tunnel traffic using IPsec and the path between the two routers has a latency of 20 msec or more, the performance of the TCP sessions as seen by end systems may show degradation versus performance seen in low latency environments. This is most likely to be observed when the vEOS routers are running at full capacity with respect

to their IPsec load. This problem is fixed with the software modification associated with this bug. (378110)

- When multiple IPSec connections are created, some with AES encryption and some with NULL encryption and NULL authentication, bessd will crash when decrypting the IPsec packets encrypted with NULL encryption and NULL authentication. (378590)
- Having a vEOS L3 EVPN VXLAN connection pass over an IPSec tunnel results in packets not being able to be fragmented properly. Since MTU setting is not supported on VXLAN interface the MTU has to be setup on the network to 50 byts less than the MTU of the IPsec tunnel. (378858)
- When removing a Vxlan1 interface from a vEOS device, Sfe agent might crash but it will recover (379861)
- When a number of vrfs are removed via 'no vrf defintion <vrfName>' and then re-added on a vEOS device, it is possible that the Sfe Agent will crash but it will recover. (379908)

Limitations and Restrictions in 4.22.0F

Accelerated System Update

- Aborting 'reload fast-boot' command with Ctrl-C is unsupported (344574)
- Rib agent can restart when the scale of IS-IS adjacencies are huge and multiple IS-IS adjacencies flap and/or multiple prefix-segments are flapped (345657)
- On systems with 4GB of /mnt/flash space, there may be insufficient space to store both an old swi image and a new swi image to perform an SSU upgrade. The workaround is to delete the old swi image from /mnt/flash prior to performing the upgrade. (385364)

CVX

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

  If a switch has more than 128 directly connected TaskTracker nodes with MapReduce Tracer deployed, then the MapReduceTracer Agent can increase CPU utilization significantly. It is recommended to keep the directly connected TaskTracker count to below 128 when MapReduce Tracer feature is deployed. (80403)
- Switches and CVX instances participating in a CVX setup must either all run 4.15.4F or later, or all run images from before 4.15.4F. (146664)

- If a CVX service agent connects to a client switch running a newer software version, the service agent may fail to connect to the client.

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

- The Macro-Segmentation Service (MSS), requires MLAG fast redirection be enabled on the service MLAG TOR pair, when configured to work with a L2 transparent FW. This configuration minimizes traffic loss when any individual interface in the MLAG port channel goes down. (268291)

- The MssClient agent, can restart unexpectedly, when the MacroSegmentation Service, deployed with a L2-transparent firewall and is configured to intercept more than 10,000 end-points on a single switch (a number that's far greater than the total number of intercepts possible on the switch). (283221)

- Mss does not support traffic interception for the hosts connected to Service VTEP. (354754)

- The MacroSegmentation Service (MSS) allows users to configure multiple tags on CVX that can be used in firewall policy definition, to identify policies MSS should work on. It is however not possible to delete one of the multiple tags configured.

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

- In Macro-Segmentation Service (MSS) for L3 Firewalls , Virtual SVI IP addresses cannot be treated as intercepted hosts and thereby cannot be added to any redirect policy. However, these IPs can still be part of offload policies, even though implicit redirect will not work for them. (372504)

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

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

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

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

General

- The CLI does not allow access to files whose names contain spaces. (539)
- LACP PDU fast rate ("lacp rate fast" on an Arista switch) should not be configured on any port in an MLAG pair, or any port connected to an MLAG pair. (13950)
- Copying boot-config containing a SHA-512 boot secret to switches with Aboot 2.0.7 or older will reject all boot passwords upon reload and the password recovery procedure will be required. (50695)
- When VmTracer is used in a network that includes Cisco devices, ESX hosts connected to the Arista switch will see CDP frames from the Arista switch as well as from the other Cisco devices. VmTracer will display the VM info when the ESX host reports Arista's CDP info, but then will delete it when the ESX host reports other CDP info.

  The workaround is to drop all inbound CDP on all ports using MAC ACLs. (101756)
- Running the CVX server on physical switches is no longer supported (149359)
- EOS SDK-based applications must be run via the 'daemon' command configuration. They cannot be run directly from bash. (158431)
- EOS extension scripts and agents which do not use EOS-SDK will need to be modified to work in 4.16.6M and beyond. (158467)
- Private VLAN is not supported in the EOS image distributed to US customers. (159408)
- EOS swix extensions containing RPMs which were based out of or procured from Fedora distributions prior to Fedora 18 may fail to install due to dependency issues. Suggested workaround is to recreate the swix files with the corresponding RPMs from Fedora 18 distribution. (162325)

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

  To access those commands exit the "router bgp" config mode to switch to the global config mode. (219390)
- Changes made to running-config after a config session is created might be reverted when the config session is committed which contains changes in similar or related areas. Use 'show session-config diff' inside the config session before committing to verify what changes will be made. (226850)
- The uptime in the output of "show module" may not match the uptime as shown by "show version" or "show uptime". The uptime in "show module" represents the time the module has been up since the last change in the "Status" column in the output of "show module", which differs from the meaning of uptime in "show version" and "show uptime". (248230)
- Intercept hosts on trunk ports with native vlan are ony supported on 7050X2 platforms (257580)
- To avoid interoperability problems when using strong SNMPv3 encryption algorithms, follow the following minimums:

  - When using AES192 for privacy, use a minimum of SHA224 for authentication.
  - When using AES256 for privacy, use a minimum of SHA256 for authentication. (268290)
- If /mnt/flash gets full and some EOS agent is stuck in the crash loop, the Fossil service will end up filling up /var/tmp. The work around is to free up some space on /mnt/flash. (375563)

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-F, DCS-7280SRAM-48C6-R, DCS-7280SRM-40CX2-F and DCS-7280SRM-40CX2-R

Layer 2

- 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)
- In MACsec, currently there is no show command to reflect the status of "sci" configuration. (277751)
- Port Security Protect does not support trunk ports. If a port is secured to source MAC <A> in VLAN <X>, traffic from source MAC <A> in other VLANs of the trunk port will also be allowed. (344773)
- While performing SSO, a device's peers might not receive LLDP packets within default TTL (120 seconds) and those peer devices may not maintain LLDP status information for the device undergoing SSO. After SSO is complete, the peers will again have LLDP state available.

  To avoid losing LLDP state during SSO, consider increasing LLDP hold time at the remote. (347459)
- If 'traffic unprotected allow' is configured in a Macsec profile, agent restart or Stateful Switchover may result in traffic loss (362766)

Layer 3

- ECMP is not supported for routes learned via RIP. (34773)
- DHCP relay is not supported when running virtual machines on EOS (101754)
- In a multi-topology IS-IS router, changing the address family or IPv4/IPv6 addresses of an IS-IS interface results in a flap of the IS-IS adjacency. (115993)
- CLI does not detect TCAM exhaustion and automatically revert the change when a PBR policy is applied to an L3 interface in the "shutdown" state, since the actual hardware programming happens when the interface comes up. (122651)
- OSPFv3 support for multiple address families (RFC 5838) in EOS introduces support for IPv4 routing with OSPFv3. OSPFv3 for IPv4 AF in EOS requires configuring neighboring OSPFv3 IPv4 routers on the same link with primary IPv4

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

- When BFD is unconfigured with Fhrp as the only client and/or BFD is administratively shut, both VRRP routers will become Master. This transition is
temporary and Master-Backup relationship will be restored. (159647)
- BFD echo functionality is not available on L2 Port-Channels with BFD RFC7130 mode enabled. (190418)
- If a port channel is configured with "bfd echo" and "bfd per-link rfc-7130" and is used for OSPFv3 or PIM, BFD sessions will not come up with echo functionality. BFD with SSO for OSPFv3 or PIM is not supported. (190997)
- BMP is not supported for router bgp instances configured in non-default VRFs. (210363)
- SSO is not supported when BFD is configured for ISIS IPv6 sessions. (239704)
- An LACP port channel configured with BFD per-link and echo may experience BFD session flap on existing member links if new member link(s) are added to the port channel while the BFD peer device undergoes Stateful Switchover (SSO). (255042)
- The DSCP value of BGP control plane packets set by the command 'bgp transport qos
dscp <value>' will be overwritten when the platform 'qos rewrite dscp' is enabled
and the outgoing interface is a SVI based on traffic-class to dscp map present on
the platform. (269764)
- VRF support is 256 VRFs on 32G memory platforms, 64 VRFs on >= 8G memory platforms, 32 VRFs on < 8G memory platforms. (271478)
- In the output of the "show ipv6 dhcp relay counters" or "show ip dhcp relay counters" command, interface-specific DHCP relay counters might not sum up to "All Req" and "All Resp" counters. (278786)
- Tunnel interfaces flap when configuring ttl, tos and path mtu discovery on them. (281464)
- IPv6 addresses are not supported as tunnel source and destination. (283912)

- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289217)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289224)
- When BFD is configured with per-link mode rfc-7130 for an IPv6 link-local address on a Port-Channel and the Port-Channel transitions to the Down state, the Port-channel may not transition back to the UP state.

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

Multi-agent

- When the multi-agent routing protocol model is configured, the maximum number of next hops for an OSPF ECMP route is 128. (274888)
- BGP session over indirectly connected interface that was down due to lack of IGP path to peering address may take around one minute to re-establish once the IGP path is available. (287628)
- Next hop groups cannot be used in a PBR match criterion. (294334)
- Route-map specified as match-map for dynamic prefix-list is evaluated only against BGP routes. (345444)
- The IpRib agent runs out of addressable memory when a large number of routes are leaked to a large number of VRFs. (346575)

Rib

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

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

MLAG

- Spanning-tree mode backup does not work in an MLAG configuration. (15151)

- Configuring a device connected to an MLAG with a round-robin LAG distribution algorithm is not supported if the device is going to participate in IGMP. (28370)
- On a scaled MLAG setup, Lag+LacpAgent agent may restart unexpectedly if MLAG is reinitialized due to configuration changes. (116125)
- MLAG ISSU is not supported for a direct upgrade/downgrade from/to 4.15.x & below to/from this release. For further details, please refer to the MLAG ISSU compatibility matrix. (159409)
- On an MLAG system configured with dual primary detection and with default control plane ACL enforced, MLAG will not form if the management interfaces of the MLAG peers are connected through extra hops (e.g., on different subnets), due to the default control plane ACL dropping MLAG control packets with TTL < 255.

  The work around is to apply a control plane ACL which permits MLAG control packets with the correct TTL. (271308)
- TCAM profile switch in a MLAG switch may result in some packet duplication for broadcast or multicast packets (341353)

Multicast

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

  Use "no ip igmp snooping VLAN <vlanId> immediate-leave". Otherwise, some desired multicast traffic might be lost. (21367)
- After ASU reload, if the first PIM hello packet received from the neighbor gets lost and If the neighbor has a high hello query-interval, it will result in traffic loss. Workaround is to use default PIM hello query-interval and increase the PIM hello query-count. (341447)

DCS-7150 Series

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

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

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

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

- Source port loopback suppression at high traffic rate can reduce the egress bandwidth of a port. This will happen when a port forwards incoming multicast traffic at a high rate to a multicast group which also has this port as a member. This port will receive back copies of the replicated traffic and drop them due to loopback suppression if there are no receivers on it at the egress stage of the port. (40294)

- NAT and VXLAN cannot be configured on the switch at the same time. (47258)

- When agile ports are configured for 40G with clause 73 autoneg (speed auto 40gfull) while using CR cables, the ports cannot be used in a port channel.

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

- Systems with A0 silicon, packets are not forwarded to the egress port when egress truncation and timestamping are enabled together. Removing the configurations resumes packet forwarding. Truncation and timestamping are not supported together. (72756)

- The VXLAN flood VTEP list cannot be more than 128k VLAN-VTEPs, where VLAN-VTEPs is the number of VLANs multiplied by the number of VTEPs. (73228)

- Ingress port access control list (PACL) on a VXLAN-enabled interface is not supported. If ingress PACL is configured on a VXLAN-enabled interface, in-hardware VXLAN encapsulation of incoming packets is not expected to work correctly. VXLAN encapsulation of packets in software is not affected. Removing ingress PACL configuration on the affected interface restores correct VXLAN hardware encapsulation behavior on that interface. (76085)

- Multicast-routed frames get tx-mirrored even when they do not egress the tx-mirrored multicast destination due to VLAN suppression. This is a limitation of the switch ASIC. (87933)

DCS-7170 Series

- If an ingress packet contains more than one 802.1Q tag and the egress port is a trunk port, the egressing packet may get corrupted. (180436)
  Series Affected: DCS-7170 Series

- BFD sessions are not supported on Port-channel interfaces. The sessions will not come up. (250315)

- Changing port speed from 10G/25G/50G to 40G/100G will cause the forwarding agent (BfnSliceAgent) to restart. This can result in a brief traffic outage in the order of 100ms on all ports. The link state of other ports will not change. (256519)

- After applying or removing shaping configuration on an L2 sub-interface, the L2 sub-interface must be flapped for the configuration to take effect. (281312)
  Series Affected: DCS-7170 Series

- After setting CoPP rate = 0, a few packets (<10) might be let into the CPU, but eveything will be dropped afterwards. (350676)

DCS-7300 and DCS-7500 Series

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

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

- Non per-link BFD sessions over port-channels may flap when a linecard where some of
  the port-channel's members reside is powered off with "no power enable". (363122)

- When a system shuts down due to a power overload, EOS does not log a 'power overload' reload cause (371160)

DCS-7020RA, DCS-7160, DCS-7170, DCS-7280R2A, DCS-7280RA and DCS-7500R with 7500RA or 7500R2A linecards Series

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

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

- The L3 MTU on interfaces is not enforced on SVIs. (11659)

- Disabling IPv6 routing on an L3 interface is not supported when IPv4 routing is also enabled on that L3 interface. (43093)

- Linecards with HW revision < 02.00, when configured for 100G and a physical layer problem is present, a one-way link may result. The link experiencing the physical layer problem will be down. However, the peer link may still be up.

The workaround is to detect this situation, by, for example, looking at the link status on both sides of the 100G link, and perform remedial actions; for example, replacing the transceiver or cable. Running a keepalive protocol like LACP on these links is advised. (67325)
SKUs Affected: 7500E-12CM-LC and 7500E-72S-LC

- show interfaces counters rates' for port channels may show a slightly inaccurate data rate (including above 100%). (67367)

- Linecards with HW revision < 02.00, under extreme scenarios where we receive linerate small packets on all interfaces connected to the same forwarding ASIC and ingress mirror or sflow dangerous is turned on many of these interfaces, some packets (both the original and the mirror/sFlow copy) can be dropped. (67942)
  SKUs Affected: 7500E-12CM-LC, 7500E-36Q-LC, 7500E-48S-LC and 7500E-72S-LC

- A VXLAN encapsulated packet after de-encapsulation is not forwarded back to the port on which the original packet was received.
  Similarly, a native ethernet packet after VXLAN encapsulation is not forwarded back to the original receive port towards the remote VTEP. (75075)
  Series Affected: DCS-7500E Series

- VARP and VRRP cannot be enabled at the same time. If VRRP is configured, there can be up to 128 combinations of interface and Virtual Router ID (VRID), but there can only be at most 8 different VRIDs and the VRIDs must differ only in the 3 least significant bits. VRRP for IPv4 and IPv6 cannot be configured at the same time. (75752)

- Due to hardware limitation, if an egress IP ACL is applied on an interface which is part of a VLAN with an egress router ACL, the filtering behavior for egress traffic of that interface may be undefined. (82094)

- For packets that hit rules across multiple ACLs, the hit counts are accounted for only in a single ACL. PACL counters take precedence over RACLs and MAC ACLs, and RACL counters take precedence over MAC ACLs. (85440)

- When an egress IPv6 RACL is applied to a VLAN, IPv6 packets may not be subjected to MTU checks. (88778)

- An egress ACL on a destination port of a monitor session only takes effect for Rx mirrored packets which are IP forwarded. The egress ACL will not work for bridged packets or Tx mirrored routed packets. (89915)

- When a shaper is configured the output of the shaper has a tolerance of +30% with 100-byte frames. This reduces to +5% as the frame size is increased to 600 bytes. The shaper variance reduces when the port happens to be the endpoint of a tunnel. (93546)

- The egress per-queue byte counter is not exact for certain packet types. For routed packets the egress per-queue byte counter may be offset by +/-2 bytes. For any packets originating from the CPU, the per-queue byte counter may be offset by -12 bytes. (97660)
- Packets that attempt to do a GRE key lookup where the GRE key lookup misses will have egress IPv6 ACLs applied with the ingress VLAN (102455)
- GRE key forwarding packets that are recirculated will have no QoS policy applied to them. (102458)
- If a double tagged frame is received on an untrusted or DSCP-trusted interface, the CoS is retained if the outgoing frame is tagged, unless the frame is routed (131614)
- Mirroring to GRE tunnels does not support GRE destinations that are reachable by sub-interfaces. (136130)
- IPv6 Egress ACL deny logging is not supported on subinterfaces. (146487)
- PBR policy is supported only on interfaces in the default VRF. If a PBR policy is configured on an interface in a user defined VRF, the behavior is undefined. (148895)
- Disabling IPv4 routing on an L3 interface is not supported when IPv6 routing is enabled on that L3 interface. When "no ip routing ipv6 interfaces" is configured to enable the forwarding of IPv4 packets over IPv6 nexthops over interfaces that do not have IPv4 address, but only a IPv6 address, it will also allow routing IPv4 traffic over these interfaces. (151356)
- VXLAN routing will work only with /32 attached hosts in LEM. This impacts traffic flows that involve route and encap operations. In particular, we do not support LPM lookups with VXLAN routing. Consequently, VXLAN routing is mutually exclusive with the "ip hardware fib optimize prefixes profile internet" configuration. (179037)
- When the 'vxlan-routing'
  hardware TCAM profile is configured, VXLAN traffic flows in overlay will not work
  alongside IPv6 traffic flows in underlay. (184930)
- When the 'vxlan-routing' hardware TCAM profile is configured, VXLAN traffic flows in overlay may not work alongside IPv6 traffic flows in underlay. The workaround is to use the CLI command 'no platform sand ipv6 host-route exact-match'. (190017)
- L2-ECMP over VXLAN is not supported for multi-homed hosts under EVPN. (193475)

- Linecards with HW revision < 02.00, IP-in-IP encapsulated packets might not get forwarded after termination if the packet size falls within the range 223-287 bytes. (193573)
  SKUs Affected: 7500E-12CM-LC, 7500E-36Q-LC, 7500E-48S-LC and 7500E-72S-LC

- In Tap Aggregation mode, Tap Aggregation features may not work for 802.1BR/VN tagged packets when 802.1BR/VN tag stripping is not configured. (198798)

- Links are not compatible with systems running EOS 4.18.1F

  The workaround is to upgrade both sides of the link. (199152)
  SKUs Affected: 7500R-8CFPX-LC

- Traceroute to a IP route where the IP route is getting forwarded through a MPLS nexthop-group may incorrectly display the first hop as being unreachable. The actual first hop will appear in the second position. (209273)

- When ingress packet truncation is enabled in Tap Aggregation mode, truncated packets are not counted as packets received on the ingress interface. (215346)

- Frames less than 64 bytes are dropped (as expected),
  but not counted.

  Command "hardware phy cmx42550 <intf > diag read64 mac line 0x138"
  can be used to dump the counter for the purpose of debugging". (218429)
  SKUs Affected: 7500R2M-36CQ-LC

- Vxlan decapped packets will not match rules that match on 'tcp established' in egress acls (226582)

- Native VLAN setting on Tap interfaces does not take effect. (227770)

- While in Tap Aggregation mode, ingress VLAN membership filtering does not work when packets are destined to a tool interface configured with egress packet truncation. (227841)

- While in Tap Aggregation mode, interfaces configured for egress truncation will continue to consume hardware resources even when they are down. (228384)

- LLDP functionality is not supported on interfaces with ingress truncation enabled. (229983)

- If multiple tx-queues of a port are configured with low shape-rate values, traffic egressing out of that port on tx-queues that are not shaped, may get impacted. (241118)

- With MTU increased beyond 9100B, there can be some multicast packets of size greater than 9100B that are dropped as a port gets close to linerate. (261446)
- Changing PMF profile on one linecard may cause a brief disruption of PMF-based features on other linecards which should not be affected by the same change. (270849)
- A QoS policy, with a class map which has either set-tc or set-dscp on the ingress port and which has DSCP rewrite map on the egress port, results in indeterministic behavior with respect to DSCP remarking of the outgoing packet. (273320)
- Rate values for subinterfaces shown by "show interface counter rates" may deviate up to 15% from corresponding rate on parent interface. (278700)
- After applying or removing shaping configuration on an L2 subinterface, the L2 subinterface must be flapped for the configuration to take effect. (278730)
- Packets larger than 10196 bytes will be discarded by the forwarding ASICs. (280459)
- PFC on DSCP trusted ports is not supported. (281214)
- L3 forwarding through SVIs is not support on L2 sub-interfaces. (293410)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- When a port becomes a LAG member, there may be a very brief moment where a few packets may be duplicated on the LAG. (295260)
- Fragmented IP packets will not match on flow-spec rules that include a match on TCP/UDP port. (297309)
- When all the L2 subinterfaces belonging to a parent interface are shutdown, packets are bridged and mac addresses are learnt on parent interface (300202)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series
- VXLAN routing is not supported on a modular system on which one of the linecards is configured for TAP aggregation. (306180)
- Configuring policy-map on lag l2 sub-interface with member ports of the lag also having l2 sub-interfaces may cause SandTunnel to crash. Work around is to remove the l2 sub-interfaces on the member ports. (338794)
- When a BGP peer session is cleared, traffic which hits existing flow-spec rules may increment the wrong counter for a brief period of time. (339523)
- If non shaped sub interfaces are configured under same interface having shaped sub interfaces, traffic will egress out of the non shaped sub

interfaces in a round robin fashion. It is recommended to not have shaped and unshaped sub interfaces under the same parent. (341851)
  Series Affected: DCS-7280R, DCS-7280R2, DCS-7500R and DCS-7500R2 Series

- Stale telemetry flow-spec statistics may be seen until the next "show flow-spec" command is issued. (345195)
- Filtered mirroring is not supported on linecards or fixed-systems operating in TAP Aggregation mode when traffic steering is enabled. (349883)
- If more than maximum supported Ipsec tunnels are configured, it is random as to which ones are assigned flow entries within on-chip encrypt/decrypt engine. The set of Ipsec tunnel interfaces assigned flow entries may not be the same as the set of tunnel interfaces which have IKE connections established. As a result, some tunnel interfaces with IKE connection established may not get link up. (376431)

DCS-7280R and DCS-7500R Series

- On the system with DCS-7500R linecards, if DCS-7500R2 linecards are inserted, routes may consume more resources in the hardware routing table. (215393)
- Jericho chip buffering DRAM test at agent startup does not catch bad parts while it should fail to force RMA of bad parts. (218971)
- IPv6 multicast routes are stored in two separate TCAM resources. The group IP is stored in one TCAM while the source IP is stored in another. Since one TCAM contains the source IP it will contain an entry for every multicast route. However, the group IP TCAM will reuse entries for duplicate group IPs. For example, if you have routes (S1,G1) and(S2,G1), one TCAM entry is required for the group IP and two TCAM entries are required for the source IP. This means it is possible that more TCAM resources are allocated to the source IP lookup than the group IP lookup. The scaling achieved during the initial distribution of multicast routes will be preserved. However, if the distribution later changes, optimal scaling for the new distribution may not be reached. (257364)
- MPLS ping and traceroute are not supported for VRF label routes. (343526)
- If mirroring to CPU is configured for an interface and an incoming MPLS packet on that interface matches a configured "mpls static vrf-label" entry, the packet will not be mirrored to CPU. (359362)
- Ingress ACLs and PBR are not supported for MPLS tunnel terminated traffic using VRF label. (366695)
- VRF terminated traffic is not sampled with sFlow. (369010)

DCS-7020 Series

- The interface bin counter for both in and out packets with sizes 1523-Max may be inaccurate. Packets with sizes between 9217-9236 are not counted by the switch for interfaces Ethernet 1-48. (188121)

- Mirroring and sFlow can not be supported on the same source interface at the same time due to a hardware limitation. (209060)

- The current implementation of IPSec supports only encryption/authentication algorithm combinations of AES128/SHA1 or AES256/SHA256. No other combinations are currently supported. (359986)
  SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

- Non-default VRF assignment for IPsec tunnel interface is not supported. Assigning a non-default VRF to IPsec tunnel interface will cause all decrypted packets for tunnel interface to be forwarded by the CPU (slow path). (377965)
  SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

- Using an Ipsec tunnel interface for policy-based routing nexthop is not supported. (378861)
  SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

- Ipsec tunnel interfaces are not supported as part of nexthop-group. (378890)
  SKUs Affected: DCS-7020SRG-24C2-F and DCS-7020SRG-24C2-R

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

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

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

- System does not stop transmitting traffic for the exact amount of time quanta received in PAUSE or PFC frame. This can lead to packet loss during congestion if the peer switch does not have enough buffers. The workaround is to use XON/XOFF instead of time based flow control, or to increase the time to account for one MTU sized packet. (27190)

- Running the 10G SFP+ ports of the DCS-7050Q at 100Mbps is not supported (30130)
  Series Affected: DCS-7050Q Series

- QoS shaping and guaranteed bandwidth will only work in store-and-forward mode. They are not supported in cut-through mode. This includes 'shape rate X' and the 'bandwidth guaranteed X' commands. (56790)

- L3 packets destined to a next hop whose MAC address is not learned and which fail the uRPF check configured on the incoming SVI are flooded in the ingress VLAN. (75660)
  Series Affected: DCS-7050 Series

- Accelerated Software Upgrade (ASU) from 4.13.x (x < 6) to 4.13.6 and above requires a special upgrade procedure due to a software image format change. (90097)

- 10 Gbps packet replication on all SFP+ ports at the same time may lead to packet discard. (91149)

- Egress L2 MTU violations in cut-through mode may result in unexpected discards of other frames. (95577)

- VXLAN encapsulated packets cannot be TX mirrored at the egress properly. The mirrored packets will have an incorrect MAC header. (97272)

- Forwarding-table partition mode 4 is not supported. (100862)
  Series Affected: DCS-7010T Series

- When IGMP snooping is enabled a single unknown multicast floodset is shared across all VLANs with IGMP snooping enabled. This can cause unexpected flooding to a trunk port that has no multicast routers on a given VLAN, but has multicast router attached to the same trunk port on another VLAN. (105188)

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

  This affects ip access-list, ip standard access-list, and ipv6 access-list. ipv6 standard access-list
  is not impacted. (105498)
  Series Affected: DCS-7010T, DCS-7050X and DCS-7050X2 Series

- "hardware access-list resource sharing vlan in" configuration will not share resources required by IPv6 standard access-lists. (105502)
  Series Affected: DCS-7010T and DCS-7050X Series

- When "hardware access-list resource sharing vlan in" is configured, Link Local Multicast traffic will not be intercepted by the ACL attached to the VLAN interface. (105504)
- Subinterfaces cannot be used to send or receive VXLAN encapsulated packets. (105767)
- When "hardware access-list resource sharing vlan in" is configured, changing the ACL on a VLAN interface:

  1. From an ACL that is not shared with other VLAN interfaces to one that is shared with other VLAN interfaces can cause deny traffic to get permitted until the new ACL takes effect.

  2. To another ACL that doesn't fit in the TCAM can cause deny traffic to get permitted until the system reverts back to the previously applied ACL. (106646)
- If IP-in-IP decap groups are configured, "qos trust dscp" configuration is not supported on traffic matching the decap groups. (107579)
- Vxlan and MPLS features may not be configured at the same time. (109330)
- When "hardware access-list resource sharing vlan in"" is configured and TCAM is close to being full, restart of the forwarding agent or a reboot of the box, can cause some ACLs to not get programmed properly. (109876)
- The Accelerated Server Upgrade feature is not supported with the "hardware access-list resource sharing vlan in" configuration. (110114)
- IGMP snooping is not supported on Vxlan enabled VLANs. The multicast packets will always be flooded locally and to remote VTEPs. (110479)
- A front-panel or fabric port may get stuck in an errored state and drop traffic egressing the port. When this happens, the forwarding agent log will contain "port <internal port #> start error detected". This condition will persist until it is manually cleared. A forwarding agent restart is required to clear the error and forward traffic normally. (112051)
- MPLS is not supported on subinterfaces. (113110)
- Packets being VXLAN encapsulated are constrained to a single underlay physical nexthop per physical egress port.

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

  Topologies involving virtual VTEPs connected via a downstream L2 switch

will not work, as the receiving vswitch will not have the capability to route the packet to the right VTEP. (113722)

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

  To recover from this state, delete/remove the access-list and then reconfigure the access-list. (114028)
- Configuring "hardware access-list resource sharing vlan in" in a config replace session or in config session is not supported (114291)
  Series Affected: DCS-7010T and DCS-7050X Series
- L2 flooding of BUM packets on Vxlan vlans uses L3 replication tables, which are also used by IP multicast routing. The replication tables have an entry for each vlan, physical port combination (lag or non lag). We will run out of this table resources if we need more than 64k entries. (114517)
- Toggling "hardware access-list resource sharing vlan in" configuration could cause ACLs that were already programmed on VLAN interfaces to not get programmed. This could traffic loss on the VLANs where the ACLs were originally applied. (115348)
  Series Affected: DCS-7010T and DCS-7050X Series
- If the switch is in cut-through mode, L2 MTU violations will not be counted if the packet headers are changed while forwarding. (116760)
- QoS policies are not supported for L3 Unicast flooded packets (routed packets with nexthop MAC not learned) (123883)
- When a PBR is attached to any interface, Layer-3 packets which are getting flooded in the egress VLAN will not be treated with Egress Router ACL if one is configured on that egress VLAN. (133144)
- VXLAN routing is not supported on systems with
  BCM56750 fabric chips. Such switches can be identified using the CLI command "show
  platform trident l3 summary", which will show bcm56750 for "Fabric chip model". (134625)
- Mirroring of incoming multicast traffic to a GRE tunnel may result in data traffic loss on the egress interface in an oversubscription scenario. (139591)
- If hardware tables for VLAN translation overflows, the entries need to be un-configured and re-configured after the overflow condition is removed due to appropriate configuration changes. (157196)

- Up to 4 mirroring sessions are supported.

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

- 1) Multicast traffic matching reserved IPv4 multicast address range 224.0.0.X is not counted by L3 interface counters.
  2) L3 interface counters not supported for IPv6 multicast traffic.
  3) Unicast traffic to CPU is counted by L3 interface counters only if routing is enabled (i.e. ip routing for IPv4 traffic and ipv6 unicast-routing for IPv6 traffic) (160930)

- Mirroring of Vxlan encapsulated packets in the egress direction is not supported. (167189)

- When a packet that is routed and flooded on the egress vlan is sent to a port in congestion, it will not result in PFC frames being sent on the source port of the packet. (174399)

- "switchport dot1q ethertype" command is disabled. (186577)

- If ingress ACL is applied on an interface, traffic sent to CPU from that interface is falsely counted as InAclDrop in "show int counters ingress acl drop". (203465)

- When "reload fast-boot" is performed from a release prior to 4.17.2, switches with BGP peering connections with the switch undergoing "reload fast-boot" could experience BGP KeepAlive timeouts before they see port flaps. This could result in longer than expected data plane downtime during the upgrade. (214700)
  Series Affected: DCS-7050X, DCS-7060X and DCS-7060X2 Series

- A multi-bit ECC error in forwarding ASIC packet memory cannot be automatically corrected. A hitful restart of the slice agent managing the forwarding ASIC or a reboot of the system is required to clear the condition. (214950)

- If ip-ttl is the only enabled field for port channel hashing, after reload hitless, packets may get hashed to a different link than before the reload. (223568)
  Series Affected: DCS-7050TX, DCS-7050X and DCS-7060X Series

- Sflow output interface determination will not work when the output interface is a subinterface. (250424)

- A switch may not be able to bridge or route VXLAN traffic if it hits one or all of following conditions:
  - Nexthop resources are exhausted as indicated by syslog VXLAN_HWHER_NEXTHOP_RESOURCE_FULL.

- MAC table overflow as indicated by show platform trident mac-address-table missing.

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

- Flexible port-channel hashing may not hash accurately when the inner L4 header option is configured for IPV6 GRE packets. (267487)
- If a packet stream is policed by multiple policers, the policed rate may be lower than any of the configured policer rates. (271046)
- To perform filtered mirroring of a packet based on VLAN from an interface configured to perform VLAN translation, the ACL must be programmed with the translated VLAN. The mirrored packet will be the same as the packet at the source interface. (278599)
- BGP neighbor sessions may flap when URPF is configured on an interface. (279113)
- To perform filtered mirroring of a packet based on VLAN from a subinterface source, the ACL must be programmed with the internal vlan of the subinterface. The mirrored packet will be the same as the packet at the source interface. (280943)
- Matching on inner VLANs is not supported in mirror ACLs. If an entry is specified matching on inner VLAN, it will be ignored and permit all packets. (284371)
- Added a new Cli command 'show platform trident mmu buffers all' to display the buffer usage and percentage utilization. (285602)
- SSU is supported when upgrading from the associated release in each release train: 4.18.8, 4.19.8, 4.20.7, 4.21.0. (289548)
- Issuing "no shutdown" on a member link of port-channel may send duplicate packets for sub-second on that member link (291514)
- DirectFlow (and features such as the MacroSegmentation Service that use DirectFlow) are unsupported on the 7260, 7300 series of switches. (296715)
- When primary vlan and secondary vlan are part of different MST instances, the hardware programming of lag will not happen. (300119)
- Not all interfaces in the groups Ethernet49-Ethernet52 and Ethernet53-Ethernet56 can be broken out in to 4x25G, 4x10G or 2x50G mode due to MAC resource limitation on the system. Interfaces without MAC resource will be marked inactive and 'STRATA-4-HW_PORT_RESOURCE_FULL' will be logged in syslog.

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

- When MLAG peer MAC routing is enabled, OSPF neighborship over VXLAN overlay may never get established. (349242)
- Ip-length based rules are not supported on Egress ACLs. (353247)
- Packets dropped in the egress pipeline may still be egress mirrored successfuly. (357609)
- Some ports (esp. Et73-Et96 and Et103-Et104) might blackhole traffic or fail to linkup. (358362)
  SKUs Affected: DCS-7050TX2-128-F and DCS-7050TX2-128-R

DCS-7010T Series

- IEEE 1588 PTP transparent clock does not decrement TTL for PTP routed packets (55078)
- Enabling counters for Static or Twice NAT connection
  can cause FocalPointV2 agent restarts, when the number of connections exceed 275. (109048)

DCS-7060X and DCS-7300 Series

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

  The workaround for this is to switch to store-and-forward mode where the limitation is not present. (154721)
- Cut-through mode is not supported on the SFP+ interfaces. (182728)
- Cut-through mode is not supported at 1G speed. (219427)

DCS-7260X3 Series

- LANZ is not supported. (217543)
- Autonegotiation is not supported (219223)
- Configuring 100G, 50G, 40G, 25G and 10G simultaneously might result in unexpected packet forwarding behavior. (235075)

7368 Series

- Multicast replication resources in TH3 are only 10% of their unicast counterparts. Due to this multicast replication in TH3 is limited to 1.2 Tbps. (297317)
- Layer2 source MAC learning may occur on packets received with CRC errors. (325507)
- 802.3X pause is not supported (348758)
- Routed unicast traffic egressing a SVI for which destination MAC is not learnt are sent to CPU and then flooded on the egress VLAN. (368228)

DCS-7050X and DCS-7300 Series

- The mirroring to GRE Tunnel feature does not support sub-interfaces as egress interfaces. (136514)
- Packets sourced from the CPU to be VXLAN encapsulated will not have the right DSCP marking in the outer header based on the global QOS map configuration. (139248)
- Tcam entries used by IPv6 standard access list is not shared across interfaces (159827)
- EOS will fail to detect modules, such as some 100GBASE-DWDM4 modules, that require deassertion Low Power Mode (LPMODE) pin to be de-asserted. (160016)
  Series Affected: DCS-7050S and DCS-7050T Series
- In MLAG VTEP configuration, the switch incorrectly drops VXLAN encapsulated packets that are destined to the VTEP IP address with the destination mac as peer switch's bridge MAC address. The workaround is to configure VARP in such scenarios. (173716)
- Multicast traffic destined to boundary or prefix mroutes configured may be forwarded out of order if fabric priority flow control is enabled for the given priority. (185981)
- If an SVI is used as a core interface to reach remote Vteps, any L3 EVPN traffic using that SVI will be dropped if the underlay mac becomes unlearnt. A workaround is to configure MAC timeout to a higher value than ARP timeout to ensure macs don't get aged out. (192419)
- During "reload hitless", ECMP traffic flows may get hashed to a different ECMP path than before the reload. (218700)
  Series Affected: DCS-7050X Series

- Direct flow on T2+ has the following limitations
  1. Only IP packets can have their source mac , destination mac and vlan changed.
  2. When a packet has its smac, dmac and vlan changed then it can only be forwarded to one interface.
  3. When a packet has its vlan changed it can only be forwarded untagged on access ports.
  4. An output interface must be specified whenever a packet's smac,dmac and/or vlan is changed. (251477)

DCS-7050X2 and DCS-7300 Series

- When IEEE 1588/PTP is configured on interfaces running at 100Mb/s speeds, the master offset and mean path delay calculations are inaccurate due to a limitation in the internal PHY (122816)
- Egress VLAN translation will not work on routed multicast packets when the VLAN to be translated is part of the outgoing interface list. (151093)
- VLAN translation is not supported on vxlan core ports (154309)
  Series Affected: DCS-7050X2 Series
- Due to an increase in the number of default entries programmed in the TCAM, 'Out of TCAM entries', might be seen while configuring ACLs. This happens due to the TCAM hardware resources being full. The workaround is to unconfigure any unused feature that requires TCAM resources. (182178)

DCS-7050CX3, DCS-7050SX3 and DCS-7300 Series

- Configuring any speed less than 1G will cause continuous restart of forwarding agent. (225977)
  Series Affected: DCS-7050X3 and DCS-7260X3 Series
- When configuring a DirectFlow action, specifying "interface hardware-loopback" as the destination of "action output" is not supported. (243031)
- IPv6 packets with WESP payload (next header 141) bypass the IPv6 security ACLs and service policies. (243387)
- Cut-through mode is not supported on SFP+ interfaces. (252731)
- Directflow configuration with "set vlan", but no "action set output interface" configured is not supported. (345852)
- NAT is not supported on subinterfaces (353236)
  SKUs Affected: 7300X3-32C-LC and 7300X3-48YC4-LC
- NAT is not supported on ECMP interfaces (353704)
  SKUs Affected: 7300X3-32C-LC and 7300X3-48YC4-LC

DCS-7160 Series

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

  The workaround is to ensure that the underlay MAC of the next-hop to reach remote VTEPs are always present (which is the normal operation). (178395)
- Control plane Policy Map counters are not supported. (180303)
- A maximum of 3840 remote VTEPS are supported when the switch is configured as a Vxlan Vtep. (188553)
- The IGP IS-IS protocol is not supported in overlay VXLAN networks. (216353)
- The DSCP value of BGP control plane packets set by the command 'bgp transport qos dscp <value>' is overwritten when 'qos rewrite dscp' is enabled and the outgoing routed interface is a SVI. The rewritten value is based on traffic-class to DSCP map configuration. (288617)

Transceiver Series

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

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

vEOS Router Series

- The /mnt/flash/cloud_ha_config.json based configuration for the CloudHa agent is not supported from this release. If you are upgrading from an older image (< 4.20.5) please reconfigure HA using EOS CLI.

Please see information on converting json file based config to CLI commands in the vEOS Router Configuration Guide. (264660)

- While vEOS instances in an AWS cloud can be created using either a Bring Your Own License (BYOL) or Pay As You Go (PAYG) format, there currently is no show command in vEOS to easily help the user or support determine which mode this instance uses. (272649)
- Tx/Rx ring parameters of an "Elastic Network Adapter (ENA)" intterface cannot be set through config.json in vEOS. (276382)
- vEOS supports up a maximum of 8 physical cores. With hyperthreading enabled it is 16 vCPUs (278286)
- In dpdk based vEOS(MODE=sfe in /mnt/flash/veos-config), SR-IOV mode doesn't work with Intel x520/82599 on ESXi. As a workaround use mixed mode instead of the SRIOV only mode in ESXi (296718)
- Hot detaching network interface instances is unsupported. Reboot the instance after deattching a network interface. (306132)

Conditions and Impacts

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

Condition: Rib agent restart
Impact: When the Rib agent goes down, all routing sessions are terminated; all BGP sessions drop, all OSPF adjacencies are lost, and neighbors stop forwarding traffic to us. When the Rib agent restarts, adjacencies are re-formed, and, after protocol convergence, traffic flows through us again. In most cases where there is adequate networt-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.