

Arista Networks EOS 4.34.3.1M Release Notes

Version: 0.2

04 November 2025

Table of Contents

EOS 4.34.3.1M Release Highlights

New Platforms and Hardware

SNMP MIBs

SNMP Traps

Supported Hardware

Reference to MLAG ISSU Upgrade/Downgrade Table

Resolved Caveats in 4.34.3.1M

7300X3, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7060X6, DCS-7260X3 and DCS-7300X Series

7300X3 and DCS-7060X Series

DCS-7060X2 Series

DCS-7260X3 Series

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

Known Software Caveats in 4.34.3.1M

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

Containerized EOS

vEOS Router / CloudEOS

CVP

Network Provisioning

Telemetry

CVX

Enterprise routing

SDWAN

General

IPSEC

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MLAG

MPLS

Multi-domain Segmentation Services

Multicast

NAT

SwitchApp

DCS-7130 and DCS-7170 Series

DCS-7130 and DCS-7170 Series

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

AWE-5000, AWE-7000 and ZTX-7000 Series

DCS-7170 Series

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

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

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

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

DCS-7280R3A and DCS-7800R3A Series

DCS-7280R3A Series

DCS-7280R3 Series

DCS-7280R3 Series

DCS-7280R3 Series

DCS-7280R3A Series

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

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

7300X3, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7060X6, DCS-7260X3 and DCS-7300X Series

CCS-750 Series

7300X3 and DCS-7060X Series

DCS-7060X2 Series

DCS-7260X3 Series

7368 and DCS-7060X4 Series

DCS-7060X5 Series

7388 and DCS-7060X5 Series

DCS-7060CX5 and DCS-7060DX5 Series

DCS-7060X6 Series

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

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

CCS-720DF, CCS-720DP and CCS-720XP Series

CCS-750 Series

DCS-7050X3 Series

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

7358 and DCS-7050X4 Series

DCS-7050X4 Series

7388 and DCS-7060X5 Series

DCS-7060X6 Series

Transceiver Series

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

7358 and DCS-7050X4 Series

Limitations and Restrictions in 4.34.3.1M

Accelerated System Update

BGP EVPN L2/L3 VXLAN/MPLS

Containerized EOS

vEOS Router / CloudEOS

CVX

Enterprise routing

SDWAN

EOS SDK

General

IPSEC

Layer 1

Layer 2

Layer 3

Multi-agent

Rib

MetaMux

MetaWatch

MLAG

MPLS

MultiAccess

Multicast

NAT

Switch Aggregation Group

SwitchApp

DCS-7130 and DCS-7170 Series

DCS-7130 and DCS-7170 Series

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

AWE-5000, AWE-7000 and ZTX-7000 Series

CCS-750 and DCS-7060X4 Series

DCS-7170 Series

DCS-7130, DCS-7132 and DCS-7135 Series

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

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

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

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

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

DCS-7280R3A and DCS-7800R3A Series

DCS-7280R3A Series

DCS-7280R3 Series

DCS-7280R3 Series

DCS-7280R3A Series

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

7300X3, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7060X6, DCS-7260X3 and DCS-7300X Series

CCS-750 Series

7300X3 and DCS-7060X Series

DCS-7060X2 Series

DCS-7260X3 Series

7368 and DCS-7060X4 Series

DCS-7060X5 Series

7388 and DCS-7060X5 Series

DCS-7060CX5 and DCS-7060DX5 Series

DCS-7060X6 Series

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

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

CCS-720DF, CCS-720DP and CCS-720XP Series

CCS-750 Series

DCS-7050X3 Series

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

7358 and DCS-7050X4 Series

DCS-7050X4 Series

7388 and DCS-7060X5 Series

DCS-7060X6 Series

Transceiver Series

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

7358 and DCS-7050X4 Series

Conditions and Impacts

EOS 4.34.3.1M Release Highlights

New Platforms and Hardware

- None

SNMP MIBs

- ARISTA-ACL-MIB
 - aristaAclDpSupportFlags, aristaIpAclRuleStatsTable, aristaIpAclRuleTable, aristaIpAclTable, aristaIpv6AclRuleStatsTable, aristaIpv6AclRuleTable, aristaIpv6AclTable, aristaMacAclRuleStatsTable, aristaMacAclRuleTable, aristaMacAclTable
- ARISTA-ASIC-COUNTERS-MIB
 - aristaAsicInternalDropStatsRatesSupported, aristaAsicInternalDropStatsTable
- ARISTA-BGP4V2-MIB
 - aristaBgp4V2BackwardTransitionNotification, aristaBgp4V2EstablishedNotification, aristaBgp4V2PeerConfiguredTimersTable, aristaBgp4V2PeerCountersTable, aristaBgp4V2PeerErrorsTable, aristaBgp4V2PeerEventTimesTable, aristaBgp4V2PeerNegotiatedTimersTable, aristaBgp4V2PeerTable, aristaBgp4V2PrefixEvpnNlriGaugesTable, aristaBgp4V2PrefixGaugesTable
- ARISTA-BRIDGE-EXT-MIB
 - aristaDot1qTpFdbTable, aristaMacAge, aristaMacLearn, aristaMacMove, aristaMacStatsTable
- ARISTA-CLB-MIB
 - aristaClbDestinationInterfaceTable, aristaClbPortGroupAllocatedFlowAtOrAboveLimit, aristaClbPortGroupAllocatedFlowAtOrAboveThresholdWarning, aristaClbPortGroupAllocatedFlowBelowLimit, aristaClbPortGroupAllocatedFlowBelowThresholdWarning, aristaClbPortGroupNameMappingTable, aristaClbPortGroupTable, aristaClbStatus, aristaClbTotalAllocatedFlows, aristaClbTotalLearnedFlows, aristaClbTotalUnallocatedFlows
- ARISTA-CONFIG-COPY-MIB
 - aristaConfigCopyCommandTable
- ARISTA-CONFIG-MAN-MIB
 - aristaCmdHistoryEventTable, aristaCmdHistoryRunningLastChanged, aristaConfigManEvent
- ARISTA-CONMON-MIB
 - aristaHttpProbeTable, aristaIcmpProbeTable, aristaTcpProbeTable
- ARISTA-DAEMON-MIB
 - aristaDaemonDataTable, aristaDaemonEnabledTable, aristaDaemonOptionTable, aristaDaemonRunningTable

- ARISTA-ENTITY-SENSOR-MIB
 - aristaEntSensorAlarm, aristaEntSensorThresholdTable, aristaFanStateTable
- ARISTA-EXTERNAL-ALARM-MIB
 - aristaExternalAlarmAssertedNotif, aristaExternalAlarmDeassertedNotif, aristaExternalAlarmTable
- ARISTA-FIB-STATS-MIB
 - aristaFIBStatsByPrefixLenTable, aristaFIBStatsByRouteTypeTable, aristaFIBStatsSummaryTable
- ARISTA-GENERAL-MIB
 - aristaReloadCauseTable
- ARISTA-HARDWARE-UTILIZATION-MIB
 - aristaHardwareUtilizationAlert, aristaHardwareUtilizationTable
- ARISTA-IF-MIB
 - aristaIfTable
- ARISTA-IP-MIB
 - aristaIpIfStatsTable
- ARISTA-LOGMGR-MIB
 - aristaLogForwardingStatsTable
- ARISTA-MAU-MIB
- ARISTA-NEXTHOP-GROUP-MIB
 - aristaNexthopGroupCounterTable, aristaNexthopGroupTable
- ARISTA-PFC-MIB
 - aristaPfcPriorityTable, aristaPfcWatchdogTxQueueTable
- ARISTA-PHY-MIB
 - aristaPhyIntfTable
- ARISTA-PRODUCTS-MIB
- ARISTA-QOS-MIB
 - aristaClassMapMatchTable, aristaClassMapTable, aristaEcnCounterTable, aristaEcnQueueCounterTable, aristaPolicyMapActionTable, aristaPolicyMapClassTable, aristaPolicyMapTable, aristaQosPerIntfStatsTable, aristaQosPolicerStatsTable, aristaQosStatsTable, aristaServicePolicyTable
- ARISTA-QUEUE-MIB
 - aristaEgressQueueTable, aristaIngressQueueTable
- ARISTA-REDUNDANCY-MIB
 - aristaRedundancyLastSwOverReason, aristaRedundancyProtocolConfig, aristaRedundancyProtocolOper, aristaRedundancySwitchOverNotif, aristaRedundancyUnitStateTable
- ARISTA-SMI-MIB
- ARISTA-SNMP-TRANSPORTS-MIB
 - aristaAuthFailTrapSrcTAddress, aristaAuthFailTrapTDomain
- ARISTA-STORM-CONTROL-MIB
 - aristaStormControlIfTable
- ARISTA-SW-IP-FORWARDING-MIB
 - aristaSwFwdIpStatsTable

- ARISTA-TAPAGG-MIB
 - aristaTapaggPolicyClassTable, aristaTapaggPolicyTable
- ARISTA-TEST-MIB
 - aristaTestNotification
- ARISTA-TUNNEL-MIB
 - aristaIpsecTunnelIfTable
- ARISTA-VRF-MIB
 - aristaVrfIfTable, aristaVrfTable
- ARISTA-VXLAN-MIB
 - aristaVxlanVniCountersTable, aristaVxlanVtepCountersTable, aristaVxlanVtiVniCountersTable
- ARISTA-XCVR-DWDM-MIB
 - aristaXcvrDwdmTable
- ARISTA-XGS-MIB
 - aristaXgsCpuQueueStatsTable, aristaXgsIfTable, aristaXgsNexthopEcmpInUseEntries, aristaXgsNexthopEcmpInUseSets, aristaXgsNexthopEcmpMaxEntries, aristaXgsNexthopEcmpMaxSets, aristaXgsNexthopEcmpOverlayInUseEntries, aristaXgsNexthopEcmpOverlayInUseSets, aristaXgsNexthopEcmpOverlayMaxEntries, aristaXgsNexthopEcmpOverlayMaxSets, aristaXgsNexthopEcmpUnderlayInUseEntries, aristaXgsNexthopEcmpUnderlayInUseSets, aristaXgsNexthopEcmpUnderlayMaxEntries, aristaXgsNexthopEcmpUnderlayMaxSets, aristaXgsQueueWatermarkTable
- BGP4-MIB
 - bgpBackwardTransNotification, bgpEstablishedNotification, bgpIdentifier, bgpLocalAs, bgpPeerTable, bgpVersion
- BRIDGE-MIB
 - dot1dBaseBridgeAddress, dot1dBaseNumPorts, dot1dBasePortTable, dot1dBaseType, dot1dTpAgingTime, dot1dTpFdbTable, dot1dTpLearnedEntryDiscards
- ENTITY-MIB
 - entAliasMappingTable, entConfigChange, entLastChangeTime, entPhysicalContainsTable, entPhysicalTable
- ENTITY-SENSOR-MIB
 - entPhySensorTable
- ENTITY-STATE-MIB
 - entStateOperDisabled, entStateOperEnabled, entStateTable
- EtherLike-MIB
 - dot3HCStatsTable, dot3PauseTable, dot3StatsTable
- HC-RMON-MIB
 - etherStatsHighCapacityTable, hcRMONCapabilities
- HOST-RESOURCES-MIB
 - hrDeviceTable, hrMemorySize, hrProcessorTable, hrStorageTable, hrSystemDate, hrSystemMaxProcesses, hrSystemNumUsers, hrSystemProcesses, hrSystemUptime

- IEEE8021-CFM-MIB
 - dotlagCfmMaNetTable, dotlagCfmMdTable, dotlagCfmMepTable
- IEEE8021-PFC-MIB
 - ieee8021PfcIfTable
- IEEE8021-SECY-MIB
 - secyCipherSuiteTable, secyIfTable, secyRxSCStatsTable, secyStatsTable, secyTxSCStatsTable
- IEEE8023-LAG-MIB
 - dot3adAggPortDebugTable, dot3adAggPortListTable, dot3adAggPortStatsTable, dot3adAggPortTable, dot3adAggTable, dot3adTablesLastChanged
- IF-INVERTED-STACK-MIB
 - ifInvStackTable
- IF-MIB
 - ifNumber, ifStackLastChange, ifStackTable, ifTable, ifTableLastChange, ifXTable, linkDown, linkUp
- IGMP-STD-MIB
 - igmpCacheTable, igmpInterfaceTable
- IP-FORWARD-MIB
 - inetCidrRouteNumber, inetCidrRouteTable, ipCidrRouteNumber, ipCidrRouteTable
- IP-MIB
 - icmpInAddrMaskReps, icmpInAddrMasks, icmpInDestUnreaches, icmpInEchoReps, icmpInEchos, icmpInErrors, icmpInMsgs, icmpInParmProbs, icmpInRedirects, icmpInSrcQuenchs, icmpInTimeExcds, icmpInTimestampReps, icmpInTimestamps, icmpMsgStatsTable, icmpOutAddrMaskReps, icmpOutAddrMasks, icmpOutDestUnreaches, icmpOutEchoReps, icmpOutEchos, icmpOutErrors, icmpOutMsgs, icmpOutParmProbs, icmpOutRedirects, icmpOutSrcQuenchs, icmpOutTimeExcds, icmpOutTimestampReps, icmpOutTimestamps, icmpStatsTable, ipAddrTable, ipAddressPrefixTable, ipAddressTable, ipForwarding, ipNetToMediaTable, ipNetToPhysicalTable, ipv6IpForwarding
- IPMCAST-MIB
 - ipMcastDeviceConfigStorageType, ipMcastEnabled, ipMcastInterfaceTable, ipMcastRouteEntryCount, ipMcastRouteNextHopTable, ipMcastRouteTable
- IPMROUTE-STD-MIB
 - ipMRouteEnable, ipMRouteEntryCount, ipMRouteInterfaceTable, ipMRouteNextHopTable, ipMRouteTable
- ISIS-MIB
 - isisAdjacencyChange, isisAreaAddrTable, isisAreaMismatch, isisAttemptToExceedMaxSequence, isisAuthenticationTypeFailure, isisCircLevelTable, isisCircTable, isisDatabaseOverload, isisISAdjAreaAddrTable, isisISAdjIPAddrTable, isisISAdjProtSuppTable, isisISAdjTable, isisManAreaAddrTable, isisOwnLSPPurge, isisRejectedAdjacency, isisRouterTable, isisSequenceNumberSkip, isisSummAddrTable, isisSysAdminState, isisSysID, isisSysL2toL1Leaking, isisSysLevelTable, isisSysLevelType, isisSysMaxAge,

- isisSysMaxLSPGenInt, isisSysMaxPathSplits, isisSysNotificationEnable, isisSysPollESHelloRate, isisSysProtSupported, isisSysReceiveLSPBufferSize, isisSysVersion, isisSysWaitTime, isisSystemCounterTable
- LLDP-EXT-DOT1-MIB
 - lldpXdot1ConfigPortVlanTable, lldpXdot1LocTable, lldpXdot1RemProtoVlanTable, lldpXdot1RemProtocolTable, lldpXdot1RemTable, lldpXdot1RemVlanNameTable
- LLDP-EXT-DOT3-MIB
 - lldpXdot3LocLinkAggTable, lldpXdot3LocMaxFrameSizeTable, lldpXdot3PortConfigTable, lldpXdot3RemLinkAggTable, lldpXdot3RemMaxFrameSizeTable, lldpXdot3RemPortTable, lldpXdot3RemPowerTable
- LLDP-MIB
 - lldpConfigManAddrTable, lldpLocChassisId, lldpLocChassisIdSubtype, lldpLocManAddrTable, lldpLocPortTable, lldpLocSysCapEnabled, lldpLocSysCapSupported, lldpLocSysDesc, lldpLocSysName, lldpMessageTxHoldMultiplier, lldpMessageTxInterval, lldpNotificationInterval, lldpPortConfigTable, lldpReinitDelay, lldpRemManAddrTable, lldpRemOrgDefInfoTable, lldpRemTable, lldpRemTablesChange, lldpRemUnknownTLVTable, lldpStatsRemTablesAgeouts, lldpStatsRemTablesDeletes, lldpStatsRemTablesDrops, lldpStatsRemTablesInserts, lldpStatsRemTablesLastChangeTime, lldpStatsRxPortTable, lldpStatsTxPortTable, lldpTxDelay
- MAU-MIB
 - ifMauAutoNegTable, ifMauTable
- MEF-SOAM-PM-MIB
 - mefSoamDmCfgTable, mefSoamDmHistoryStatsXTable, mefSoamLmCfgTable, mefSoamLmHistoryStatsTable
- MPLS-LDP-STD-MIB
 - mplsLdpEntityTable, mplsLdpHelloAdjacencyTable, mplsLdpLsrId, mplsLdpPeerTable, mplsLdpSessionDown, mplsLdpSessionStatsTable, mplsLdpSessionTable, mplsLdpSessionUp
- MPLS-TE-STD-MIB
 - mplsTunnelHopTable, mplsTunnelPerfTable, mplsTunnelTable
- MSDP-MIB
 - msdpBackwardTransition, msdpCacheLifetime, msdpEnabled, msdpEstablished, msdpMeshGroupTable, msdpNumSACacheEntries, msdpPeerTable, msdpRPAddress, msdpSACacheTable
- NET-SNMP-AGENT-MIB
 - nsCacheDefaultTimeout, nsCacheEnabled, nsCacheTable, nsDebugDumpPdu, nsDebugEnabled, nsDebugOutputAll, nsDebugTokenTable, nsLoggingTable, nsModuleTable, nsNotifyRestart, nsTransactionTable
- NOTIFICATION-LOG-MIB
 - nlmConfigGlobalAgeOut, nlmConfigGlobalEntryLimit, nlmLogTable, nlmLogVariableTable, nlmStatsGlobalNotificationsBumped, nlmStatsGlobalNotificationsLogged

- OSPF-MIB
 - ospfASBdrRtrStatus, ospfAdminStat, ospfAreaBdrRtrStatus, ospfAreaTable, ospfAsLsdbTable, ospfDemandExtensions, ospfExitOverflowInterval, ospfExtLsdbLimit, ospfExtLsdbTable, ospfExternLsaChecksumSum, ospfExternLsaCount, ospfLsdbTable, ospfMulticastExtensions, ospfNbrTable, ospfOpaqueLsaSupport, ospfOriginateNewLsas, ospfRFC1583Compatibility, ospfRouterId, ospfRxNewLsas, ospfTOSSupport, ospfVersionNumber
- OSPF-TRAP-MIB
 - ospfIfAuthFailure, ospfIfConfigError, ospfIfStateChange, ospfNbrStateChange
- OSPFV3-MIB
 - ospfv3AdminStatus, ospfv3AreaAggregateTable, ospfv3AreaBdrRtrStatus, ospfv3AreaLsdbTable, ospfv3AreaTable, ospfv3AsLsdbTable, ospfv3AsScopeLsaCount, ospfv3ExtLsaCount, ospfv3HostTable, ospfv3IfConfigError, ospfv3IfRxBadPacket, ospfv3IfStateChange, ospfv3IfTable, ospfv3LinkLsdbTable, ospfv3NbrRestartHelperStatusChange, ospfv3NbrStateChange, ospfv3NbrTable, ospfv3NssaTranslatorStatusChange, ospfv3OriginateNewLsas, ospfv3ReferenceBandwidth, ospfv3RestartAge, ospfv3RestartExitReason, ospfv3RestartInterval, ospfv3RestartStatus, ospfv3RestartStatusChange, ospfv3RestartSupport, ospfv3RestartTime, ospfv3RouterId, ospfv3RxNewLsas, ospfv3VersionNumber, ospfv3VirtIfTable, ospfv3VirtNbrTable
- PIM-MIB
 - pimCandidateRPTTable, pimComponentTable, pimInterfaceTable, pimIpMRouteNextHopTable, pimIpMRouteTable, pimJoinPruneInterval, pimNeighborLoss, pimNeighborTable, pimRPSetTable
- PTPBASE-MIB
 - ptpbaseClockCurrentDSTable, ptpbaseClockDefaultDSTable, ptpbaseClockParentDSTable, ptpbaseClockPortDSTable, ptpbaseClockPortRunningTable, ptpbaseClockPortTable, ptpbaseClockPortTransDSTable, ptpbaseClockRunningTable, ptpbaseClockTimePropertiesDSTable, ptpbaseClockTransDefaultDSTable, ptpbaseSystemDomainTable, ptpbaseSystemTable
- Q-BRIDGE-MIB
 - dot1qFdbTable, dot1qGvrpStatus, dot1qMaxSupportedVlans, dot1qMaxVlanId, dot1qNextFreeLocalVlanIndex, dot1qNumVlans, dot1qPortVlanTable, dot1qTpFdbTable, dot1qVlanCurrentTable, dot1qVlanNumDeletes, dot1qVlanStaticTable, dot1qVlanVersionNumber
- RMON-MIB
 - etherStatsTable
- RMON2-MIB
 - etherStats2Table, probeCapabilities
- SNMP-FRAMEWORK-MIB
 - snmpEngineBoots, snmpEngineID, snmpEngineMaxMessageSize, snmpEngineTime
- SNMP-MPD-MIB
 - snmpInvalidMsgs, snmpUnknownPDUHandlers, snmpUnknownSecurityModels

- SNMP-USER-BASED-SM-MIB
 - usmStatsDecryptionErrors, usmStatsNotInTimeWindows, usmStatsUnknownEngineIDs, usmStatsUnknownUserNames, usmStatsUnsupportedSecLevels, usmStatsWrongDigests, usmUserSpinLock, usmUserTable
- SNMP-VIEW-BASED-ACM-MIB
 - vacmAccessTable, vacmContextTable, vacmSecurityToGroupTable, vacmViewSpinLock, vacmViewTreeFamilyTable
- SNMPv2-MIB
 - authenticationFailure, coldStart, snmpEnableAuthenTraps, snmpEnableAuthenTraps.0, snmpInASNParseErrs, snmpInBadCommunityNames, snmpInBadCommunityUses, snmpInBadValues, snmpInBadVersions, snmpInGenErrs, snmpInGetNexts, snmpInGetRequests, snmpInGetResponses, snmpInNoSuchNames, snmpInPkts, snmpInReadOnlys, snmpInSetRequests, snmpInTooBigs, snmpInTotalReqVars, snmpInTotalSetVars, snmpInTraps, snmpOutBadValues, snmpOutGenErrs, snmpOutGetNexts, snmpOutGetRequests, snmpOutGetResponses, snmpOutNoSuchNames, snmpOutPkts, snmpOutSetRequests, snmpOutTooBigs, snmpOutTraps, snmpProxyDrops, snmpSetSerialNo, snmpSilentDrops, sysContact, sysDescr, sysLocation, sysName, sysORLastChange, sysORTable, sysObjectID, sysServices, sysUpTime
- TCP-MIB
 - tcpActiveOpens, tcpAttemptFails, tcpConnTable, tcpConnectionTable, tcpCurrEstab, tcpEstabResets, tcpHCInSegs, tcpHCOutSegs, tcpInErrs, tcpInSegs, tcpListenerTable, tcpMaxConn, tcpOutRsts, tcpOutSegs, tcpPassiveOpens, tcpRetransSegs, tcpRtoAlgorithm, tcpRtoMax, tcpRtoMin
- TOKEN-RING-RMON-MIB
- TUNNEL-MIB
 - tunnelIfTable, tunnelInetConfigTable
- UCD-SNMP-MIB
 - laTable, ssCpuIdle, ssCpuNumCpus, ssCpuRawGuest, ssCpuRawGuestNice, ssCpuRawIdle, ssCpuRawInterrupt, ssCpuRawKernel, ssCpuRawNice, ssCpuRawSoftIRQ, ssCpuRawSteal, ssCpuRawSystem, ssCpuRawUser, ssCpuRawWait, ssCpuSystem, ssCpuUser, ssErrorName, ssIORawReceived, ssIORawSent, ssIOReceive, ssIOSent, ssIndex, ssRawContexts, ssRawInterrupts, ssRawSwapIn, ssRawSwapOut, ssSwapIn, ssSwapOut, ssSysContext, ssSysInterrupts
- UDP-MIB
 - udpEndpointTable, udpHCInDatagrams, udpHCOutDatagrams, udpInDatagrams, udpInErrors, udpNoPorts, udpOutDatagrams, udpTable
- VRRP-MIB
 - vrrpAssoIpAddrTable, vrrpNodeVersion, vrrpNotificationCntl, vrrpOperTable, vrrpTrapNewMaster

SNMP Traps

- ARISTA-BGP4V2-MIB
 - aristaBgp4V2BackwardTransitionNotification, aristaBgp4V2EstablishedNotification
- ARISTA-BRIDGE-EXT-MIB
 - aristaMacAge, aristaMacLearn, aristaMacMove
- ARISTA-CLB-MIB
 - aristaClbPortGroupAllocatedFlowAtOrAboveLimit, aristaClbPortGroupAllocatedFlowAtOrAboveThresholdWarning, aristaClbPortGroupAllocatedFlowBelowLimit, aristaClbPortGroupAllocatedFlowBelowThresholdWarning
- ARISTA-CONFIG-MAN-MIB
 - aristaConfigManEvent
- ARISTA-ENTITY-SENSOR-MIB
 - aristaEntSensorAlarm
- ARISTA-EXTERNAL-ALARM-MIB
 - aristaExternalAlarmAssertedNotif, aristaExternalAlarmDeassertedNotif
- ARISTA-HARDWARE-UTILIZATION-MIB
 - aristaHardwareUtilizationAlert
- ARISTA-REDUNDANCY-MIB
 - aristaRedundancySwitchOverNotif
- ARISTA-TEST-MIB
 - aristaTestNotification
- BGP4-MIB
 - bgpBackwardTransNotification, bgpEstablishedNotification
- ENTITY-MIB
 - entConfigChange
- ENTITY-STATE-MIB
 - entStateOperDisabled, entStateOperEnabled
- IF-MIB
 - linkDown, linkUp
- ISIS-MIB
 - isisAdjacencyChange, isisAreaMismatch, isisAttemptToExceedMaxSequence, isisAuthenticationTypeFailure, isisDatabaseOverload, isisOwnLSPPurge, isisRejectedAdjacency, isisSequenceNumberSkip
- LLDP-MIB
 - lldpRemTablesChange
- MPLS-LDP-STD-MIB
 - mplsLdpSessionDown, mplsLdpSessionUp
- MSDP-MIB
 - msdpBackwardTransition, msdpEstablished
- NET-SNMP-AGENT-MIB
 - nsNotifyRestart
- OSPF-TRAP-MIB
 - ospfIfAuthFailure, ospfIfConfigError, ospfIfStateChange, ospfNbrStateChange

- OSPFV3-MIB
 - ospfv3IfConfigError, ospfv3IfRxBadPacket, ospfv3IfStateChange, ospfv3NbrRestartHelperStatusChange, ospfv3NbrStateChange, ospfv3NssaTranslatorStatusChange, ospfv3RestartStatusChange
- PIM-MIB
 - pimNeighborLoss
- SNMPv2-MIB
 - authenticationFailure, coldStart
- VRRP-MIB
 - vrrpTrapNewMaster

Supported Hardware

Fixed System

* 7326-56X-O-AC	* DCS-7130-48EHS-R (HwRev:B1)	* DCS-7280CR3MK-32D4S-R
* 7326-56X-O-AC-F	* DCS-7130-48EHS-R (HwRev:B2)	* DCS-7280CR3MK-32P4
* 7726-32X-O-AC	* DCS-7130-48G3	* DCS-7280CR3MK-32P4-F
* 7726-32X-O-AC-F	* DCS-7130-48G3-F (HwRev:B1)	* DCS-7280CR3MK-32P4-R
* AWE-5300-550-A-PS	* DCS-7130-48G3-F (HwRev:B2)	* DCS-7280CR3MK-32P4S
* AWE-5310	* DCS-7130-48G3-R (HwRev:B1)	* DCS-7280CR3MK-32P4S-F
* AWE-5310-F	* DCS-7130-48G3-R (HwRev:B2)	* DCS-7280CR3MK-32P4S-R
* AWE-5500-800-A-PS	* DCS-7130-48G3S	* DCS-7280DR3-24
* AWE-5500-A-FAN	* DCS-7130-48G3S-F (HwRev:C1)	* DCS-7280DR3-24-F
* AWE-5510	* DCS-7130-48G3S-R (HwRev:C1)	* DCS-7280DR3-24-M-F
* AWE-5510-F	* DCS-7130-48L	* DCS-7280DR3A-36
* AWE-7220RP-5TH-2S	* DCS-7130-48L-F (HwRev:A2)	* DCS-7280DR3A-36-F
* AWE-7220RP-5TH-2S-F	* DCS-7130-48L-F (HwRev:A3)	* DCS-7280DR3A-54
* AWE-7230R-4TX-4S	* DCS-7130-48L-R (HwRev:A2)	* DCS-7280DR3A-54-F
* AWE-7230R-4TX-4S-F	* DCS-7130-48L-R (HwRev:A3)	* DCS-7280DR3AK-36
* AWE-7250R-16S	* DCS-7130-48LA	* DCS-7280DR3AK-36-F
* AWE-7250R-16S-F	* DCS-7130-48LA-F (HwRev:A2)	* DCS-7280DR3AK-36S
* CCS-710P-12	* DCS-7130-48LA-F (HwRev:A3)	* DCS-7280DR3AK-36S-F
* CCS-710P-16P	* DCS-7130-48LA-R (HwRev:A2)	* DCS-7280DR3AK-54
* CCS-720DF-48Y	* DCS-7130-48LA-R (HwRev:A3)	* DCS-7280DR3AK-54-F
* CCS-720DF-48Y-2	* DCS-7130-48LAS	* DCS-7280DR3AM-36
* CCS-720DF-48Y-2F	* DCS-7130-48LAS-F (HwRev:B2)	* DCS-7280DR3AM-36-F
* CCS-720DF-48Y-DC	* DCS-7130-48LAS-F (HwRev:B3)	* DCS-7280DR3AM-54
* CCS-720DF-48Y-DC-2	* DCS-7130-48LAS-R (HwRev:B2)	* DCS-7280DR3AM-54-F
* CCS-720DF-48Y-DC-2F	* DCS-7130-48LAS-R (HwRev:B3)	* DCS-7280DR3K-24
* CCS-720DF-48Y-DC-F	* DCS-7130-48LB	* DCS-7280DR3K-24-F
* CCS-720DF-48Y-F	* DCS-7130-48LB-F (HwRev:A2)	* DCS-7280PR3-24
* CCS-720DF-48Y-M-S-2F	* DCS-7130-48LB-F (HwRev:A3)	* DCS-7280PR3-24-F

* CCS-720DF-48Y-S-2	* DCS-7130-48LB-R (HwRev:A2)	* DCS-7280PR3-24-M-F
* CCS-720DP-24S	* DCS-7130-48LB-R (HwRev:A3)	* DCS-7280PR3K-24
* CCS-720DP-24S-2	* DCS-7130-48LBA	* DCS-7280PR3K-24-F
* CCS-720DP-24S-2F	* DCS-7130-48LBA-F (HwRev:A2)	* DCS-7280SR3-40YC6
* CCS-720DP-24S-F	* DCS-7130-48LBA-F (HwRev:A3)	* DCS-7280SR3-40YC6-F
* CCS-720DP-24S-M-S-2F	* DCS-7130-48LBA-R (HwRev:A2)	* DCS-7280SR3-40YC6-R
* CCS-720DP-24S-S-2	* DCS-7130-48LBA-R (HwRev:A3)	* DCS-7280SR3-48YC8
* CCS-720DP-24ZS	* DCS-7130-48LBAS	* DCS-7280SR3-48YC8-F
* CCS-720DP-24ZS-2	* DCS-7130-48LBAS-F (HwRev:B2)	* DCS-7280SR3-48YC8-R
* CCS-720DP-24ZS-2F	* DCS-7130-48LBAS-F (HwRev:B3)	* DCS-7280SR3A-48YC8
* CCS-720DP-24ZS-F	* DCS-7130-48LBAS-R (HwRev:B2)	* DCS-7280SR3A-48YC8-F
* CCS-720DP-48S	* DCS-7130-48LBAS-R (HwRev:B3)	* DCS-7280SR3A-48YC8-R
* CCS-720DP-48S-2	* DCS-7130-48LBS	* DCS-7280SR3AK-48YC8
* CCS-720DP-48S-2F	* DCS-7130-48LBS-F (HwRev:B2)	* DCS-7280SR3AK-48YC8-F
* CCS-720DP-48S-F	* DCS-7130-48LBS-F (HwRev:B3)	* DCS-7280SR3AK-48YC8-R
* CCS-720DP-48S-M-S-2F	* DCS-7130-48LBS-R (HwRev:B2)	* DCS-7280SR3AM-48YC8
* CCS-720DP-48S-S-2	* DCS-7130-48LBS-R (HwRev:B3)	* DCS-7280SR3AM-48YC8-F
* CCS-720DP-48ZS	* DCS-7130-48LS	* DCS-7280SR3AM-48YC8-R
* CCS-720DP-48ZS-2	* DCS-7130-48LS-F (HwRev:B2)	* DCS-7280SR3E-40YC6
* CCS-720DP-48ZS-2F	* DCS-7130-48LS-F (HwRev:B3)	* DCS-7280SR3E-40YC6-F
* CCS-720DP-48ZS-F	* DCS-7130-48LS-R (HwRev:B2)	* DCS-7280SR3E-40YC6-M-F
* CCS-720DT-24S	* DCS-7130-48LS-R (HwRev:B3)	* DCS-7280SR3E-40YC6-M-R
* CCS-720DT-24S-2	* DCS-7130-96	* DCS-7280SR3E-40YC6-R
* CCS-720DT-24S-2F	* DCS-7130-96-F (HwRev:A7)	* DCS-7280SR3K-48YC8
* CCS-720DT-24S-2R	* DCS-7130-96-F (HwRev:A8)	* DCS-7280SR3K-48YC8-F
* CCS-720DT-24S-F	* DCS-7130-96-R (HwRev:A7)	* DCS-7280SR3K-48YC8-R
* CCS-720DT-24S-M-S-2F	* DCS-7130-96-R (HwRev:A8)	* DCS-7280SR3K-48YC8A
* CCS-720DT-24S-S-2	* DCS-7130-96L	* DCS-7280SR3K-48YC8A-F
* CCS-720DT-48S-2	* DCS-7130-96L-F (HwRev:A2)	* DCS-7280SR3K-48YC8A-R
* CCS-720DT-48S-2F	* DCS-7130-96L-F (HwRev:A3)	* DCS-7280SR3M-48YC8
* CCS-720DT-48S-2R	* DCS-7130-96L-R (HwRev:A2)	* DCS-7280SR3M-48YC8-F
* CCS-720XP-24Y6	* DCS-7130-96L-R (HwRev:A3)	* DCS-7280SR3M-48YC8-R
* CCS-720XP-24Y6-F	* DCS-7130-96LA	* DCS-7280SR3MK-48YC8A-S
* CCS-720XP-24Y6-R	* DCS-7130-96LA-F (HwRev:A2)	* DCS-7280SR3MK-48YC8A-S-F
* CCS-720XP-24ZY4	* DCS-7130-96LA-F (HwRev:A3)	* DCS-7280SR3MK-48YC8A-S-R
* CCS-720XP-24ZY4-F	* DCS-7130-96LA-R (HwRev:A2)	* DCS-7280TR3-40C6
* CCS-720XP-48TXH-2C-S	* DCS-7130-96LA-R (HwRev:A3)	* DCS-7280TR3-40C6-F
* CCS-720XP-48TXH-2C-S-F	* DCS-7130-96LAS	* DCS-7280TR3-40C6-R
* CCS-720XP-48Y6	* DCS-7130-96LAS-F (HwRev:B1)	* FAN-7000
* CCS-720XP-48Y6-F	* DCS-7130-96LAS-F (HwRev:B2)	* FAN-7000-F
* CCS-720XP-48Y6-R	* DCS-7130-96LAS-R (HwRev:B1)	* FAN-7000-R

* CCS-720XP-48ZC2	* DCS-7130-96LAS-R (HwRev:B2)	* FAN-7000H-R
* CCS-720XP-48ZC2-F	* DCS-7130-96LB	* FAN-7001C
* CCS-720XP-96ZC2	* DCS-7130-96LB-F (HwRev:A2)	* FAN-7001C-F
* CCS-720XP-96ZC2-F	* DCS-7130-96LB-F (HwRev:A3)	* FAN-7001C-R
* CCS-720XP-96ZC2-M-S-F	* DCS-7130-96LB-R (HwRev:A2)	* FAN-7001D
* CCS-720XP-96ZC2-S	* DCS-7130-96LB-R (HwRev:A3)	* FAN-7001D-F
* CCS-720XPM-48TH-6SY	* DCS-7130-96LBA	* FAN-7001DH
* CCS-720XPM-48TH-6SY-F	* DCS-7130-96LBA-F (HwRev:A2)	* FAN-7001DH-F
* CCS-722XPM-48Y4	* DCS-7130-96LBA-F (HwRev:A3)	* FAN-7002H
* CCS-722XPM-48Y4-F	* DCS-7130-96LBA-R (HwRev:A2)	* FAN-7002H-R
* CCS-722XPM-48ZY8	* DCS-7130-96LBA-R (HwRev:A3)	* FAN-7010
* CCS-722XPM-48ZY8-F	* DCS-7130-96LBAS	* FAN-7010X
* DCS-7010TX-48	* DCS-7130-96LBAS-F (HwRev:B1)	* FAN-7011H
* DCS-7010TX-48-DC	* DCS-7130-96LBAS-R (HwRev:B1)	* FAN-7011H-F
* DCS-7010TX-48-DC-F	* DCS-7130-96LBS	* FAN-7011H-R
* DCS-7010TX-48-DC-R	* DCS-7130-96LBS-F (HwRev:B1)	* FAN-7011M
* DCS-7010TX-48-F	* DCS-7130-96LBS-F (HwRev:B2)	* FAN-7011M-F
* DCS-7010TX-48-R	* DCS-7130-96LBS-R (HwRev:B1)	* FAN-7011M-R
* DCS-7010TX-48C	* DCS-7130-96LBS-R (HwRev:B2)	* FAN-7012H
* DCS-7010TX-48C-DC	* DCS-7130-96LS	* FAN-7012H-BLUE
* DCS-7010TX-48C-DC-RV3	* DCS-7130-96LS-F (HwRev:B1)	* FAN-7012H-RED
* DCS-7010TX-48C-F	* DCS-7130-96LS-F (HwRev:B2)	* FAN-7012M
* DCS-7010TX-48C-R	* DCS-7130-96LS-R (HwRev:B1)	* FAN-7012M-BLUE
* DCS-7050CX3-32C	* DCS-7130-96LS-R (HwRev:B2)	* FAN-7012M-RED
* DCS-7050CX3-32C-F	* DCS-7130-96S	* FAN-7021H
* DCS-7050CX3-32C-R	* DCS-7130-96S-F (HwRev:B1)	* FAN-7021H-RED
* DCS-7050CX3-32S	* DCS-7130-96S-R (HwRev:B1)	* FAN-7130-1U
* DCS-7050CX3-32S-F	* DCS-7130B-32QD	* FAN-7130-1U-F
* DCS-7050CX3-32S-R	* DCS-7130B-32QD-F	* FAN-7130-1U-R
* DCS-7050CX3-32S-SSD-F	* DCS-7130B-32QD-R	* FAN-7130-2U
* DCS-7050CX3-32S-SSD-R	* DCS-7130LBR-48S6QD	* FAN-7130-2U-F
* DCS-7050CX3M-32S	* DCS-7130LBR-48S6QD-F	* FAN-7130-2U-R
* DCS-7050CX3M-32S-F	* DCS-7130LBR-48S6QD-R	* FAN-7202
* DCS-7050CX3M-32S-R	* DCS-7132LB-48Y4C	* FAN-7202-RED
* DCS-7050CX4-24D8	* DCS-7132LB-48Y4C-F	* NIM-1QC
* DCS-7050CX4-24D8-F	* DCS-7132LB-48Y4C-R	* NIM-4S
* DCS-7050CX4-24D8-R	* DCS-7135LB-48Y4C	* PWR-1011-AC
* DCS-7050CX4M-48D8	* DCS-7135LB-48Y4C-F	* PWR-1011-AC-BLUE
* DCS-7050CX4M-48D8-F	* DCS-7135LB-48Y4C-R	* PWR-1011-AC-RED
* DCS-7050CX4M-48D8-R	* DCS-7170-32C	* PWR-1011-DC
* DCS-7050DX4-32S	* DCS-7170-32C-F	* PWR-1011-DC-BLUE

* DCS-7050DX4-32S-F	* DCS-7170-32C-M-F	* PWR-1011-DC-RED
* DCS-7050DX4M-32S	* DCS-7170-32C-M-R	* PWR-1021-AC
* DCS-7050DX4M-32S-F	* DCS-7170-32C-R	* PWR-1021-AC-RED
* DCS-7050PX4-32S	* DCS-7170-32CD	* PWR-1100AC
* DCS-7050PX4-32S-F	* DCS-7170-32CD-F	* PWR-1100AC-F
* DCS-7050SDX4-48D8	* DCS-7170-32CD-R	* PWR-1100AC-R
* DCS-7050SDX4-48D8-F	* DCS-7170-64C	* PWR-1511-AC
* DCS-7050SDX4-48D8-R	* DCS-7170-64C-F	* PWR-1511-AC-BLUE
* DCS-7050SPX4-48D8	* DCS-7170-64C-M-F	* PWR-1511-AC-RED
* DCS-7050SPX4-48D8-F	* DCS-7170-64C-M-R	* PWR-1511-DC
* DCS-7050SPX4-48D8-R	* DCS-7170-64C-R	* PWR-1511-DC-BLUE
* DCS-7050SX3-24YC4C-S	* DCS-7170B-64C	* PWR-1511-DC-RED
* DCS-7050SX3-24YC4C-S-F	* DCS-7170B-64C-F	* PWR-1512-AC
* DCS-7050SX3-24YC4C-S-R	* DCS-7260CX3-64	* PWR-1512-AC-BLUE
* DCS-7050SX3-48C8	* DCS-7260CX3-64-F	* PWR-1512-AC-RED
* DCS-7050SX3-48C8-F	* DCS-7260CX3-64-R	* PWR-1513-AC
* DCS-7050SX3-48C8-R	* DCS-7260CX3-64E	* PWR-1513-AC-BLUE
* DCS-7050SX3-48C8C	* DCS-7260CX3-64E-F	* PWR-1513-AC-RED
* DCS-7050SX3-48C8C-F	* DCS-7260CX3-64E-R	* PWR-1600AC
* DCS-7050SX3-48C8C-R	* DCS-7260CX3-64LQ	* PWR-1600AC-F
* DCS-7050SX3-48YC12	* DCS-7260CX3-64LQ-F	* PWR-1611-AC
* DCS-7050SX3-48YC12-F	* DCS-7260CX3-64LQ-R	* PWR-1611-AC-RED
* DCS-7050SX3-48YC8	* DCS-7280CR3-32D4	* PWR-1611-DC
* DCS-7050SX3-48YC8-F	* DCS-7280CR3-32D4-F	* PWR-1611-DC-RED
* DCS-7050SX3-48YC8-R	* DCS-7280CR3-32D4-M-F	* PWR-1900-DC
* DCS-7050SX3-48YC8C	* DCS-7280CR3-32D4-M-R	* PWR-1900-DC-F
* DCS-7050SX3-48YC8C-F	* DCS-7280CR3-32D4-R	* PWR-1900-DC-R
* DCS-7050SX3-48YC8C-R	* DCS-7280CR3-32P4	* PWR-1900AC
* DCS-7050SX3-96YC8	* DCS-7280CR3-32P4-F	* PWR-1900AC-F
* DCS-7050SX3-96YC8-F	* DCS-7280CR3-32P4-M-F	* PWR-1900AC-R
* DCS-7050SX3-96YC8-R	* DCS-7280CR3-32P4-M-R	* PWR-2021-AC
* DCS-7050TX3-48C8	* DCS-7280CR3-32P4-R	* PWR-2021-AC-RED
* DCS-7050TX3-48C8-F	* DCS-7280CR3-36S	* PWR-2411-AC
* DCS-7050TX3-48C8-R	* DCS-7280CR3-36S-F	* PWR-2411-AC-BLUE
* DCS-7050X4-48Y-4DF	* DCS-7280CR3-36S-R	* PWR-2411-AC-RED
* DCS-7050X4-48Y-4DF-F	* DCS-7280CR3-96	* PWR-2411-DC
* DCS-7050X4-48Y-4DF-R	* DCS-7280CR3-96-F	* PWR-2411-DC-BLUE
* DCS-7060CX-32C	* DCS-7280CR3A-24D12	* PWR-2411-DC-RED
* DCS-7060CX-32C-F	* DCS-7280CR3A-24D12-F	* PWR-2411-MC-RED
* DCS-7060CX-32C-R	* DCS-7280CR3A-24D12-R	* PWR-2411C
* DCS-7060CX-32S	* DCS-7280CR3A-32S	* PWR-3001-AC

* DCS-7060CX-32S-F	* DCS-7280CR3A-32S-F	* PWR-3001-AC-RED
* DCS-7060CX-32S-F (HwRev:13.20)	* DCS-7280CR3A-32S-R	* PWR-3001-DC
* DCS-7060CX-32S-R	* DCS-7280CR3A-48D6	* PWR-3001-DC-RED
* DCS-7060CX-32S-R (HwRev:13.20)	* DCS-7280CR3A-48D6-F	* PWR-400-DC
* DCS-7060CX2-32S	* DCS-7280CR3A-72	* PWR-400-DC-BLUE
* DCS-7060CX2-32S-F	* DCS-7280CR3A-72-F	* PWR-400-DC-RED
* DCS-7060CX2-32S-R	* DCS-7280CR3A-72-R	* PWR-400AC
* DCS-7060CX5-56D8	* DCS-7280CR3AK-24D12	* PWR-400AC-BLUE
* DCS-7060CX5-56D8-F	* DCS-7280CR3AK-24D12-F	* PWR-400AC-RED
* DCS-7060CX5-56D8-R	* DCS-7280CR3AK-24D12-R	* PWR-460AC
* DCS-7060DX4-32	* DCS-7280CR3AK-24D12-S	* PWR-460AC-F
* DCS-7060DX4-32-F	* DCS-7280CR3AK-24D12-S-F	* PWR-460AC-R
* DCS-7060DX4-32C-RV3	* DCS-7280CR3AK-24D12-S-R	* PWR-460DC
* DCS-7060DX4-32C-RV3-F	* DCS-7280CR3AK-32S	* PWR-460DC-F
* DCS-7060DX4-32S	* DCS-7280CR3AK-32S-F	* PWR-460DC-R
* DCS-7060DX4-32S-F	* DCS-7280CR3AK-32S-R	* PWR-500-DC
* DCS-7060DX4-32S-MON	* DCS-7280CR3AK-48D6	* PWR-500-DC-F
* DCS-7060DX4-32S-MON2	* DCS-7280CR3AK-48D6-F	* PWR-500-DC-R
* DCS-7060DX4-32S-PEM	* DCS-7280CR3AK-72	* PWR-500AC
* DCS-7060DX4-32SB-MON3	* DCS-7280CR3AK-72-F	* PWR-500AC-F
* DCS-7060DX4-32SB-RV3	* DCS-7280CR3AK-72-R	* PWR-500AC-R
* DCS-7060DX4-32SB-RV3-F	* DCS-7280CR3AM-24D12	* PWR-501AC
* DCS-7060DX4-32SBON3	* DCS-7280CR3AM-24D12-F	* PWR-501AC-F
* DCS-7060DX4-32SON	* DCS-7280CR3AM-24D12-R	* PWR-501AC-R
* DCS-7060DX4-32SON2	* DCS-7280CR3AM-32S	* PWR-511-AC
* DCS-7060DX5-32	* DCS-7280CR3AM-32S-F	* PWR-511-AC-BLUE
* DCS-7060DX5-32-F	* DCS-7280CR3AM-32S-R	* PWR-511-AC-RED
* DCS-7060DX5-32-R	* DCS-7280CR3AM-48D6	* PWR-511-DC
* DCS-7060DX5-64	* DCS-7280CR3AM-48D6-F	* PWR-511-DC-BLUE
* DCS-7060DX5-64-F	* DCS-7280CR3AM-72	* PWR-511-DC-RED
* DCS-7060DX5-64E	* DCS-7280CR3AM-72-F	* PWR-561-AC
* DCS-7060DX5-64E-F	* DCS-7280CR3AM-72-R	* PWR-561-AC-RED
* DCS-7060DX5-64S	* DCS-7280CR3K-32D4	* PWR-621-AC
* DCS-7060DX5-64S-F	* DCS-7280CR3K-32D4-F	* PWR-621-AC-BLUE
* DCS-7060DX5-64S-F (HwRev: 12.00)	* DCS-7280CR3K-32D4-R	* PWR-621-AC-RED
* DCS-7060DX5-64S-R	* DCS-7280CR3K-32D4A	* PWR-721-DC
* DCS-7060DX5-64S-R (HwRev: 12.00)	* DCS-7280CR3K-32D4A-F	* PWR-721-DC-RED
* DCS-7060PX4-32	* DCS-7280CR3K-32D4A-R	* PWR-745AC
* DCS-7060PX4-32-F	* DCS-7280CR3K-32P4	* PWR-745AC-F
* DCS-7060PX5-64E	* DCS-7280CR3K-32P4-F	* PWR-745AC-R
* DCS-7060PX5-64E-F	* DCS-7280CR3K-32P4-R	* PWR-747AC

* DCS-7060SX2-48YC6	* DCS-7280CR3K-32P4A	* PWR-747AC-F
* DCS-7060SX2-48YC6-F	* DCS-7280CR3K-32P4A-F	* PWR-747AC-R
* DCS-7060SX2-48YC6-R	* DCS-7280CR3K-32P4A-R	* PWR-750AC
* DCS-7060X6-32PE	* DCS-7280CR3K-36A	* PWR-750AC-F
* DCS-7060X6-32PE-F	* DCS-7280CR3K-36A-F	* PWR-750AC-R
* DCS-7060X6-32PE-N	* DCS-7280CR3K-36A-R	* PWR-861-AC
* DCS-7060X6-64PE	* DCS-7280CR3K-36S	* PWR-861-AC-RED
* DCS-7060X6-64PE-F	* DCS-7280CR3K-36S-F	* ZTX-7230S-4TX-4S
* DCS-7130-16G3	* DCS-7280CR3K-36S-R	* ZTX-7230S-4TX-4S-F
* DCS-7130-16G3-F (HwRev:C1)	* DCS-7280CR3K-96	* ZTX-7250S-16S
* DCS-7130-16G3-R (HwRev:C1)	* DCS-7280CR3K-96-F	* ZTX-7250S-16S-F
* DCS-7130-16G3S	* DCS-7280CR3MK-32D4	* FAN-7000H
* DCS-7130-16G3S-F (HwRev:D1)	* DCS-7280CR3MK-32D4-F	* FAN-7000H-F
* DCS-7130-16G3S-R (HwRev:D1)	* DCS-7280CR3MK-32D4-R	* FAN-7002
* DCS-7130-48EH	* DCS-7280CR3MK-32D4A-S	* FAN-7002-F
* DCS-7130-48EH-F (HwRev:A4)	* DCS-7280CR3MK-32D4A-S-F	* FAN-7002-R
* DCS-7130-48EH-R (HwRev:A4)	* DCS-7280CR3MK-32D4A-S-R	* FAN-7002H-F
* DCS-7130-48EHS	* DCS-7280CR3MK-32D4S	* FAN-7012MP
* DCS-7130-48EHS-F (HwRev:B1)	* DCS-7280CR3MK-32D4S-F	* FAN-7012MP-RED
* DCS-7130-48EHS-F (HwRev:B2)		

Modular

* FAN-7000H	* 7388X5-SC	* CCS-758-CH
* FAN-7000H-F	* 7500R3-24D-LC	* CCS-758-X3-SC
* FAN-7002	* 7500R3-24P-LC	* DCS-7001-SUP-A
* FAN-7002-F	* 7500R3-36CQ-LC	* DCS-7304
* FAN-7002-R	* 7500R3-48Y4D-LC	* DCS-7308
* FAN-7002H-F	* 7500R3K-36CQ-LC	* DCS-7500-SUP2
* FAN-7012MP	* 7500R3K-48Y4D-LC	* DCS-7500-SUP2-D
* FAN-7012MP-RED	* 7504R3-FM	* DCS-7504
* 7289	* 7508R3-FM	* DCS-7504N
* 7289-SUP	* 7512R3-FM	* DCS-7508
* 7289-SUP-D	* 7800R3-36D-LC	* DCS-7508N
* 7289-SUP-S	* 7800R3-36P-LC	* DCS-7512N
* 7289-SUP-S-D	* 7800R3-48CQ-LC	* DCS-7800-SUP
* 7289R3A-SC	* 7800R3-48CQ2-LC	* DCS-7800-SUP1A
* 7289R3AK-SC	* 7800R3-48CQM-LC	* DCS-7800-SUP1S
* 7289R3AM-SC	* 7800R3-48CQM2-LC	* DCS-7804-CH
* 7300-SUP	* 7800R3-48CQMS-LC	* DCS-7808-CH
* 7300-SUP-D	* 7800R3A-36D-LC	* DCS-7812-CH
* 7300-SUP2	* 7800R3A-36D2-LC	* DCS-7816-CH
* 7300-SUP2-D	* 7800R3A-36DM-LC	* DCS-7816-SUP
* 7300X3-32C-LC	* 7800R3A-36DM2-LC	* DCS-7816-SUP1S
* 7300X3-48TC4-LC	* 7800R3A-36P-LC	* DCS-7816L-CH
* 7300X3-48YC4-LC	* 7800R3A-36PM-LC	* FAN-752M
* 7304X3-FM	* 7800R3AK-36DM-LC	* FAN-752M-RED
* 7304X3-FM2	* 7800R3AK-36DM2-LC	* PWR-2422-HV
* 7308X3-FM	* 7800R3AK-36PM-LC	* PWR-2422-HV-RED
* 7308X3-FM2	* 7800R3K-36DM-LC	* PWR-2700-DC
* 7358	* 7800R3K-48CQ-LC	* PWR-2700-DC-F
* 7358-16C	* 7800R3K-72Y-LC	* PWR-2700-DC-R
* 7358X4-SC	* 7804R3-FM	* PWR-2900AC
* 7368	* 7808R3-FM	* PWR-3351-AC
* 7368-16C	* 7808R3-FM2	* PWR-3351-AC-RED
* 7368-16S	* 7812R3-FM	* PWR-3K-AC
* 7368-4D	* 7816LR3-FM	* PWR-3K-AC-F
* 7368-4P	* 7816R3-FM	* PWR-3K-AC-R
* 7368-SUP	* CCS-750-SUP100	* PWR-3K-DC
* 7368-SUP-D	* CCS-750-SUP25	* PWR-3K-DC-BLUE
* 7368X4-SC	* CCS-750X-48SX-LC	* PWR-3K-DC-RED
* 7388	* CCS-750X-48THP-LC	* PWR-3KT-AC

* 7388-16CD	* CCS-750X-48TP-LC	* PWR-3KT-AC-BLUE
* 7388-16CD2	* CCS-750X-48ZP-LC	* PWR-3KT-AC-RED
* 7388-8D	* CCS-750X-48ZXP-LC	* PWR-D1-3041-AC
* 7388-8DR	* CCS-750XPM-48TXH-LC	* PWR-D1-3041-AC-BLUE
* 7388-SUP	* CCS-755-CH	* PWR-D4-3041-AC
* 7388-SUP-D	* CCS-755-X3-SC	* PWR-D4-3041-AC-BLUE

Transceiver

* 1000BASE-BX10	* 40GBASE-PLR4	* OSFP-800G-2FR4
* 1000BASE-BX10-D	* 40GBASE-PLRL4	* OSFP-800G-2LR4
* 1000BASE-BX10-U	* 40GBASE-SR4	* OSFP-800G-2PLR4
* 1000BASE-LX	* 40GBASE-SRBD	* OSFP-800G-2VSR4
* 1000BASE-SX	* 40GBASE-UNIV	* OSFP-800G-2XDR4
* 1000BASE-T	* 40GBASE-XSR4	* OSFP-800G-AR8
* 100BASE-FX	* 50GBASE-CR	* OSFP-800G-VSR8
* 100GBASE-AOC	* A-D400-2Q200	* OSFP-AMP-ZR
* 100GBASE-CR4	* A-D400-Q400	* QDD-200G-2LR4
* 100GBASE-CWDM4	* A-O400-2Q200	* QDD-400G-DR4
* 100GBASE-DR	* A-O400-Q400	* QDD-400G-FR4
* 100GBASE-ERL4	* ADPT-O-Q-100G	* QDD-400G-LR4
* 100GBASE-FR	* AOC-D-D-400G	* QDD-400G-LR8
* 100GBASE-LR	* AOC-O-O-400G	* QDD-400G-PLR4
* 100GBASE-LR4	* CAB-D-2Q-200G	* QDD-400G-PLR4-S
* 100GBASE-LRL4	* CAB-D-2Q-400G	* QDD-400G-SR8
* 100GBASE-PLR4	* CAB-D-4Q-200G	* QDD-400G-SR8-C
* 100GBASE-PLRL4	* CAB-D-4Q-400G	* QDD-400G-SRBD
* 100GBASE-PSM4	* CAB-D-8S-200G	* QDD-400G-VSR4
* 100GBASE-SR4	* CAB-D-8S-400G	* QDD-400G-XDR4
* 100GBASE-SRBD	* CAB-D-D-400G	* QDD-400G-XDR4-S
* 100GBASE-SWDM4	* CAB-O-2Q-200G	* QDD-400G-ZR
* 100GBASE-XCWM4	* CAB-O-2Q-400G	* QDD-400G-ZRP
* 100GBASE-XSR4	* CAB-O-4Q-200G	* QDD-800G-2FR4
* 100GBASE-ZR4	* CAB-O-4Q-400G	* QDD-800G-2LR4
* 100GE-DWDM2	* CAB-O-8S-200G	* QDD-800G-2PLR4
* 10GBASE-AOC	* CAB-O-8S-400G	* QDD-800G-2VSR4
* 10GBASE-CR	* CAB-O-O-400G	* QDD-800G-2XDR4
* 10GBASE-DWDM	* CAB-Q-2Q-100G	* QDD-800G-AR8
* 10GBASE-DWDM-T	* CFP2-100G-DWDM-DCO	* QDD-800G-VSR8
* 10GBASE-DWDM-ZR	* CFP2-100GBASE-ER4	* QDD-8X25-MR-SR-E
* 10GBASE-ER	* CFP2-100GBASE-LR4	* QDD-8X25R-SR-E
* 10GBASE-ERBD	* CFP2-100GBASE-XSR10	* QSFP-100G-ER
* 10GBASE-ERBD-D	* CFPX-100G-DWDM	* QSFP-100G-SR1.2
* 10GBASE-ERBD-U	* CFPX-200G-DWDM	* QSFP-100G-SR4-E
* 10GBASE-ERLBD	* E-D800-D800	* QSFP-200G-AR4
* 10GBASE-ERLBD-D	* E-O800-O800	* QSFP-200G-FR4
* 10GBASE-ERLBD-U	* H-D400-4Q100	* QSFP-AMP-ZR
* 10GBASE-LR	* H-O400-4Q100	* SFP-10G-MRA-T

* 10GBASE-LRL	* LPO-800G-2DR4	* SFP-10G-RA-1G-LX
* 10GBASE-SR	* OSFP-200G-2LR4	* SFP-10G-RA-1G-SX
* 10GBASE-SRL	* OSFP-400-2FR4-LC	* SFP-10G-T
* 10GBASE-ZR	* OSFP-400G-2FR4	* SFP-10G-T-RP
* 200GBASE-CR4	* OSFP-400G-DR4	* SFP-10GRA-T
* 200GBASE-SR4	* OSFP-400G-FR4	* SFP-25G-LR-E
* 25GBASE-AOC	* OSFP-400G-LR4	* SFP-25G-MR-LR
* 25GBASE-CR	* OSFP-400G-LR8	* SFP-25G-MR-SR
* 25GBASE-ER	* OSFP-400G-PLR4	* SFP-25G-MR-XSR
* 25GBASE-LR	* OSFP-400G-PLR4-S	* SFP-25G-SR-E
* 25GBASE-SR	* OSFP-400G-SR8	* SFP-25GR-LR
* 25GBASE-XSR	* OSFP-400G-SRBD	* SFP-25GR-SR
* 40GBASE-AOC	* OSFP-400G-VSR4	* SFP-25GR-XSR
* 40GBASE-CR4	* OSFP-400G-XDR4	* SFP-50G-MR-LR
* 40GBASE-ER4	* OSFP-400G-XDR4-S	* SFP-50G-MR-SR
* 40GBASE-LR4	* OSFP-400G-ZR	* SFP-50GR-LR
* 40GBASE-LRL4	* OSFP-400G-ZRP	* SFP-50GR-SR

Reference to MLAG ISSU Upgrade/Downgrade Table

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

Resolved Caveats in 4.34.3.1M

7300X3, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7060X6, DCS-7260X3 and DCS-7300X Series

7300X3 and DCS-7060X Series

- MAC address learning may fail to initialise during system startup due to a memory allocation error during driver initialisation. When this occurs, kernel messages about page allocation failures will appear in system logs.

Workaround: Reboot the switch. (1217098) (1)

DCS-7060X2 Series

- MAC address learning may fail to initialise during system startup due to a memory allocation error during driver initialisation. When this occurs, kernel messages about page allocation failures will appear in system logs.

Workaround: Reboot the switch. (1217098) (1)

DCS-7260X3 Series

- MAC address learning may fail to initialise during system startup due to a memory allocation error during driver initialisation. When this occurs, kernel messages about page allocation failures will appear in system logs.

Workaround: Reboot the switch. (1217098) (1)

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

- MAC address learning may fail to initialise during system startup due to a memory allocation error during driver initialisation. When this occurs, kernel messages about page allocation failures will appear in system logs.

Workaround: Reboot the switch. (1217098) (1)

Known Software Caveats in 4.34.3.1M

Accelerated System Update

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

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

- SSU may fail and result in cold reboot leading to an extended traffic outage. (856802)
SKUs Affected: DCS-7060DX5-64S, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- Port-channel members and the port-channels that are configured with ``no switchport`` may flap during SSU.
The workaround is to configure ``default switchport`` on port-channels, and port-channel members. (860647)
- Connections using tunable SFP transceivers like the 10GBASE-DWDM-T and 10GBASE-DWDM-ZR may flap during 'reload fastboot'. There is no workaround, but after the flap(s), the link should be operational again. (878971)
SKUs Affected: 10GBASE-DWDM-T and 10GBASE-DWDM-ZR
- SSU will fail if a linecard or secondary switchcard are powered down prior to performing SSU. (914302)
SKUs Affected: CCS-755-CH and CCS-758-CH
- During SSU, network churn events (such as remote VTEP flaps and route churns) may result in extended traffic outage. (927699)
- SSU for the below release combinations is not supported:
 - From the source release 4.28.8.3M and later releases in the 4.28 train to

- the target releases in the range between 4.29.0F (inclusive) and 4.29.8 (exclusive) <- this is the first release the patch was casted to 4.29 train
- From the source release 4.28.8.3M and later releases in the 4.28 train to the target releases in the range between 4.30.0F (inclusive) and 4.30.6 (inclusive)
 - From the source release 4.28.8.3M and later releases in the 4.28 train to the target releases in the range between 4.31.0F (inclusive) and 4.31.2 (inclusive)
 - From the source release 4.28.8.3M and later releases in the 4.28 train to the target releases in the range between 4.32.0F (inclusive) and 4.32.1F (exclusive)
 - From the source release 4.29.8M and later releases in the 4.29 train to the target releases in the range between 4.30.0F (inclusive) and 4.30.6 (inclusive)
 - From the source release 4.29.8M and later releases in the 4.29 train to the target releases in the range between 4.31.0F (inclusive) and 4.31.2 (inclusive)
 - From the source release 4.29.8M and later releases in the 4.29 train to the target releases in the range between 4.32.0F (inclusive) and 4.32.1F (exclusive)
 - From the source release 4.31.3M and later releases in the 4.31 train to the target releases in the range between 4.32.0F (inclusive) and 4.32.1F (exclusive) (960833)
- SSU is not supported when upgrading from 4.21.2.9F and older releases to this release (1073003)
 - When performing SSU upgrade from releases older than EOS-4.29.2F to the affected releases, SSU upgrade may fail with fatal error, if at least one of the interfaces is errdisabled with the speed-misconfigured cause. Fatal error will result in a regular reload, causing link flaps and traffic disruption.

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

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

BGP EVPN L2/L3 VXLAN/MPLS

- A policy change (such as changing the color extended communities) which causes an L2 EVPN MAC-IP or IMET route to resolve over a different SR-TE policy or tunnel may cause a brief traffic loss. (521090)
- A remote VTEP reached through unnumbered BGP over IPv6 will be incorrectly displayed as a configuration error with "FAIL" state due to "No route" (525842)

- In an EVPN multi-homing deployment, designated forwarder election for a given Ethernet segment and EVPN instance pair is based on the complete set of VLANs within the EVPN instance. This differs from the standard in which only the instance VLANs that are in the Ethernet segment interface's allowed VLAN list are used in the election algorithm. This deviation may cause interoperability issues with other vendors. It may also cause devices to be elected DF for an EVI on which they can't forward, if there is disjoint configuration of allowed VLANs for the ES interface on different peers in the same ES.

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

- On Multi-Domain EVPN VXLAN deployment using anycast gateways, administratively enabling and disabling the VXLAN interface may cause and ARP entry to be lost and trigger traffic lost. (632955)
- In an L2 EVPN MPLS configuration, disabling MAC address learning on a VLAN that is part of the EVPN bridging domain does not take effect. (677182)
Series Affected: 7500R3, DCS-7280CR3, DCS-7280DR3, DCS-7280PR3, DCS-7280R3, DCS-7280SR3, DCS-7500R, DCS-7500R3 and DCS-7800R3 Series
- When configuring "set_nextthop_self()" using RCF in the EVPN outbound context, the BGP agent can restart. (717957)
- In an EVPN VXLAN deployment, the BGP process can leak memory when the VLAN to VNI mapping is continuously provisioned and unprovisioned. The amount of memory leaked does not scale with the number of routes but instead with the number of MAC-VRFs. (774066)
- In a EVPN VXLAN deployment with L2 ARP or ND proxy enabled ("arp learning bridged" for IPv4 hosts and/or "nd learning bridged" for IPv6 hosts), when a GARP is generated by a host to update its MAC in an ARP or ND binding, the multi-homed PEs may advertise stale and conflicting MAC to ARP/ND binding. (838924)
- On an EVPN router, the system might run out of dynamic VLANs available, in the event of a VXLAN VRF VNI mapping configurations churn.

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

- When one of the multi-homing gateways is reloaded, there may be a period of time during which traffic loss occurs.
To mitigate its effect, configure an event handler that disables the I-ES for a brief period of time after the reload. (854714)

- In a EVPN VXLAN deployment where L2 ARP and ND proxy is enabled, executing the "clear arp|nd bridged" or the "no arp|nd learning bridged" command might not clear all the locally learned L2 ARP entries in the "show arp bridged" or "show ipv6 neighbors bridged" output. (882754)
- In a EVPN VXLAN deployment, if an L2 VTEP is configured with "arp|nd learning bridged" receives a VXLAN encapsulated NA in response to a IPv6 DAD request, the response may not be decapsulated and bridged.

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

- When transitioning IPv6 L3 VARP VTEP functionality from one VTEP to another, IPv6 neighbor entries may not be properly updated from being remotely learnt through EVPN, to being locally learnt. This may cause traffic to be encapsulated with stale MAC addresses. (1005904)
- In EVPN deployments where learned or a static MAC pointing to a front panel port conflicts with the system MAC or VARP MAC (if configured) or VRRP MAC, the EVPN MAC-IP route might be incorrectly withdrawn or advertised with the sticky bit unset.

As a workaround, if the route is withdrawn unexpectedly, MAC redistribution can be re-triggered either by relearning the MAC if the front panel learned MAC redistribution is expected, otherwise, the `redistribute router-mac..` config can be removed and added back again to re-trigger Router MAC redistribution if that's expected instead. (1030035)

- When a multihoming group is configured with the `designated-forwarder election algorithm hrw`, all VLANs in an Ethernet segment are assigned to the same designated forwarder. This prevents load balancing and can lead to uneven traffic distribution.

To work around this issue, configure the `designated-forwarder election algorithm modulus` to distribute VLANs automatically among designated forwarder candidates. (1065539)

- In an EVPN multihoming setup, if only one peer advertises an EVPN Type 2 MAC IP route for a host, the other peers on the same ethernet segment may temporarily advertise that route when the original peer withdraws its Type 1 or Type 2 route even if the host is unreachable, and then withdraw the route immediately afterward. This behavior can also occur during a MAC move. (1073012)
- In an EVPN MPLS setup, a PE may never age out an ARP entry as long as the corresponding MAC address exists in the MAC table. In some cases, if the entry was originally on a different PE and moved to this PE, the PE may also repeatedly send out ARP requests for an entry.

Note that this does not affect IPv6 neighbors.

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

- In EVPN Multicast transit network with a RP on a stick topology, some multicast receivers connected to the PIM External Gateway that is on the path to the source from the RP may not get traffic. The workaround is to add a receiver on another PIM External Gateway. (1079151) (1)
- In an EVPN MPLS multihome active-active setup, if an IPv6 device changes its address, all active multihome peers will repeatedly send Neighbor Solicitation (NS) packets for the old IPv6 address. This causes them to repeatedly switch between advertising proxy and non-proxy MAC-IP routes. As a result, traffic may be incorrectly sent to these peers, even if a new device with a different MAC has taken the IP address on a different PE.

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

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

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

- IPv4 ICMP messages generated by the system in VRFs without IPv4 interface for TTL expiry and unreachable destination may incorrectly use `192.0.0.8` source IP address instead of the VTEP IP address, despite `icmp helper vrf` being configured under VXLAN interface.

As a workaround, a loopback interface with an IPv4 address may be configured in the VRFs that are without any IPv4 interfaces. (1153892)

- In EVPN DCI deployment, if a remote domain IP VRF RT is not configured while using remote domain EVPN peers that exchange EVPN Type-2 routes that are imported into L3 VRFs, we may observe a one time BGP agent restart or memory leak in the BGP agent when the peering goes down.

Removing and re-adding the local domain RT for the affected L3 VRFs will help address the memory leak. (1196798)

- In a EVPN VXLAN MH Active-Active IRB deployment, if a device's ARP entry is learned on all MH peers and that device is later bound to a new MAC and announces it through it a GARP to one of the Active-Active PEs, the PE receiving the ARP packet may delete its local ARP entry and replace it with a remote ARP entry with the old MAC instead of updating the ARP entry with the new MAC. This will cause traffic destined to the IP to continue to be sent to the old MAC

Workaround is to add an RCF rule to not advertise EVPN type 2 routes for hosts whose MACs are expected to change or have the host resend multiple GARP packet upon changing MACs. (1226818)

- In EVPN DCI deployment, if both remote and local domain IP VRF RT is configured while using remote domain EVPN peers that exchange EVPN Type-2 routes that are imported into L3 VRFs, EVPN may not be able to import paths into the L3 VRFs using the local domain RT when the remote domain RT is removed.

Removing and re-adding the local domain RT for the affected L3 VRFs will help EVPN import paths into the IP VRFs from the remote domain peer using local domain RT (1248513)

- In EVPN DCI deployment with remote domain EVPN peers configured, if a remote domain RT were to be configured on an IP VRF that has imported routes from the remote domain EVPN peer with local domain RT, the paths imported with local domain RT may not be removed. On the other hand, if remote domain RT were to be un-configured with local domain RT left intact on an IP VRF that has imported routes from the remote domain EVPN peer with remote domain RT, the paths imported with remote domain RT may not be removed.

Removing and re-adding the import RT for the affected L3 VRFs on the domain with stale paths will help remove the stale imported paths from EVPN (1248596)

- On devices running EVPN OISM with underlay VXLAN multicast, having only an S,G Exclude overlay receiver might result in not joining the underlay tree.

Workaround: add a static IGMP/MLD snooping receiver to force joining the underlay tree. (1258360) (1)

- In a EVPN VXLAN IRB deployment with dual-stack VLAN interfaces, the IPv6 remote bindings from EVPN may not be re-learned against the dual-stacked VLAN interfaces after agent or device reboot.

A work around is to "shutdown" and "no shutdown" the VLAN interfaces.
(1270132)

- If VRRP is configured on an Ethernet port, and an EVPN MAC-VRF is configured, the Bgp agent can restart unexpectedly. (1289709)

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

Containerized EOS

- Multicast traffic stops flowing in a EVPN OISM Vxlan network after reload of ceoslab or veoslab switches (1218937)

vEOS Router / CloudEOS

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

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

- A reboot due to a kernel panic can happen during first 10 minutes of vEOS bootup on Microsoft Azure. The kernel panic occurs because of a task blocking without being scheduled for more than 300 seconds in "D" state. (249618)
- Interface may be allowed to be configured with un-supported speed of NIC using "speed <speed>", which will cause an unexpected PhyEthtool agent restart. Do not use "speed <speed>" command to configure interface speed that NIC does not support. (264395)
- When loading a vEOS instance in an Azure cloud, it is possible for the device to experience a kernel panic just after bootup, causing the device to reboot a second time. (354411)
- In an EVPN IRB deployment on vEOS, changing the VXLAN source-interface interface may result in EVPN ip-prefix routes not getting installed into one or more VRFs for one or more remote VTEPs. (369816)
- CloudEOS supports upto 600K BGP routes with 8G of system memory; If 800K BGP routes are programmed, out-of-memory condition will kill random processes.

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

- When an IPsec connection is established with a primary IP address of an interface and a secondary IP address is configured on an interface and that secondary IP address is used to establish an IPsec connection, the IPsec connection with the secondary IP address doesn't come up.
The workaround is to restart the IPsec agent. (572808)
- When running CloudEOS on ESXI hypervisor, while using Intel Ethernet Controller X710 NIC in SRIOV mode, the TX traffic out of the interface, may stop working after either the router reload and/or SFE agent (Software Forwarding Engine) restart for any reason.
The workaround is to set "trust-mode" ON and "spoof-check" OFF for the SRIOV interface on the VMware hypervisor.
Another possible workaround is to log into the ESXI VCenter console and change the SRIOV port's assigned VLAN to any other VLAN and then revert it back to the original VLAN. (647410)
- SSH provisioning fails while creating a VM due to infra changes on Azure's platform.
A workaround is to set up the VM using a username and password and then set up SSH afterwards. (681168)
- The /var/log/ filesystem on AWS CloudEOS VMs may get filled up with awslogs.*backup files. The workaround is to delete the /etc/cron.d/awslogs_log_rotate file. This workaround is better applied using an on-boot event handler that deletes this file on every reboot of the VM. (713562)
- CloudHA uses older Boto SDK version than some of the AWS regions allow. CloudHA may not work in such AWS regions.

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

Boto config file content with the exact path:

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

Event handler configuration:

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

```
delay 0
---- (758301)
```

- STUN clients may not receive responses from the STUN server if there is an ECMP path from the server to the client. (1016192) (1)
- The "show adaptive-virtual-topology path" command may hit an endless loop when the device receives a multihop path that doesn't belong to it from the Pathfinder. The workaround is to restart the Sfe agent. (1276475)

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

CVP

Network Provisioning

- During Zero Touch Provisioning(ZTP), repeated failures to execute a bootstrap script that forks other scripts may result in the switch running out of memory.
The workaround is to rewrite the bootstrap script to either avoid forking or properly cleanup forked child processes when the script exits, and then reloading the switch to restart ZTP. (820844)
- Device enrollment and consequently ZTPaaS onboarding could fail for a device, that was enrolled to CVaaS before, with the error: "Failure when receiving data from the peer". The workaround is to delete the TerminAttr certs stored at /persist/secure/ssl/terminattr and attempt ZTPaaS onboarding again. (916124)

Telemetry

- OpenConfig may return inaccurate PHY Pre-FEC bit error rate, FEC uncorrected codewords, or FEC corrected codewords. (1290065) (1)

CVX

- If a VmTracer session configures all the available VLANs on the switch, then the switch may errdisable the routed ports by freeing up the internal VLANs. (139294)
- On all switch series configured as a MLAG pair where CVX is used as the VXLAN control plane, after a MAC address move between switches the MAC address might be advertised to the CVX server by two different VTEPS. This can cause packets destined to the MAC address to be blackholed. The workaround is to clear the mac address on both the advertising VTEPS so that the mac address will be relearned properly. (158130)

- The MacroSegmentation Service (MSS), working with L2 transparent firewalls, requires the firewall interface be statically identified on CVX via configuration commands (as opposed to using LLDP to detect them dynamically). This prevents issues such as traffic loss in the event that a member port of a aggregated link goes down. (275384)
- Macro-segmentation Service (MSS) might stop intercepting traffic when the last member of the Near service LAG goes down.

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

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

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

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

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

- In a VCS deployment with MLAG, a rapid MAC move may result in a MAC being learned locally on multiple VTEPs. The workaround is to run "clear mac address-table dynamic vlan VLAN" on the MLAG secondary which should not have the MAC learned locally. (702518)
- In a VXLAN deployment using VXLAN Controller Service (VCS) to distribute MAC addresses, the MAC address of an orphan host behind a secondary MLAG switch may not get published to VCS. (770762)
- The VXLAN Controller Service (VCS) feature available on CVX deployments will no longer be supported.

Please migrate the deployment to using an BGP/EVPN control plane. If the deployment must continue to use VCS, please continue to upgrade EOS to the

latest version in the train where this defect remains open (that is a release where the feature has not been removed). (950887)

- In a Vxlan Controller Service deployment where VRRP is configured on the MLAG peer, VRRP MAC and VRRP MAC-IP binding maybe not be advertised if the VRRP Master is not the MLAG primary. (1089309)
- MCS Client could crash and restart if we receive a IGMP join with an empty group address. (1263601)
- If the Media Controller Service (MCS) in a CVX HA cluster restarts before its clients are configured on the switches, it may not function correctly.

Reloading the CVX leader can resolve this issue. (1267007)

Enterprise routing

- The XcvrAgent may restart unexpectedly when the AWE-7220RP-5TH-2S-F system is operating under high load. There is no workaround. (957413)
SKUs Affected: AWE-7220RP-5TH-2S-F
- Policy based IPsec tunnels flap continuously when sa idle-timeout is configured.

A workaround could be removing sa idle-timeout configuration if configured. (1184208)

- Update your release note here.
Policy based VPN tunnel does not get established over tunnels with directly connected tunnel destination.
A workaround would be to add a static ARP entry for the tunnel destination. (1281098)
Series Affected: AWE-5000 and AWE-7000 Series

SDWAN

- Hardware encrypted DPS IPsec packets in tunnel mode are sent without UDP header.

The workaround is configure the IPsec in transport mode. (878530)
Series Affected: AWE-5000 and AWE-7000 Series

- When auto-discovery of remote DPS peers is enabled in a path-group using the 'peer dynamic' CLI, receiving delete notification for a remote peer may cause the Sfe agent to restart unexpectedly. (1093407)
Series Affected: AWE-5000, AWE-7000 and CloudEOS VM Series

General

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

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

- The OpenFlow agent fails to return counters associated with a flow, when queried by the OpenFlow controller using protocol version 1.0 (132812)
- On switches using the tg3 driver for the management interface a kernel panic can happen resulting in a switch reload. This is due to a Single Event Upset (SEU) affecting the management interface ethernet controller, and is a rare occurrence. (136944)
Series Affected: 7368 Series
SKUs Affected: DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050SX3-48YC12-F, DCS-7060CX-32S-F, DCS-7060CX-32S-R, DCS-7060CX2-32S-F, DCS-7060CX2-32S-R, DCS-7060SX2-48YC6-F, DCS-7060SX2-48YC6-R, DCS-7170-32C-F, DCS-7170-32C-M-F, DCS-7170-32C-M-R, DCS-7170-32C-R, DCS-7170-32CD-F, DCS-7170-32CD-R, DCS-7170-64C-F, DCS-7170-64C-M-F, DCS-7170-64C-M-R, DCS-7170-64C-R, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7304, DCS-7308, DCS-7504N, DCS-7508N and DCS-7512N
- When configuration of a service-policy of type pdp on an interface fails due to lack of hardware resources, the interface may not be protected by any PDP policy. The workaround is to remove the configuration of the PDP service-policy and to configure the PDP service-policy default-pdp-policy on the interface. (216365)
- The MacroSegmentation Service (MSS) feature doesn't work in conjunction with the VARP VTEP IP configuration required for VXLAN routing where a subset of VTEPs are L2 VTEPs. (263532)
- When a MLAG peer (configured as a service VTEP for MSS) reloads, MSS re-installs intercept flows which may cause traffic to drop for a short duration. (349254)
- If multiple front-panel ethernet interfaces are configured as reflector interfaces for the same flow of traffic to be reflected, loops can be formed and hosts can flap from one interface to another when the reflector configuration is subsequently changed. The workaround is to only configure one reflector interface to handle traffic reflection. (360869)

Series Affected: DCS-7500R Series

- When configuring VXLANs from the EOS CLI, "ip address virtual" is supported and mapped to YANG, but "ip virtual-router address" is not supported. (372193)
- If the number of PTP vlan's enabled on a switch using 'ptp vlan' command exceeds 65536, PTP agent will become unresponsive and eventually restart. Reducing the VLANs below the limit will avoid getting in this state. (408199)
- When in transparent two-step mode PTP packet sequenceId collision may cause transient jump in offsetFromMaster. Workaround is to use one-step transparent clock mode. (408202)
- DirectFlow 'action output nexthop <ip addr>' ignores TTL of incoming packet and forwards the packet even if the TTL has expired. (417994)
- System may unexpectedly reboot with an Unknown reload cause due to Uncorrectable ECC errors. (419257)
 SKUs Affected: DCS-7050CX3-32S, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050SX3-48YC12, DCS-7050SX3-48YC12-F, 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
- In rare conditions, the kernel may panic with the message containing "BUG at ../mm/slub.c:3334" (476225)
- ZeroTouch provisioning will fail to fetch configuration on QSFP-DD/OSFP ports with inserted QSFP or SFP transceivers or QSFP ports with inserted SFP transceivers when these ports require a non-default speed-group configuration to link up.

 The workaround is to use a different type of transceiver or a different port to connect to the bootstrap server. (488491) (1)
- On system reload, the kernel may crash with the "PANIC: double fault" message. (493516)
- Local (admin) user is allowed to login with empty password when RADIUS over TLS is enabled (569426)
- "reload peer" configured with a delay such as "reload peer in 1 force" results in both supervisors being reloaded after the specified delay. (621572)
- BOM change on PCA-01601-03 and change of riser pol image from SFT-01392-01 to SFT-01392-02 resolves issue. (732299)
 SKUs Affected: DCS-7280DR3A-54-F, DCS-7280DR3AK-54 and DCS-7280DR3AM-54-F
- In an EVPN VXLAN multihoming setup, administratively shutting down and re-enabling the VXLAN source loopback interface on a remote VTEP can result in

MAC-IP routes of some hosts, received from the multihoming VTEPs, not being installed causing the traffic to fail to those hosts, in the absence of other covering routes. (804491)

- Units may be unable to complete exiting from maintenance mode if entering maintenance mode is cancelled before reaching Under Maintenance state. (813375)
- The KernelNetworkInfo IP address shared memory collection can grow over time if interfaces and IP addresses are repeatedly removed and added back again. If this causes an issue, the KernelNetworkInfo agent can be restarted and the IP address collection will return to its original size. (824906)
- The device may be unable to finish exiting from maintenance mode if entering maintenance mode is cancelled before reaching Under Maintenance state. (933059)
- OpenConfig path /network-instances/network-instance[name=]/protocols/protocol[identifier=][name=]/pim/interfaces/interface[interface-id=]/neighbors/neighbor[neighbor-address=]/state/mode incorrectly reports PIM_MODE_DENSE when neighbor is in bidirectional pim mode only. (960478)
- In case of a GM switchover when the GMs' times are significantly different, downstream S-clock, connected to the M-clock experiencing the said switchover, might enter a brief period (~120 Sync message intervals) of incorrect elevated OFM/skew values. The workaround is to ensure that GMs are reliable sources of time. (983258)
- The Sfe agent may be unresponsive for up to 30 seconds when initializing or reconfiguring certain Ethernet ports, resulting in loss of traffic for that period. (1009668)
Series Affected: CloudEOS VM Series
SKUs Affected: AWE-7220RP-5TH-2S and AWE-7220RP-5TH-2S-F
- In some rare cases, systems with a SMART eMMC may experience a transient disconnect resulting in -110 (timeout) errors. This leads to the flash filesystem becoming read-only in EOS. When this issue occurs, the eMMC becomes totally disconnected and no longer replies to commands. Even though the file system shows as read only, reads and writes will fail.

This disconnect is transient and does not relate to any permanent condition or issue with the eMMC. A system experiencing this eMMC transient disconnect can be fixed by reloading the system. (1018819)

SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720XP-24ZY4-F, CCS-720XP-48TXH-2C-S-F, CCS-720XP-48Y6-R, CCS-720XP-48ZC2-F, CCS-720XP-96ZC2-F, CCS-722XPM-48Y4, CCS-722XPM-48ZY8, DCS-7010TX-48, DCS-7010TX-48-DC, DCS-7050CX3-32S-F, DCS-7050CX3-32S-R, DCS-7050CX3M-32S, DCS-7050SX3-48C8, DCS-7050SX3-48YC8, DCS-7050TX3-48C8, DCS-7060CX-32C,

DCS-7060CX-32S, DCS-7060SX2-48YC6, DCS-7280CR3-32D4, DCS-7280CR3-96,
DCS-7280CR3K-96, DCS-7280DR3-24, DCS-7280DR3K-24, DCS-7280PR3K-24,
DCS-7280SR3-40YC6, DCS-7280SR3-48YC8, DCS-7280SR3E-40YC6,
DCS-7280SR3K-48YC8, DCS-7280SR3M-48YC8 and DCS-7280TR3-40C6

- When first booting up or after a SSO, one or more linecards may remain in the "Initialized" state in the "show ptp hardware-sync" command output, causing the device to be unable to participate in PTP.

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

Series Affected: CCS-750 Series

- ACL Agent can terminate on network namespace change on Layer3 interfaces (1038673)
- If an EOS event-handler has these 3 conditions:
 - The action is triggered on config
 - The handler has a delay
 - The first event is detected within the delay

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

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

- SSL profiles become "invalid" if the system clock is outside the validity period of the certificates. If the system clock is later changed to be within the validity period via NTP, SSL profiles will still remain "invalid".

To re-trigger the SSL profile validation, do any configuration changes in the affected SSL profiles or restart the SuperServer agent. (1068537)

- Updating PBR rules with new priority_mask settings does not work until the rules are deleted and added back. (1087146)
- In a traffic policy deployment, the mirror action for a rule may fail even if the target mirroring session is up. As a result, traffic matching the rule will not be mirrored. (1139348) (1)
- In certain circumstances, the kcompactd kernel thread can run continuously for a long period of time without yielding the CPU, causing a soft lockup panic. (1175555)
- The system may reboot due to a voltage glitch with debugging information showing FixedSystemcpuCardPowerController Rail#7 (VOUT_OV) (1216844)

SKUs Affected: DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7050SX3-48C8-F,

DCS-7050SX3-48C8-R, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R,
 DCS-7050SX3-96YC8-F, DCS-7050SX3-96YC8-R, DCS-7050TX3-48C8-F,
 DCS-7050TX3-48C8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F,
 DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F,
 DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R,
 DCS-7280CR3-36S-F, DCS-7280CR3-36S-R, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F,
 DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R,
 DCS-7280CR3K-36S-F, DCS-7280CR3K-36S-R, DCS-7280CR3K-96-F,
 DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F,
 DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R,
 DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F,
 DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F,
 DCS-7280PR3K-24-F, DCS-7280SR3-40YC6-F, DCS-7280SR3-40YC6-R,
 DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R, DCS-7280SR3E-40YC6-F,
 DCS-7280SR3E-40YC6-M-F, DCS-7280SR3E-40YC6-M-R, DCS-7280SR3E-40YC6-R,
 DCS-7280SR3K-48YC8-F, DCS-7280SR3K-48YC8-R, DCS-7280SR3M-48YC8-F,
 DCS-7280SR3M-48YC8-R, DCS-7280TR3-40C6-F and DCS-7280TR3-40C6-R

- Performing a config replace may cause multicast routing to be disabled in some VRFs. A workaround is to restart Igmp agent. (1219461)
- When a maintenance unit is configured to enter maintenance mode immediately after boot-up using two methods simultaneously (an "on-boot" duration in its profile and a quiesce directive in the startup configuration), some agents, such as the BGP routing agent, may fail to start. This can also cause interfaces to remain stuck in an "initializing" state.

To prevent this issue, it is recommended to remove the "on-boot" duration from the maintenance unit's profile. (1223377)

- With "vxlan flood vtep learned data-plane" configuration, clearing MAC table does not clear learnt VTEPs. (1225724)
- When running CloudEOS VM instance on VMware ESXi hypervisor, if an Intel 82599 network interface is used in SR-IOV mode, traffic may not flow through the interface after booting the VM instance.
 Restarting the Sfe agent will fix this condition. (1232621)
 Series Affected: CloudEOS VM Series
- When running CloudEOS VM instance on VMware ESXi hypervisor, if an Intel XL710 network interface is used in SR-IOV mode, the network interface may not come up after booting the VM instance.
 The workaround is to flap the affected interface using "shutdown" followed by "no shutdown" config command. (1232626)
 Series Affected: CloudEOS VM Series
- Using the the event handler with the "on-logging" trigger and "threshold" command might trigger the events that were skipped once the threshold is reached. This might cause the trigger of more actions than expected.

Workaround: None (1241235)

- BGP metadata will be omitted from the IPFIX flow records for traffic which is decapsulated from a static VXLAN tunnel and routed in a non-default VRF. (1255255)
- eAPI can unexpectedly delay responses for valid requests. (1270812)
- Adding 'goto next' action for the last match in a policy will cause continuous restarts of the traffic-policy forwarding agent. (1273277)
- In a modular system with a large number of line-cards, the EventMgr agent might be terminated with a SIGQUIT.

Workaround: Disable the built-in event VoqDeleteEventHandler handler with the config commands:

no event-handler VoqDeleteEventHandler (1276148)

- In EVPN Symmetric IRB deployments, when ipv6 virtual-router is configured, and the VLAN interface does not have a VLAN to VNI mapping configured, then the Neighbor Solicitation packets for VLAN interface's virtual IPs may not update the neighbor entries previously learned from a Neighbor Advertisement, which can result in stale neighbor resolution.

The workaround is to use "ip address virtual" for the VLAN interface instead of "ipv6 virtual-router". (1281004)

- AAA enable authentication sends empty remoteHost and tty to remote server. This may cause the server to reject the request. (1296293)

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

IPSEC

- Under some circumstances IPsec rekey time may show up as 'N/A'. (561575)
- In Jericho2 IPsec tunnel mode, both peers should either enable or disable replay protection. Asymmetric configuration will lead to complete packet loss over the tunnel. (605866)
- On a topology configured with IPsec AutoVPN, doing a shut/no shut several times on interfaces can cause the IPsec agent to restart unexpectedly. (634820)
- When IPsec ESP packets are received on a NAC port with RSS enabled on Councilbluff or Independence device, the IPsec ESP packets are not

distributed across multiple workers/cores, but all the IPsec packets are sent to a single core.

This leads to all the IPsec packets processed by a single worker/core and the rest of the free workers/cores not utilized. (1089266)

Series Affected: AWE-5000 and AWE-7000 Series

- IPsec agent may experience high CPU utilization of up to 100% when two or more IPsec tunnels which share the same underlying IPsec connection and profile are administratively shutdown.

Workaround is to remove the IPsec profile configured under the tunnel interfaces which are shut. (1233491)

Series Affected: AWE-5000 and AWE-7000 Series

SKUs Affected: DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4A-S, DCS-7280CR3MK-32D4S and DCS-7280CR3MK-32P4

- When the vrf name starts with a number or a hyphen and a number(-123), the charon log files in /var/log/agents will not be rotated and the charon log file size grows continuously. The workaround is to change the vrf name to start with a letter. (1273055)

Series Affected: AWE-5000, AWE-7000 and CloudEOS VM Series

Layer 1

- The forwarding agent may unexpectedly restart after a sequence of speed changes involving speeds of 10G, 25G, 40G, 50G and 100G. (358317)

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

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

SKUs Affected: 7300X3-32C-LC, 7300X3-48TC4-LC, 7300X3-48YC4-LC and DCS-7050CX3M-32S

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

Series Affected: DCS-7260CX3 Series

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- L1 Profiles configuration is not supported with OpenConfig. (789141)
- Executing "shutdown" interface configuration command after SSU may result in an unexpected forwarding agent restart and a brief traffic outage on all ports. (792791)

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

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

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

- Using the rc.eos script to add port numbering to the boot config will result in the port numbering not being applied.

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

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

- Interfaces with 1GBASE-T SFPs installed may not link up. Workaround is to shut/no shut the port to recover. (847035)

Series Affected: CCS-720DF Series

SKUs Affected: CCS-720DF-48Y, CCS-720DF-48Y-2, CCS-720DF-48Y-2F and CCS-720DF-48Y-F

- Interfaces configured with 1G speed may flap during an SSU reload if they were ever configured with 100M speed prior to the SSU reload. (890078)

SKUs Affected: CCS-720DF-48Y

- A link flap on BASE-T ports configured with the default speed ("no speed" or "default speed") may cause an unexpected restart of the forwarding agent. This leads to a flap of all interfaces on the line card and a brief traffic outage. (899785)

SKUs Affected: CCS-750X-48TP-LC

- A transceiver with a corrupted EEPROM may cause the interface sourcing configuration (l1 source ...) to be dropped for all four interfaces tied to the same software agent.

The workaround is to remove the corrupted transceiver and reboot the switch. (908341)

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

- MAC faults may fail to propagate to the internal CMS42550 PHY chip. This may lead to a one way link up. (917641)

SKUs Affected: DCS-7170B-64C-F

- On speed change, the PhyIsland agent can restart unexpectedly, causing the links of the system to go down momentarily. (934072)

SKUs Affected: DCS-7280CR3A-24D12, DCS-7280CR3A-32S, DCS-7280CR3AK-24D12, DCS-7280CR3AK-32S, DCS-7280CR3AM-24D12 and DCS-7280CR3AM-32S

- An interface may fail to link up after a configuration change, reload, or power cycle.

Workaround is to use "shutdown" followed by a "no shutdown" CLI command on the affected interface. Workaround may be needed to be applied multiple times to recover the interface. (992949)

SKUs Affected: DCS-7060CX5-56D8, DCS-7280CR3A-24D12, DCS-7280CR3A-32S, DCS-7280CR3AK-24D12, DCS-7280CR3AK-24D12-S, DCS-7280CR3AK-32S, DCS-7280CR3AM-24D12, DCS-7280CR3AM-32S, DCS-7280SR3A-48YC8, DCS-7280SR3AK-48YC8 and DCS-7280SR3AM-48YC8

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

SKUs Affected: DCS-7050SX3-48YC12

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

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

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

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

SKUs Affected: DCS-7060DX5-64S

- If an SFP slot contains an optic or SFP BASE-T transceiver which is down, or if an SFP slot previously contained an optic or SFP BASE-T transceiver which was down prior to its removal, and the command "reload fast-boot" is performed, the interface may fail to link up after reload.

To work around this problem, hitful restart of the forwarding agent is required, leading to the flap of all interfaces. (1089844) (1)

Series Affected: DCS-7260X3 Series

- The CLI commands "speed 100mhalf/10mfull/10mhalf" will error with the message "Speed and duplex settings are not compatible" for interfaces on the dut. (1203646)

SKUs Affected: CCS-710P-12

- ZeroTouch can fail to link up on front panel interfaces.

Workaround is to use the management port for bootstrap connectivity.

(1207819)

SKUs Affected: DCS-7280CR3-96 and DCS-7280CR3K-96

- If 50g-2 is configured with FEC disabled when the inserted transceiver does not support 50g-2, the Strata agent could restart unexpectedly, causing all of the links on the system to go down.

The workaround is to change the configuration back to a speed supported by the inserted transceiver, or to ensure 50g-2 is supported by the inserted transceiver before configuring 50g-2. (1212939)

SKUs Affected: DCS-7060CX5-56D8, DCS-7280CR3A-24D12, DCS-7280CR3A-32S, DCS-7280CR3AK-24D12, DCS-7280CR3AK-24D12-S, DCS-7280CR3AK-32S, DCS-7280CR3AM-24D12, DCS-7280CR3AM-32S, DCS-7280SR3A-48YC8, DCS-7280SR3AK-48YC8 and DCS-7280SR3AM-48YC8

- Interfaces may fail to link up with 1000BASE-X optics, if they were ever configured with 100M speed before.

Workaround is to configure the following on the affected interfaces:

speed 1g

phy media base-x parallel-detection disabled

These will force the interface speed to 1Gbps, and will disable auto-negotiation and the parallel detection feature on the interface.

If disabling auto-negotiation as a workaround is not desirable, the following command can be run once to fix this issue. Please note that this will cause link flaps and brief traffic disruption on interfaces in the range Et1-Et40.

reset phy interface Ethernet1 chip bcm54185 (1272148)

SKUs Affected: CCS-720DF-48Y

- During ZTP, customer bootstrap scripts that involve configuring interfaces may fail to work as intended causing links to not come up consistently. (1273225)
- Insertion of a BASE-X or BASE-T adapter into a port with default speed configuration may unexpectedly inactivate other interfaces.

The workaround is to explicitly configure the speed on the interface before inserting the transceivers. (1273868) (1)

SKUs Affected: DCS-7280CR3-32D4, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32D4A, DCS-7280CR3K-32D4A-F, DCS-7280CR3K-32D4A-R, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-32P4A, DCS-7280CR3K-32P4A-F, DCS-7280CR3K-32P4A-R, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F,

DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4A-S, DCS-7280CR3MK-32D4A-S-F, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

- Interfaces may fail to link up at 10G speed.

Workaround is to restart the Evora agent using the ``agent AGENTNAME terminate`` CLI command. The agent name(s) to restart on each affected platform and/or interface ranges are detailed below.

For the 7500RM-36CQ-LC SKU: Evora-LinecardX where LinecardX is the linecard with the affected interface(s).

For the DCS-7280SR3M-48YC8, DCS-7280SR3MK-48YC8A-S SKUs: Evora-FixedSystem.

For the DCS-7050CX3M-32S SKU, the agent name for the affected interface(s) will be as follows:

- For interfaces of ports Et1-Et8: Evora-FixedSystem-0
- For interfaces of ports Et9-Et16: Evora-FixedSystem-1
- For interfaces of ports Et17-Et24: Evora-FixedSystem-2
- For interfaces of ports Et25-Et32: Evora-FixedSystem-3

Please note that the Evora agent restart will cause link flaps and brief traffic disruption on the affected interface ranges. (1280270)

SKUs Affected: DCS-7050CX3M-32S, DCS-7280SR3M-48YC8 and DCS-7280SR3MK-48YC8A-S

- An upgrade or downgrade may introduce different default modes to the secondary interfaces on a qsfPdd port. These interfaces therefore may become errrdisabed or operable according to the default modes.

The workaround is to configure forced speeds on these interfaces. (1284935)

(1)

SKUs Affected: DCS-7060DX4-32-F, DCS-7060DX4-32C-RV3-F and DCS-7060PX4-32-F

- OpenConfig may return inaccurate PHY Pre-FEC bit error rate, FEC uncorrected codewords, or FEC corrected codewords. (1290065) (1)

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

Layer 2

- When running rapid-PVST at high scales of interfaces and vlans, the device may occasionally not transmit a BPDU, which can lead to a BPDU timeout on the neighbor. (442663)
- Performing an SSO switchover on an MLAG primary peer running Multiple Spanning Tree Protocol may result in STP topology reconverging. (470962)

- Restarting STP agent with a high number of interfaces and vlans in PVST mode may result in STP topology reconverging. (472675)
- Pseudowire agent may restart unexpectedly when a large number of flexible cross-connect pseudowires are configured/unconfigured. (732883)
- MACsec configured interfaces might not be able to generate SAK post SSO.

The interfaces can be recovered by flapping. (854366)

- On a MLAG switch where 'ip virtual-router' command is configured with a subnet,
if MLAG is split, Arp sync for the host ip covered by the subnet still happens.
The result is that the MAC of the host will move on the connected harness switch,
thus it can cause disruption for traffic destined to the host on the harness switch. (951951)
- In an EVPN VXLAN deployment, if a sub interface with dot1q-tunnel configuration is used for VXLAN forwarding, then the c-tag of the overlay packet will not be preserved. (968892)
- In case there are 802.1X MBA supplicants connected to the switch and an SSU is performed, then there is a possibility that these MBA device might face traffic loss due to the AAA server name not being resolved in time. A simple workaround for this is to configure the server IP directly instead of using the hostname. (976313)
- The Mvrp agent may restart when a large number of VLANs are registered on several interfaces at a frequent interval. This can lead to traffic loss. (1034789)
- The Mvrp agent may experience continuous restarts when a large number of VLAN registrations occur. This issue persists even after the VLAN Join PDUs are stopped and can cause traffic loss. (1036730)
- The Mvrp agent may experience continuous restarts when a large number of VLAN registrations are sent on several interfaces. (1045352)
- Dot1x memory usage increases gradually when periodic accounting is configured with a low interval on a switch with a large number of supplicants.
Increasing the periodic accounting interval may help reduce Dot1x memory usage. (1202455)
- The Macsec agent may not process all MKPDUs in some scenarios. (1203696)

- If supplicants are rapidly moving between routed ports configured with dot1x, they may be incorrectly assigned to the wrong internal VLAN and have their traffic dropped (1218914)
- Enabling Traffic Trap to CPU on an inactive multi-lane interface may cause the `StrataMirror` agent to unexpectedly restart continuously until the configuration is removed. This condition may occur if traffic trap is enabled on a subordinate interface of a multilane module, followed by a speed or mode change that causes these interfaces to disappear. Similarly, this may also occur if the traffic trap mode is forcibly configured on an inactive interface using either a configuration replace command, or by exposing these interfaces with the `service interface inactive expose` command.

To work around this issue, use the `service interface inactive expose` to expose the inactive interfaces, and then remove traffic trap mode on them with the `no forwarding action trap` command. (1261276) (1)

SKUs Affected: CCS-720XP-48TXH-2C-S-F, CCS-720XP-48ZC2-F, CCS-720XP-96ZC2-F and CCS-720XP-96ZC2-M-S-F

- In an EVPN VXLAN deployment with extended VLANs, an incoming VXLAN-encapsulated ARP packet with a VNI mapped to an extended VLAN can cause the VxlanSwFwd agent to restart. (1283418)
- Customers using RADIUS servers connecting over TLS, may see RADIUS requests not going out of the switch. The symptoms includes "Messages received:" counter not changing and the "Requests timeout:" counter keeps increasing in "show radius" command output.
The customer might also see one of the following error logs in /var/log/radsecproxy-<vrf>.log
 1. dossslwrite: SSL: error:0A00007F:SSL routines::bad write retry
 2. dossslwrite: SSL: error:0A00010F:SSL routines::bad length
 3. dossslwrite: SSL: error:0A0C0103:SSL routines::internal error
 4. dossslwrite: SSL: error:1C800066:Provider routines::cipher operation failed

The workaround is to disable and enable the RADIUS server configuration by doing

```
switch(config)#no radius-server host <hostname/ip-addr> tls ssl-profile <profile>
switch(config)#radius-server host <hostname/ip-addr> tls ssl-profile <profile>
(1288640)
```

- Access request with invalid password length is not marked as bad request by Radius Proxy. (1295643)

Layer 3

- With scaled VRF config, in the order of 1K VRFs, if the Arp agent restarts, it may fail to come up. (275514)
- In certain scale situations upon a sysdb restart, arp entries can be learned in kernel but not show up under 'show arp'. When this happens, the work around is to restart the Arp agent. (275749)
- Configuring 'max-metric router-lsa' on an OSPF router may cause summary LSAs to be generated with maximum metric. (276629)
- Auto discovery of multiple IPv6 link-local BGP neighbors over a single interface does not work and will result in only one of the neighbors transitioning to established state. (289875)
- BFD configurations with large numbers of sessions may see BFD session flaps during initial session bringup shortly after boot. (375773)
- Removing SVI configuration at scale via a config-replace operation may cause the BFD agent to unexpectedly restart. (414702)
- The MLAG agent may exit unexpectedly after upgrading from 4.22.0 or an older release if the MLAG peer switch is still running the older software version and DHCPv6 prefix delegation routes are being sync'd across the peers. (425810)
- The output of 'show ip nat sync peer' may incorrectly indicate a connected state, even though an incorrect peer IP address has been configured. (464011)
- DHCP relay agent might exit unexpectedly if several VRFs are continuously deleted and re-created. (471169)
- When large number of distribute lists are configured in RIP then it may result in route flaps and the protocol agent CPU utilization can increase significantly. (536500)
- KernelFib agent may restart on interface flap, if there are a large number of route resolving through that interface (539757)
- KernelFib agent may restart, if EVPN config is toggled in quick succession. (557358)
- RSVP tunnels might not get established or may use non-best paths when IS-IS system ID or OSPF router ID of any two neighbouring nodes in the IGP

network is changed back-to-back and the best path goes via that node.
(572080)

- KernelFib agent may restart when interface flaps lead to a large number of routes (entire routing table) being withdrawn (597557)
- Moving the managementactive(ma0) interface to a VRF and restarting Sysdb would lead to managementactive interface staying down in the default VRF after Sysdb(and other agents) come back up. (613014)
- During Zero Touch Provisioning(ZTP) in an IPv6 environment, even if the Router Advertisement(RA) packets do not have the M and the O bits set, the Arista switch still tries to contact a DHCPv6 server for information other than the IPv6 address. (691387)
- Upon the Arp agent starting up in setups with a large number of routes, the Arp agent may restart. (764196)
- After reload, static ARP or ND entries may not get correctly restored.

Workaround is to restart the Arp agent. (790370)

- Router Solicitation (RS) packets with the source address that is not the unspecified address and no source link-layer address option, causes a Router Advertisement (RA) to be sent with the destination address as unicast address and the destination link-layer address as all-nodes multicast address. (802490)
- After SSU, management routes may not be programmed in the kernel. (911548)
- OSPF may incorrectly export its topology after being disabled and enabled. Restart the OSPF agent to clear the error state. (964629)
- If there is a high churn in IPv4/IPv6 routes in the FIB while the L3 forwarding agent is down, the L3 forwarding agent might restart unexpectedly when the agent comes back up. (985291)
- OpenConfig/Octa agent may restart continuously if there's at least one static route that recursively resolves to a DROP/Null0 route in the FIB.

Workaround:

Directly set the next hop to DROP/Null0 for the static route that resolves over another prefix whose next hop is DROP/Null0.

For example, if "ip route <pfx1> null0" is the first prefix and "pfx2" resolves over "pfx1", ("ip route <pfx2> <pfx1>"), make pfx2 point to Null0 directly, ie. "ip route <pfx2> null0" (1003004)

- The CLI "ipv6 nd dad [DAD_VALUE]" may not work as expected.
The workaround is to apply this config before enabling IPv6 processing on

the interface, or disable and then enable IPv6 processing on the interface after applying this config. (1019418)

- With TI-LFA node protection configured along with SRLG protection, if a destination has a path other than the primary path and backup path and if the first link of this path has the same SRLG as that of the first link of the primary path and with lower cost than the backup path, then when the primary path fails there may be microloops causing traffic loss. (1028897)
- IS-IS authentication will stop working after reload of the router if "isis enable" command is part of "interface profile" configuration and "isis authentication" command is part of interface configuration

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

- RSVP and TI-LFA tunnels may not come up when the OSPF router id is changed. Restarting the OSPF agent will clear the error state. (1042820)
- With TI-LFA FRR enabled along with local convergence delay config and when the protected resource fails, LFIB routes may get updated immediately instead of applying local convergence delay and so can cause microloops. Affected LFIB routes are the ones which have TI-LFA backup paths programmed for local or remote adjacency segments which are reachable via ECMP and these ECMP paths share the same failure risk, like same SRLG or same nexthop router. (1086677)
- When IS-IS SR node segments are leaked across levels then the TI-LFA backup path may not get calculated for such leaked segments. (1156591)
- If the CLB flow is deleted and re-added because of its destination reachability changing then EOS may not delete the CLB flows correctly when aging timer expires.
As a workaround the flows can be manually cleared from the CLI. (1167368)
- OSPF may remove an inter area route when there are multiple LSAs advertising the same prefix with a distribute list configured. (1193137)
- When MPLS label stack size of TI-LFA backup path (till Q node) along with service label (if used) matches maximum sid depth (MSD) and "hardware nexthop fast-failover" config is enabled, then the primary path itself may get un-programmed in HW. (1200552)
- If a BGP peering that is covered by a BFD session becomes reestablished after going down due to the BFD session going down, BGP routes learned from this peer may be installed as drops in hardware while the BFD session is still down. As a workaround, users can configure "no ip hardware fib next-hop update event bfd" to disable BFD event tracking in the L3 forwarding agent

followed by running "start hardware ale review adj l3" to reprocess all affected adjacencies. (1210154)

- SRv6 Agent can restart when EOS is configured with large number of locators. (1227379)
- With SR and Flex-Algo features enabled, Isis agent may repeatedly restart if an IS-IS instance is configured to participate in ten or more flexible algorithms. Receiving at least ten non-zero flexible algorithms at the start of SR-Algorithm sub-TLV of Router Capability TLV from another router can also cause restart of the Isis agent. (1257695)
- Changing the IS-IS preference will not affect the preference of IS-IS FlexAlgo tunnels until the path of the flex-algo changes and routes are recomputed (1264996)
- In case CSPF agent restarts while local-convergence-delay timer is running, the IS-IS TI-LFA backup paths that are currently in use and became invalid while Cspf was not running may still be used, causing traffic loss until the timer expires. (1271765)
- If nexthop group counters are checked immediately after updating a nexthop group, there is a window in which a subset of nexthop group entries may not have their packet counters available. (1273796)
- In case CSPF agent restarts while local-convergence-delay timer is running, the OSPF TI-LFA backup paths that are currently in use and became invalid while Cspf was not running may still be used, causing traffic loss until the timer expires. (1288281)

Multi-agent

- Performing a stateful switchover in a switch with a large number of VRFs configured, with BGP config in all VRFs and IGP configured in a large number of VRFs, may cause the system to restart unexpectedly, leading to traffic loss. (376773)
- SSO with MPLS L3 VPN configured may result in temporary loss of traffic flowing from MPLS core towards the tenant sites (507353)
- Performing SSU may lead to traffic drop if the dynamic BGP peers are configured.

There is no workaround for this issue. Configuring the following may reduce the chances of traffic drop:

- configure router-id under 'router bgp' mode
- configure 'graceful-restart restart-time' and 'graceful-restart stalepath-time' value 420 sec or higher.
- configure 'bgp convergence slow-peer time' value 300 sec or higher.

(843834)

Series Affected: CCS-720XP and DCS-7010TX Series

- Under rare circumstances, Bgp Agent may restart when VRF configuration is removed (845945)
- Running "show mpls bgp labeled-unicast bindings detail" may cause ConfigAgent to restart unexpectedly if there are too many bindings. (918712)
- The BGP inheritance model requires BGP to be configured in default VRF before it can be configured in non-default VRFs.
The workaround is to enable BGP in the default network-instance (AS number at minimum) before enabling BGP in non-default network-instance. (927063)
- If using BGP MPLS VPN and the neighbor received attributes discard command is enabled to discard attribute 16 (BGP Extended Community), then all previously received and imported VPN routes should be removed from any VRF where they were imported, since the Route Target Extended Community is now removed. However, on scaled setups, it is possible that some routes may remain imported in some VRFs. (937337)
- If MPLS IP is deconfigured and then reconfigured, the encap MPLS labels of one or more VPLS BGP-signaled pseudowires may not be correctly programmed. Any affected VPLS BGP-signaled pseudowire groups need to be deconfigured and then reconfigured to resolve the issue. (976290)
- The TI-LFA backup path will no longer be computed if the nodes along the path from PLR to the Q node are not segment-routing enabled. (986875)
- If the Bgp agent is repeatedly restarted while a VPLS pseudowire group using BGP signaling is configured, the Bgp agent restart may be hitfull for all AFI-SAFIs. (1004708)
- The OSPF3 agent may restart when an extended community list with many entries is configured. (1022316)
- In rare situations, when addresses in the 127.0.0.0/8 reserved block are transmitted by protocols or used in certain configurations, the Ira agent may unexpectedly restart until the offending address is removed.
Consider using addresses from private-use address blocks like 10.0.0.0/8 (see <https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml>) instead. (1064362)
- Rcf policy for redistributed routes into non-default Isis VRF instance may sometimes not work, especially after Rcf agent restart (1070669)
- During SSU, link-local BGP peering sessions - configured using 'neighbor interface ...' command - may not be reestablished before BGP convergence timeout which may result in traffic loss until the router neighbors are learned via router advertisement messages. (1084729)

- The OSPF agent may restart when segment routing is configured and an interface flaps when an OSPF adjacency is in the INIT state. (1173862)
- With Flex-Algo configured, Cspf agent may sometimes restart unexpectedly when the Isis agent restarts. (1174057)
- When a hierarchical BGP aggregate, A, is configured for a prefix, P, after BGP has started, the covering aggregate may not get updated correctly if a more preferable competing path, R, already exists for P. If the competing aggregate, A, is later deconfigured, then R may fail to be treated as a contributor for the covering aggregate. At this point, if the covering aggregate originally used the "summary-only" option and this option is deconfigured, R may fail to get advertised.
These problems may be resolved by causing the path R to be removing and re-added, or by deconfiguring and reconfiguring the covering aggregate. (1174495)
- When 'show tech-support bgp' command is running and waiting at 'more' prompt, if BGP agent restarts, all BGP routes become invalid.
The workaround is to end or quit the 'show tech-support bgp' command. (1223491)
- Cspf agent may repeatedly restart when the first include-hop which is present on the source node itself (the include hop could be an interface IP address or TE Router ID of the source node) is specified under the path constraints for an RSVP LER tunnel.

Workaround is to remove the first include-hop and the Cspf agent may need to be manually restarted. (1229036)

- SR-TE ENLP (explicit-null label policy) config change might result in an explicit null label to not be present at the bottom of the label stack when it is expected. (1246649)
- VPLS-BGP routes may be incorrectly marked as invalid if the same VPLS-BGP route is received from multiple route reflectors while ECMP is enabled (1262973)
- When IS-IS SR is configured and
management api models
 provider isis
 link-state-database
is configured, Isis agent memory may grow without bound. The rate of growth is proportional to the rate of IS-IS LSDB churn. (1280448)
- After the outbound policy of a BGP peer is modified through configuration, if during policy reapplication a route is updated and the new path to be advertised for the route is one that is received from the peer undergoing the policy change, then a withdrawal for the route may not be sent to the peer. The peer will be left with a stale path that will not be removed

until a new path for the same route needs to be advertised to the peer.

Workaround: A soft clear of the BGP session should delete the stale path.
(1295165)

Rib

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

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

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

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

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

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

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

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

resulting in traffic loss despite the affected routes appearing correctly in "show ip(v6) route" output.

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

MLAG

- If non-MLAG reload-delay is configured to be lower than the MLAG reload-delay when LACP standby is used, traffic entering the non-MLAG interfaces destined towards hosts on the MLAG interfaces will be lost during the LACP standby period. (83330)
- MLAG ISSU breaks on a switch running a version of EOS < 4.20.5 if it receives ARP entries from the VXLAN Control Service (VCS) while the peer has been upgraded to EOS version >= 4.20.5 (497776)
- When the MLAG agent restarts and the MLAG domain ID is changed at roughly the same time (within 250ms or so), MLAG may report an Active and Connected state in "show mlag", despite the differing domain IDs between peers. This will break the exchange of all MLAG state between peers.

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

MPLS

- MLDP mappings may be temporarily withdrawn and then sent again when a FEC transitions from bud role to leaf role. (979220)
- Traffic loss may be observed when the Mpls agent restarts due to route vias being temporarily removed. (1168777)
- When a local link failure happens MPLS traffic destined to a global adjacency segment might get dropped for the duration of local convergence delay. (1202973)

Multi-domain Segmentation Services

- Interface traffic-policies cannot match on VXLAN decapsulated bridge packets. (913467) (1)

- The throughput of the device may be reduced if traffic is hashed to interfaces that are not bound to dedicated CPU cores. (1273660)
SKUs Affected: ZTX-7230S-4TX-4S and ZTX-7230S-4TX-4S-F

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

Multicast

- An S,G route may not properly cleanup previously sent S,G joins. This would prevent S,Gs from being deleted in the upstream router. A workaround is to flap the interface in which the join is sent on. (291756)
- If large number of IGMP reports are received in a very short interval, some of them might get dropped resulting in the multicast route not being created. The routes will be re-learned on the subsequent query interval. If "show cpu counters queue" shows consistent drops under "CoppSystemIgmp", consider increasing the bandwidth for this Copp class. (356675)
- show ip igmp vrf <name> groups command does not filter the IGMP interfaces as expected. (368071)
- High CPU utilization might be observed for the forwarding agent when PIM Bidir routes are present. (479119)
- Removing router pim bsr also removes router pim sparse mode configuration. Workaround is to just remove specific configuration under router pim bsr mode. (516188)
- Configuring "ip igmp host-proxy <ACL-NAME>" where the access list has a large multicast group prefix length may cause the IgmpHostProxy to be repeatedly restarted resulting in traffic loss.
Use smaller prefix lengths to work around the problem. (586865)
- In certain topologies with EVPN Multicast deployments and Multicast Transit feature enabled, with the RP in external network, learns about the source through PIM SGRPT prunes rather than PIM Register messages, RP may never join the source tree later upon receiving IGMP joins.

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

- MVPN related multicast route might be programmed incorrectly after BGP restart on a scaled MVPN PE (677280)
- On devices running EVPN VXLAN multicast OISM, Overlay multicast traffic with TTL1 arriving on VXLAN enabled VLAN in a VRF with `evpn multicast` enabled is not VXLAN encapsulated when there is a receiver in the same VLAN on remote VTEP.

Workaround

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

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

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

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

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

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

- On 64 bit platforms, with Bess and EVPN Multicast configurations, scale of greater than 20 VRFs might cause system to run out of memory.

Workaround is to keep the VRF scale to below 20 VRFs. (763516)

- In an MVPN setup, traffic at ingress PE might get dropped if multiple flaps occur on the source interface. (777298)
- Using the MLD L2 Querier at scale larger than 30k multicast groups may result in some missing group memberships. (783159)
- On devices running multicast with MRIB lookup enabled (`rpf resolution ribs multicast-rib`), when the MRIB lookup yields a connected route it is not used as the RPF and results in traffic not being forwarded. (783549)
- BessMgr agent might restart when the multicast route scale exceeds 128K routes. (815858)
- Pimsm agent might restart when processing leaves for over 60K groups at the same time. (816562)

- IPsec Tunnels may Flap when disabling Multicast Routing with Traffic flowing over IPsec tunnels (882022)
- When a router is acting as DR + RP + MSDP peer, the receivers connected to the peer may experience delay in receiving the traffic. (883260)
- When SSU is attempted with EVPN multicast configuration, the traffic loss for the multicast traffic might last for a duration of about 10 mins or more. (884878)
- When multicast routing is disabled on default VRF, a small loss of traffic for 2-3sec might be seen for unrelated traffic flows flowing through IPsec tunnels configured on non-default VRF. (892478)
- After reload, at scale and if multicast is enabled, there might be brief flapping of IPsec tunnels. (892513)
- On devices running EVPN L3 multicast with PIM -ASM underlay, persistent multicast traffic duplication may be seen.
To avoid this issue, consider having PIM-SSM in underlay by configuring `vxlan multicast protocol pim ssm`. (899871)
- On devices running Multicast with Kernel based software forwarding, toggling multicast routing configuration may result in complete and permanent loss of traffic.
Workaround is to use `software-forwarding sfe`. Note that changing this config will require a reload. (914121)
- Executing "platform trident forwarding-table partition flexible ip-multicast" might result in stale multicast routes causing packet drops and RPF failures.

Work around:

Toggle `routing` under `router multicast` for the affected VRF. (918123)

- In rare situations. restarting forwarding agent might result in stale multicast routes (ones existing before the restart was executed) being programmed in the hardware resulting in traffic loss.

Workaround: Restart KernelMfib agent. (967832)

- When software multicast forwarding is set to 'sfe', and a virtual IPv6 address is configured on one of the interfaces, after the software forwarding agent, BessMgr, restart, the virtual address is marked as duplicate by the IPv6 protocol, and all multicast routes that depend on this address will send PIM joins in a wrong direction.
Toggling the affected interface resolves the issue. (989536)

- Restarting the PEG (PIM External Gateway) DR (designated router) in an EVPN Multicast OISM setup may lead to a period of duplicate traffic from the PEG DR and any PEG non-DR. (1052179)
- On devices running EVPN OISM on a PEG, change in the assignment of the dynamic VLAN (SBD VLAN) for the VRF might result in traffic loss. Workaround is to restart Pimsm agent to recover. (1056971)
- IGMP protocol agent may crash during configuration replacement when the number of IGMP interfaces added is large. The agent recovers after the initial crash. (1066748)
- With EVPN multicast configured, first few packets that are sent to CPU till the multicast route is learnt in hardware will not be routed. (1077713) (1)
- In EVPN Multicast transit network with a RP on a stick topology, some multicast receivers connected to the PIM External Gateway that is on the path to the source from the RP may not get traffic. The workaround is to add a receiver on another PIM External Gateway. (1079151) (1)
- On MLAG devices running EVPN OISM, receivers connected to orphan VLANs may not receive multicast traffic.
To work around issue, configure the VLAN and SVI and enable multicast on the MLAG peer as well. (1089742)
- On devices running EVPN OISM when the dynamic VLAN(SBD VLAN) for the VRF changes, traffic loss might be observed.
Restart BessMgr to recover. (1095059)
- On MLAG devices running EVPN OISM, receivers connected to orphan VLANs may not receive multicast traffic.
To work around issue, configure the VLAN and SVI and enable multicast on the MLAG peer as well. (1104327)
- Doing config replace with EVPN Multicast and PIM BFD configurations present may cause duplicate multicast traffic.
Restart PIM agent to fix the problem. (1105152)
- With EVPN multicast configuration present, 100% traffic loss might be observed after config replace where the switch receives the traffic but fails to forward because of missing (S,G) entry.
Workaround : restarting the BessMgr agent fixes the issue. (1108686)
- When multicast resources are temporarily exhausted in the hardware, some multicast routes may remain unprogrammed even when resources are freed.
Workaround is to clear the affected mroutes. (1150435)
- With MVPN configuration present with a scale of over 2000 (S,G)s and the underlay in constant churn, SandMcast agent might restart. Fix the churn and then remove and re-add the problem tunnel as a workaround. (1170521)

- In the case where a multicast source is directly connected to the RP and spt-threshold is configured to infinity, it is possible that the multicast traffic is dropped ending in non delivery. (1195447)
- When a burst of IGMP reports and join arrive, while the CPU is busy, it is possible to drop outgoing IGMP reports/query leading to multicast traffic loss. (1196893)
- When a box has multicast traffic ingress on a PIM MBR interface and also has a default route to a PIM neighbor, if the connectivity to the MBR interface gets shut down and then restored after the MBR route expiry, the multicast traffic may not recover.
Workaround: do 'clear ip mroute' for the fastdrop routes. (1252056)
- On devices running EVPN OISM with underlay VXLAN multicast, having only an S,G Exclude overlay receiver might result in not joining the underlay tree.
Workaround: add a static IGMP/MLD snooping receiver to force joining the underlay tree. (1258360) (1)

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

NAT

- Starting from the EOS release 4.30.2, the DCS-7170 switches have adopted the new native NAT infrastructure, which does not relies on the linux kernel for the processing and creation of NAT entries. However, the new NAT infrastructure has a problem where packets that require kernel processing and NAT translation are not processed successfully, because the kernel is now unaware of the NAT entries.
One such case is the processing of ICMP packets in NAT overflow mode. With NAT overload, the IP address used in the NAT translations is the address assigned to the layer 3 interface where NAT is configured. Given that ICMP entries are currently not programmed in hardware, in case of ping messages, when the switch receives an echo reply, it automatically forwards it to the kernel. The kernel, however, is unaware of NAT, so it cannot forward the reply back to the ping originator.
The problem has currently no fix. One simple workaround is to modify the NAT configuration and using a simple IP NAT pool with a single IP address rather than using the NAT overload configuration. (732248)
Series Affected: DCS-7170 Series
- When "ip nat tcp flow creation trigger syn" is set, if dynamic NAT configuration is modified or if the admin shutdown/no shutdown performed on an interface configured with dynamic NAT while NAT traffic is flowing, NAT connections may be left untranslated.

Workaround is to manually flush all the entries using the command "contrack

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

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

Workaround is to manually flush all the NAT entries using the command "conntrack -F" in the corresponding network namespace bash mode. If configured on a non-default VRF (vrfName, for example), run the command "ip netns exec ns-vrfName conntrack -F". (1122152) (1)

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

SwitchApp

- Flaps of interfaces with multicast NAT configured may result in a brief period of resource exhaustion. If the translations configured is at a supported scale, the system should recover automatically without user intervention. (1017562)
- Upon a large number of multicast route churn in a short duration, the MakoMcast agent may unexpectedly restart. The system is expected to stabilize and recover automatically. (1044933)
- Sub-interfaces are not supported, and their configuration will cause the MakoMcast agent to continuously restart. This will mean no new multicast routes can be programmed in the hardware, and continuously disruption to existing multicast routes.

The workaround is to remove the unsupported sub-interface configurations.

(1060292)

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

- The number of multicast (S, G) routes we are able to install in the hardware may vary between reboots - a set of routes that previously fitted in the table may not fit after a reboot and this may occur at a utilisation that is significantly lower than the expected 90%.

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

- Upon many IGMP snooping membership or multicast route churn in a small duration, the VLAN Rewrite table may reach exhaustion resulting in multicast routes not being programmed. MROUTE-3-HW_RESOURCE_FULL logging messages will also be displayed. The membership and route churn causes VLAN Rewrite table resources (VlanGroupId) to be allocated inefficiently, and exhaustion is reached even when there is a small number of trunk ports that make up outgoing interface lists (OIL) for multicast replication. (1079108)
- If the SwitchApp FPGA profile is changed from one supporting more VRFs to one supporting fewer VRFs (e.g. from layer3-mid-latency to layer3-full) and some non-default VRFs are configured then MakoL3Unicast agent may experience continuous restarts.

The workaround is to reload the device assuming the incorrect config is not saved to startup-config. (1164157)

- On the SwitchApp layer3-mid-latency profile (no NAT), TCP and UDP packets directing to a LAG could be duplicated and egress out of 2 members of the LAG. (1175465)
- If a port-channel is shutdown while the LAG agent is not operational, the port-channel will remain configured in hardware and in an active state, even after the LAG is restarted.

The workaround is to delete the affected port-channel and reconfigure it after the LAG agent is restarted. (1181640)

- SwitchApp ACL does not support matching multiple l4 ports but the CLI does not return a failure when such rule is used in an ACL. (1246021)
- SwitchApp ACL does not support keyword "neq" in rules but the CLI does not return a failure when such rule is configured. (1246022)
- SwitchApp ACL doesn't support matching TCP flags but the CLI doesn't return a failure when any TCP flag rule is configured. (1246044)
- The "unidirectional receive-only" configuration may lead to link-up failure on copper cables after a device reboot.

To work around this issue, remove and then re-apply the "unidirectional receive-only" configuration. (1266767)

SKUs Affected: DCS-7132LB-48Y4C

DCS-7130 and DCS-7170 Series

- If a LAG member has "no switchport" configured incorrectly, it may sometimes not be programmed correctly in hardware for VLAN membership checks causing traffic loss.
Workaround is to not configure "no switchport" for LAG members. (855423)
- On trunk ports, PTP packets with the native VLAN will not be dropped when the native VLAN is not in the allow VLAN list. (1226897) (1)
- When egress sampling, the SftAgent will not include the routing attributes (prefix length, nexthop, BGP AS) in the ipfix export. (1235109)
- Under rare circumstances, when Sampled Flow Tracking is enabled with hardware offload, the SftAgent can restart unexpectedly. This is possible either when the sampling rate is changed, or when the platform agents are restarted. (1283454)

DCS-7130 and DCS-7170 Series

- PTP transparent mode is currently not supported. Attempting to configure PTP transparent mode with the `ptp mode e2etransparent` or `ptp mode p2ptransparent` command will cause the BfnSlice agent to unexpectedly restart. (609143)

Series Affected: DCS-7170 Series

- Dynamic (LACP) LAGs may flap during ASU reload, which can result in delayed convergence. (1034873)
- MAP-T profile running on 7170 and performing exception processing on the local CPU, can stop processing local exception traffic if a large number fragments are received.

The only workaround is to restart the SFE agent with "agent SFE terminate". (1241674)

Series Affected: DCS-7170 Series

- A 7170 running MAP-T and handling exception traffic on the local CPU will process the packets at a lower than expected rate.

There is no workaround. (1248437)

Series Affected: DCS-7170 Series

- On 7170, when large updates to MTU configurations are invoked at once, the BfnL3 agent can restart unexpectedly. A workaround is to avoid updating many MTU configurations at once. (1255533)
SKUs Affected: DCS-7170-32C, DCS-7170-32CD, DCS-7170-64C and DCS-7170B-64C
- MAP-t profile running on 7170 and performing exception processing on the local CPU, can stop processing local exception traffic due to packet buffer leaks caused during reassembly of fragments of packets larger than 9252 bytes

The only workaround is to restart the SFE agent with "agent SFE terminate". (1262994)

Series Affected: DCS-7170 Series

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

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

- If link speed is changed on a 2.5G interface, Strata forwarding plane agent may unexpectedly restart with "Process died with signal 3 \ (SIGQUIT\)" signature. Subsequent restart of the Strata agent will resolve the issue. (368424)
- PhyIsland agent might restart unexpectedly multiple times, while configuring 2.5GBASE-T interfaces or during a reload. This will cause link flaps and traffic disruption on all front-panel BASE-T ports. (586920)
SKUs Affected: CCS-720XP-96ZC2, CCS-722XPM-48ZY8 and CCS-750X-48ZP-LC
- During hitless restart of forwarding plane agent(s), packets which were denied by security ACLs may leak through the chip(s) being managed by the agent(s) for a second. The deny security ACL rules start acting on the traffic once the forwarding plane slice agent(s) is warm.

There is no workaround for this bug. (589311)

Series Affected: CCS-750 Series

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

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

Affected ports:

CCS-720XP-48ZC2: 1-40

CCS-720XP-96ZC2: 1-40, 49-88

CCS-720XP-24ZY4: 1-16
CCS-722XPM-48ZY8: 1-48
CCS-750X-48ZP-LC: 1-48
CCS-720DP-48ZS: 1-48
CCS-720DP-24ZS: 1-16 (598487)

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

- EVPN VXLAN A-A multihoming is not supported when MACSec is enabled on the core port. Configuring it may result in undesirable double delivery of BUM traffic. (642181)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- When MACsec is disabled on an interface, multicast traffic may not seamlessly transition to clear text. The work around is to use "shutdown" followed by "no shutdown" after disabling MACsec on an interface. (645584)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- On switchcard switchover, there may be extended traffic outage. (682181)
Series Affected: CCS-750 Series
- 2.5GBASE-T ports configured with `speed 100full` (100M forced speed) configuration may fail to pass traffic even if they report linkup after a failing cable test run with the `test 11 cable base-t interfaces INTFS` CLI command.

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

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

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

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

SKUs Affected: CCS-720XP-24ZY4, CCS-720XP-48TXH-2C-S, CCS-720XP-48ZC2,
CCS-720XP-96ZC2 and CCS-750X-48ZXP-LC

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

The workaround is to reset the affected switch chip using `platform trident <x> reset`, replacing <x> with the name of the switch chip as shown by `show platform trident system port`. Alternatively, `platform trident *

reset` can be used to reset all switch chips in the system. (808236)

Series Affected: CCS-750 Series

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

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

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

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

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

Affected ports:

DCS-7010TX-48: 25-48

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

DCS-7010TX-48C: 25-48

CCS-710P-12: 1-12

CCS-710P-16P: 1-12

CCS-720DP-24S: 1-24

CCS-720DP-48S: 25-48

CCS-720DT-24S: 1-24

CCS-720DT-48S: 25-48

CCS-720XPM-48TH-6SY: 25-48

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

SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DP-24S, CCS-720DP-48S, CCS-720DT-24S, CCS-720XPM-48TH-6SY, CCS-722XPM-48Y4, DCS-7010TX-48, DCS-7010TX-48-DC and DCS-7010TX-48C

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

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

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

- PhyIsland agent restart after an SSO (stateful switchover) may be hitful, causing link flaps and brief traffic disruption on all front panel BASE-T ports of the affected linecard(s). (841719)

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

- After a restart of Strata forwarding plane agents when MACsec is used on an interface which is used for VXLAN underlay, non-unicast traffic may drop continuously. The workaround is to disable and re-enable the VXLAN interface, or restart the StrataL3 agent. (894678)

Series Affected: CCS-722XPM Series

- BASE-T ports may fail to link up when configured with `speed 100full`, even if the link partners report a linkup.

Workaround is to run `shutdown` followed by `no shutdown` on the affected interfaces. (916808)

SKUs Affected: CCS-720XP-48TXH-2C-S and CCS-750X-48ZXP-LC

- BASE-T ports may fail to link up, even if the link partners report a linkup.

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

Affected ports:

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

Series Affected: DCS-7010TX Series

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

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

SKUs Affected: CCS-720XPM-48TH-6SY

- If an SEU occurs on an FPGA and `hardware fpga error linecard action power-cycle` is configured, the card may power-cycle continuously until it is manually reseated or the chassis is rebooted. (1205682)

SKUs Affected: CCS-750X-48SX-LC and CCS-750XPM-48TXH-LC

AWE-5000, AWE-7000 and ZTX-7000 Series

- Port 1 to Port 8 might not work when it is connecting to VMware ESXi server (950484)
SKUs Affected: AWE-5510-F and AWE-7250R-16S-F
- STUN clients may not receive responses from the STUN server if there is an ECMP path from the server to the client. (1016192) (1)
- A Machine Check Error in Coreboot can cause the system to hang without further console output.

The system can be recovered by power cycling. (1044405)

SKUs Affected: AWE-5310-F

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

DCS-7170 Series

- The AlgoMatch forwarding agent will restart when config replace is used to configure access-lists with large number of rules. (240540)
- Following a physical error on link between Jericho and Algomatch FPGA the SandHalo agent may restart, causing ACLs to be removed and reinstalled. (283917)
- When an acl configuration is applied to a large number of SVIs in a configuration session, the acl configuration fails. (367513)

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

- When many IPv6 routes are configured and the BFD protocol in use, BFD sessions may flap when a linecard is removed or shut down. (279129)
- ManagementActive Agent repeatedly restarts and Management0 interface is not reachable in Modular systems. Doing switchover will address the issue. (369565)
- After linecard removal and re-insertion, lag member links may remain inactive due to "waiting for hardware to program port to RX (or receive)".

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

Series Affected: DCS-7300X Series

- Correctable PCIe errors may be seen on a forwarding chip control plane link. (772195)

Series Affected: DCS-7800 Series

- On card OIR, %HARDWARE-3-POWERCONTROLLER_BAD_OUTPUT_VOLTAGE log messages may be syslogged. The card recovers on its own. (1205710)
- In the case of EVPN Bud Role traffic, when an underlay multicast encapsulated packet is received, if the inner multicast packet is an L3 packet and results in a multicast miss, such packets are dropped, causing traffic loss.

The workaround is to issue the command "fab internal_lpbk evpn_bud 65 131". (1272387) (1)

- A malfunctioning card may cause Smbus and related system management / monitoring agents to become unresponsive for some time (1275134)
SKUs Affected: DCS-7800-SUP1A, DCS-7800-SUP1S, DCS-7816-SUP and DCS-7816-SUP1S

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

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

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

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

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

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

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

- Acl and related features that use TCAM for rule matching do not work correctly on packets with IPv6 extension headers if the rules have layer 4 match criteria such as tcp flags and source/destination ports. (221161)
- PBR policies applied on interfaces do not work if the interface has some subinterfaces that also have PBR policies applied. (243286)
Series Affected: DCS-7500R Series
- VTEP cannot learn a remote host's MAC address from VXLAN encapsulated IPv6 Neighbor Advertisement packets received from that host. This might be an issue in a pure IPv6 deployment where the VTEP does not receive any bridged traffic from the host. In such a scenario, the traffic toward the host will be flooded in the VLAN.

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

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

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

- VXLAN bridging traffic sent out of LAG interfaces may be temporarily dropped after an SSO event. (251398)
Series Affected: DCS-7500R Series

- When a DirectFlow flow is created with an action to drop the outer tag, the packet is still leaving the interface tagged. (264336)

Series Affected: DCS-7500R Series

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

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

- If one or more linecards in a modular system is running in tap aggregation mode, a small percentage of multicast traffic on linecards running in normal mode may be dropped, following a change to membership of some tap aggregation destination groups. (275286)
- When the primary and fallback policy are both changed in the same config session, traffic may stop matching on either primary or fallback policy. (287146)
- When primary and fallback policies are both changed in the same config session, some packets will not be subject to either policy for a brief period during the configuration change. (287154)
- Power cycling a 7500E linecard in a chassis with 7500E fabric modules can cause packet loss on 7500R linecards. (295825)
- Sometimes Macsec Proxy interface gets into error disabled state on reload when fips is enabled.

To workaround the issue configure errdisable max flap value; i.e., "errdisable flap-setting cause link-flap max-flaps 24 time 10". Note the setting is global and affects non Macsec Proxy interfaces as well. (396589)

- A switch reload with qos policy configured may result in SandAcl agent restart. (417935)
- Storm control maximum burst size is fixed to 10kB which can lead to storm control drops for bursty incoming traffic. (423231)
- With WRR scheduling on many ports and linerate across multiple traffic classes there can be a small amount of egress buffer drops. (425123)
- While a line card is being power cycled, corrupt VXLAN encapsulated packets may be sent. (441141)
- Performing many updates to a PBR policy in a short amount of time may cause the SandAcl agent to restart unexpectedly. (441912)

- VXLAN routing over MACsec proxy port stops working when a proxy port that is not bound to an external interface is updated to be bound to an external interface. (455759)
- "match nexthop-group" filter is not supported in PBR rules when the corresponding nexthop-group is referenced by ECMP routes. (458189)
- Once the maximum number of supported unique PBR policies per chip is exceeded, removing the PBR policy on an interface, converting the interface to a vlan member and applying a RACL on the SVI does not work. (477753)
- With a tcam profile configuration containing "feature tunnel vxlan routing", SandAcl agent restart may lead to another unexpected restart before the agent recovers. (492320)
Series Affected: DCS-7500R Series
- With MPLS VPN, during stateful switchover, some VPN traffic after MPLS decapsulation may be transiently forwarded in the default VRF instead of the appropriate tenant VRF. (541206)
- An over-voltage condition on the POS1V2_HBM_IO_JE0 power rail can cause the system to power-cycle. (544659)
SKUs Affected: DCS-7280CR3-96 and DCS-7280CR3K-96
- 'Persistent Internal Drop DeqDeletePktCnt' error messages are seen due to incorrect/missing system ACEs. Hitful restart of SandAcl agent is needed to resolve the issue. The trigger for this issue is unknown. (548524)
- On a shaped subinterface tx-queue, configuring the bandwidth doesn't take effect unless no-priority is also configured. (588133)
- If 'cpu traffic-policy <policyName> vrf all' is configured, where the traffic-policy contains rules with the policer action, and the operating TCAM profile does not support the CPU traffic-policy feature, then configuring any routed port in the switch may lead to an unexpected SandAcl agent restart. The workaround is to remove the 'cpu traffic-policy <policyName> vrf all' configuration. (619998)
- Postcard sampling doesn't work for traffic that is forwarded out of a sub-interface that has shaping applied. (621932)
- Postcard Telemetry sampling doesn't work for IPv6 packets if the egress interface has an IPv6 egress ACL applied on it. (624010)
- Continuous shut/no shut on a Port-Channel with 1000s of subinterfaces can cause the SandTunnel agent to restart unexpectedly. (718924)
- When using a large scale of PBR rules, removal and addition of member links used as nexthop groups destinations of PBR policies may cause the SandAcl agent to restart. (759846)

- On a config-replace, a change in the PBR config that qualifies for in-place update may still cause traffic loss. (760884)
- Performing many updates to a traffic policy in a short amount of time may cause the SandAcl and/or SandAegis agent to restart unexpectedly. (770478)
- Flapping the interface that is the destination for a large number of redirect actions in a traffic-policy may result in a long programming delay to updating redirect actions in hardware and to any future changes in the traffic-policy config. (814453)
- In rare instances, entries in a TCAM bank may fail to be deleted before the TCAM bank is released by one feature and then allocated to another feature, resulting in non-deterministic behavior in the new feature. (826437)
- Interface traffic-policies cannot match on VXLAN decapsulated bridge packets. (913467) (1)
- When a host moves rapidly between a local port and a remote VXLAN VTEP, then it is possible that the device stops learning the host locally, causing it to think that the host is permanently located behind the remote VTEP.

Clearing the MAC address table should help fix up the issue. (927167)

- Traffic-policies with match rules that contain OR-match field configurations (i.e., `protocol tcp flags ack not-rst syn` + `protocol udp source port 10 destination port all` in the same match rule) may cause traffic-policies with fewer than 1k rules to fail to apply while in `counter interface per-interface ingress` (counters per-interface) mode. (927775)
- Traffic-policy mac matches do not work on 'packet ipv4 vxlan eth non-ip forwarding bridged decap' and 'packet ipv6 vxlan eth non-ip forwarding bridged decap' packet types. (948990)
- CPU traffic-policy may fail to program if 'cpu traffic-policy POLICY vrf all' is configured on a switch with lots of VRFs and the applied policy contains a large number of matches with police actions. (957530)
- In an EVPN VXLAN network deploying active/active multi-homing using shared Ethernet segments, when the MAC tables on the VTEPs become momentarily inconsistent, there might a brief period of duplicate delivery or packet loss.
Typically, the MAC tables become inconsistent when a new MAC is learnt on one of the VTEPs and is yet to be advertised to other VTEPs or when a MAC is removed on one of the VTEPs and it is yet to be withdrawn from other VTEPs. (957981)

- For traffic-policy with 'update interface hitless strict' enabled, a traffic-policy update when the hardware is near resource exhaustion may fail programming and cause the SandAegis forwarding agent to restart. (985126)
- Sometimes Macsecproxy patch interface incorrectly does not terminate the internal VLAN tag resulting in an inoperational Macsecproxy subinterface.

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

- ARP/ND packets received on a port where double VLAN tag translation is configured might not be correctly headend replicated to remote VTEPs in EVPN or VCS setups when VXLAN ARP/ND proxy feature is enabled (which is default for EVPN setups and optional for VCS setups).

The workaround is to disable the ARP/ND proxy feature. This can be achieved as follows:

```
(1) On EVPN setups:
ip prefix-list disable-arp-proxy
    seq 10 deny 0.0.0.0/0
ipv6 prefix-list disable-neigh-proxy
    seq 10 deny ::0/0
router l2-vpn
    arp proxy prefix-list disable-arp-proxy
    nd proxy prefix-list disable-neigh-proxy
```

```
(2) On VCS setups (if the proxy feature was previously enabled):
interface vxlan 1
    no vxlan controller-client arp proxy
    no vxlan controller-client nd proxy (1015743)
```

- Adding or removing VLAN member interfaces while the SandMcast agent is restarting, such that the final configuration is same as the original one can prevent traffic from flooding in the VLAN correctly. (1050404)
- On a EVPN OISM Vxlan enabled network temporary loss of multicast Vxlan traffic might be observed. Traffic recovers automatically. (1081346)
- If odd-even pair ports (i.e. Et1 and Et2, Et3 and Et4) are configured to be 40G-4, transceiver insertion and removal or speed change into 40G on the even interfaces (Et2, Et4, etc) can cause odd interfaces (Et1, Et3, etc) to flap.

Workaround is to plug in transceivers into the even numbered ports first or ensure the even numbered ports are configured before the odd ports. (1165284)

SKUs Affected: DCS-7280CR3A-24D12, DCS-7280CR3A-24D12-F,
DCS-7280CR3A-24D12-R, DCS-7280CR3A-32S, DCS-7280CR3A-32S-F,
DCS-7280CR3A-32S-R, DCS-7280CR3AK-24D12, DCS-7280CR3AK-24D12-F,

DCS-7280CR3AK-24D12-R, DCS-7280CR3AK-24D12-S, DCS-7280CR3AK-24D12-S-F,
DCS-7280CR3AK-24D12-S-R, DCS-7280CR3AK-32S, DCS-7280CR3AK-32S-F,
DCS-7280CR3AK-32S-R, DCS-7280CR3AM-24D12, DCS-7280CR3AM-24D12-F,
DCS-7280CR3AM-24D12-R, DCS-7280CR3AM-32S, DCS-7280CR3AM-32S-F,
DCS-7280CR3AM-32S-R, DCS-7280SR3A-48YC8, DCS-7280SR3A-48YC8-F,
DCS-7280SR3A-48YC8-R, DCS-7280SR3AK-48YC8, DCS-7280SR3AK-48YC8-F,
DCS-7280SR3AK-48YC8-R, DCS-7280SR3AM-48YC8, DCS-7280SR3AM-48YC8-F and
DCS-7280SR3AM-48YC8-R

- In an EVPN VXLAN A/A multihoming network topology, where an L2 switch (or typically, an L2 switch pair) connects to the EVPN VTEPs, then MAC flaps might be seen on the L2 switch when it receives traffic from a new host and there is subsequent return traffic to that host within a couple of seconds. (1184408)
- When "hardware next-hop fast-failover" is configured, the hardware programming acknowledgement might be premature for certain adjacencies in the presence of any interface flaps. (1193296)
- SandFapNi or B52 may restart continuously when MACsec is enabled with a license on systems with many linecards (1206349)
- In certain EVPN VXLAN setups without OISM configured, multicast routed packets might be corrupted if PIM are enabled in overlay VLANs.

Disable PIM in overlay VLANs to avoid packet corruptions. (1207474)

- TCAM profiles where the traffic-policy feature is configured with key fields that contain a mix of raw prefix, prefix labels, and LPM prefix labels is not supported and should be rejected on policy application. (1215665)
- In EVPN setups with the "redistributed static" CLI command configured, devices will received static configured MAC addresses on peer devices. If these static MAC addresses are dynamically learned from the data plane, SandMact might restart unexpectedly. (1226981)
- When a large number of VRFs are configured, the SandAegis agent may restart after config replace from full startup-config to config reset to full startup-config gain. No action required, automatic agent restart will restore system steady state. (1234964)
- In a chassis with redundant supervisors configured with the SSO redundancy protocol and QoS features: priority-flow-control and weighted-fair-queuing enabled, after a switchover has been performed, there may be some drops on lossless queues due to DeqDelete events. (1245947)
- On a system that has Priority-Flow-Control configured, packets ingressing on a port and destined to CppSystemSflow may be dropped if there are packets

ingressing on the same port and destined to a different congested port.
(1260989)

- Insertion of a BASE-X or BASE-T adapter into a port with default speed configuration may unexpectedly inactivate other interfaces.

The workaround is to explicitly configure the speed on the interface before inserting the transceivers. (1273868) (1)

SKUs Affected: DCS-7280CR3-32D4, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32D4A, DCS-7280CR3K-32D4A-F, DCS-7280CR3K-32D4A-R, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-32P4A, DCS-7280CR3K-32P4A-F, DCS-7280CR3K-32P4A-R, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4A-S, DCS-7280CR3MK-32D4A-S-F, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

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

- During restart of SandAcl agent, packets may not be subjected to ACL rules for a short period of time. (372971)
- ZeroTouch provisioning will fail to fetch configuration on QSFP-DD/OSFP ports with inserted QSFP or SFP transceivers or QSFP ports with inserted SFP transceivers when these ports require a non-default speed-group configuration to link up.

The workaround is to use a different type of transceiver or a different port to connect to the bootstrap server. (488491) (1)

- If tap aggregation and a CPU traffic policy are configured together, either the CPU traffic policy will not apply to any traffic or the SandAcl agent will continuously restart. (492549)
- On applying/removing a Egress TC-to-CoS map on L2 LAG Subinterfaces, the packets going out of the Subinterfaces can have wrong dot1q tag. Workaround is to shut the subinterface, apply the map and then unshut it (582063)
- Following a Stateful Switchover, disabling power to either supervisor using "no power enable module Supervisor" may cause some control plane agents to restart unexpectedly and repeatedly.

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

- When EVPN MPLS is enabled on an interface with `no switchport mac address learning` configured, a packet may have still its source MAC address learned. (616937)
- On a switch configured with IPsec tunnel, restarting the B52 agent or speed change of underlay interface of the tunnel or flapping underlay interface of the tunnel may lead to complete traffic loss or undefined traffic forwarding behaviour on the tunnel.
A workaround is to flap the IPsec tunnel. (980557)
- Removing secure VXLAN profile from a remote VTEP can result in momentary traffic loss on all front panel ports. (1070042)
- The SandACL agent may restart unexpectedly if a Jericho2C+ linecard is replaced with a Jericho2 linecard. (1072612) (1)
- IPsec tunnel configuration on a switch resulting in ROUTING_TUNNEL_SIP_RESOURCE_FULL syslog may also lead to repeated restarts of the SandL3Unicast agent. (1126352) (1)
- Flapping IPsec tunnels may cause other, unrelated IPsec tunnels to flap.
A workaround is to increase the shape rate for copp-system-ipunicast copp class. (1165778) (1)
- Config replace operation to configure IPsec tunnels on a switch with pre-existing secure VXLAN or IPsec tunnel configurations may prevent the tunnels from being established.
A workaround is to remove the existing configuration on the switch before performing the config replace operation. (1184981) (1)
- With scaled route configs, link down/link up events could result in high CPU usage of SandL3Unicast agent. System would recover on its own. However this could sometimes also result in delays in convergence or BGP flaps.

Workaround: Enabling route prioritization feature reduces the impact of the issue and prevent BGP flaps. (1207772)

- In an EVPN VXLAN A/A multi-homing OISM setup, traffic received from a port-channel Ethernet segment might be incorrectly sent back to that Ethernet segment via the multi-homed peer switch. (1267127)

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

- ZeroTouch provisioning will fail to fetch configuration on QSFP-DD/OSFP ports with inserted QSFP or SFP transceivers or QSFP ports with inserted SFP transceivers when these ports require a non-default speed-group configuration to link up.

The workaround is to use a different type of transceiver or a different port to connect to the bootstrap server. (488491) (1)

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

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

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

- The SandTunnel agent can restart unexpectedly when there are more than 48 port channels with members mapped to the same forwarding ASIC core.

To avoid this issue, ensure that no single forwarding ASIC core has more than 48 port channels with members assigned to it. This can be achieved by either reducing the total number of

port channels to 48 or fewer, or by distributing port channel members across multiple cores. (1015676) (1)

DCS-7280R3A and DCS-7800R3A Series

- ZeroTouch provisioning will fail to fetch configuration on QSFP-DD/OSFP ports with inserted QSFP or SFP transceivers or QSFP ports with inserted SFP transceivers when these ports require a non-default speed-group configuration to link up.

The workaround is to use a different type of transceiver or a different port to connect to the bootstrap server. (488491) (1)

- When MACsec is enabled in hardware on a switch, then if the switch ingress experiences congestion due to a large number of small packets at line-rate, the priority drop feature may drop packets with higher priority before packets with lower priority. (685317) (1)

- Packets of size 983B-1083B are only 95% line rate when all flows are crossing both chips. (798619)

SKUs Affected: DCS-7280DR3A-36

- Shutdown of an interface on the peer containing a hardware accelerated BFD session or reload of the same peer may trigger an unexpected SandBfd agent restart causing all hardware acceleration BFD sessions to flap.

Workaround: Use "dut(config-router-bfd)# hardware acceleration disabled" to disable hardware acceleration. (830008)

- Applying speed change configuration on underlay interface of an IPsec tunnel or flapping underlay interface of an IPsec tunnel configured with anti-replay detection may lead to complete traffic loss on the tunnel.

A workaround is to flap the IPsec tunnel. (950629) (1)

- In IPsec tunnel mode, both peers should either enable or disable replay protection. Asymmetric configuration will lead to complete packet loss over the tunnel. (988725) (1)
- The SandTunnel agent can restart unexpectedly when there are more than 48 port channels with members mapped to the same forwarding ASIC core.

To avoid this issue, ensure that no single forwarding ASIC core has more than 48 port channels with members assigned to it. This can be achieved by either reducing the total number of port channels to 48 or fewer, or by distributing port channel members across multiple cores. (1015676) (1)

- The SandACL agent may restart unexpectedly if a Jericho2C+ linecard is replaced with a Jericho2 linecard. (1072612) (1)
- IPsec tunnel configuration on a switch resulting in ROUTING_TUNNEL_SIP_RESOURCE_FULL syslog may also lead to repeated restarts of the SandL3Unicast agent. (1126352) (1)
- Flapping IPsec tunnels may cause other, unrelated IPsec tunnels to flap. A workaround is to increase the shape rate for copp-system-ipunicast copp class. (1165778) (1)
- Config replace operation to configure IPsec tunnels on a switch with pre-existing secure VXLAN or IPsec tunnel configurations may prevent the tunnels from being established. A workaround is to remove the existing configuration on the switch before performing the config replace operation. (1184981) (1)
- When LOC or RDI is detected, "show cfm continuity-check end-point" will display "Interface action: transmit disabled" indicating an interface action was triggered. However, "show interfaces <INTERFACE>" will not have "CFM connectivity status: transmit disabled" which indicate transmit was not disabled on the end point's interface. (1188724)
SKUs Affected: 7289

DCS-7280R3A Series

- ZeroTouch provisioning will fail to fetch configuration on QSFP-DD/OSFP ports with inserted QSFP or SFP transceivers or QSFP ports with inserted SFP transceivers when these ports require a non-default speed-group configuration to link up.

The workaround is to use a different type of transceiver or a different port to connect to the bootstrap server. (488491) (1)

- When MACsec is enabled in hardware on a switch, then if the switch ingress experiences congestion due to a large number of small packets at line-rate, the priority drop feature may drop packets with higher priority before packets with lower priority. (685317) (1)
- Applying speed change configuration on underlay interface of an IPsec tunnel or flapping underlay interface of an IPsec tunnel configured with anti-replay detection may lead to complete traffic loss on the tunnel.
A workaround is to flap the IPsec tunnel. (950629) (1)
- In IPsec tunnel mode, both peers should either enable or disable replay protection. Asymmetric configuration will lead to complete packet loss over the tunnel. (988725) (1)
- The SandTunnel agent can restart unexpectedly when there are more than 48 port channels with members mapped to the same forwarding ASIC core.

To avoid this issue, ensure that no single forwarding ASIC core has more than 48 port channels with members assigned to it. This can be achieved by either reducing the total number of
port channels to 48 or fewer, or by distributing port channel members across multiple cores. (1015676) (1)

- IPsec tunnel configuration on a switch resulting in ROUTING_TUNNEL_SIP_RESOURCE_FULL syslog may also lead to repeated restarts of the SandL3Unicast agent. (1126352) (1)
- Flapping IPsec tunnels may cause other, unrelated IPsec tunnels to flap.
A workaround is to increase the shape rate for copp-system-ipunicast copp class. (1165778) (1)
- Config replace operation to configure IPsec tunnels on a switch with pre-existing secure VXLAN or IPsec tunnel configurations may prevent the tunnels from being established.
A workaround is to remove the existing configuration on the switch before performing the config replace operation. (1184981) (1)

DCS-7280R3 Series

- With Jumbo MTU packets we could have some FCS errors on the peer. (747305)
- The SandTunnel agent can restart unexpectedly when there are more than 48 port channels with members mapped to the same forwarding ASIC core.

To avoid this issue, ensure that no single forwarding ASIC core has more than 48 port channels with members assigned to it. This can be achieved by either reducing the total number of
port channels to 48 or fewer, or by distributing port channel members across multiple cores. (1015676) (1)

DCS-7280R3 Series

- The SandTunnel agent can restart unexpectedly when there are more than 48 port channels with members mapped to the same forwarding ASIC core.

To avoid this issue, ensure that no single forwarding ASIC core has more than 48 port channels with members assigned to it. This can be achieved by either reducing the total number of port channels to 48 or fewer, or by distributing port channel members across multiple cores. (1015676) (1)

DCS-7280R3 Series

- The SandTunnel agent can restart unexpectedly when there are more than 48 port channels with members mapped to the same forwarding ASIC core.

To avoid this issue, ensure that no single forwarding ASIC core has more than 48 port channels with members assigned to it. This can be achieved by either reducing the total number of port channels to 48 or fewer, or by distributing port channel members across multiple cores. (1015676) (1)

DCS-7280R3A Series

- The SandTunnel agent can restart unexpectedly when there are more than 48 port channels with members mapped to the same forwarding ASIC core.

To avoid this issue, ensure that no single forwarding ASIC core has more than 48 port channels with members assigned to it. This can be achieved by either reducing the total number of port channels to 48 or fewer, or by distributing port channel members across multiple cores. (1015676) (1)

- IPsec tunnel configuration on a switch resulting in ROUTING_TUNNEL_SIP_RESOURCE_FULL syslog may also lead to repeated restarts of the SandL3Unicast agent. (1126352) (1)
- Flapping IPsec tunnels may cause other, unrelated IPsec tunnels to flap. A workaround is to increase the shape rate for copp-system-ipunicast copp class. (1165778) (1)
- Config replace operation to configure IPsec tunnels on a switch with pre-existing secure VXLAN or IPsec tunnel configurations may prevent the tunnels from being established. A workaround is to remove the existing configuration on the switch before performing the config replace operation. (1184981) (1)

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

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

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

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

- If VXLAN is configured, PBR and QOS policies won't apply to traffic ingressing on a MLAG peer-link. (390804)
- When a remote MAC is learnt on an Ingress PE with L2 EVPN MPLS configuration, the traffic destined to the remote MAC may get dropped for a brief amount of time. (551995)
- During normal operation, the message 'HARDWARE-3-TIMESYNC_ERROR ... (master frequency check error)' may be logged if the internal time synchronization in a switch is disrupted. (557400)
SKUs Affected: DCS-7800-SUP
- In the presence of virtual MAC with mask configuration, packets that undergo recirculation may not be properly forwarded. (559242)
- The "Last LOC Detected" and "Last LOC Cleared" timestamps shown in the output of `show cfm continuity-check end-point` command get cleared if the SandOam agent is restarted. (629830)
- "mac address virtual-router" command no longer works on 7280R3, 7500R3, 7800R3 devices. (633307)
- Egress IPv6 ACL logging do not work on packets subject to layer 3 MPLS termination. (640491)
- If the 'update interface hitless strict' configuration is active, when the FEC associated with the routes from redirect actions used by interface traffic-policies change and traffic-policy hardware usage is near or above 50% of available resources, the FEC programmed in TCAM may not be updated, causing traffic matching a redirect rule to either misroute or drop. The most common reasons for FEC changes are SandL3Unicast restart and/or SandFapNi restart.

To prevent this issue, it is recommended that traffic-policy hardware usage does not exceed 50% of the available resources. (642948)

- SandAcl agent will restart continuously if a TCAM based feature is being configured/modified on a FAP that has linerate traffic ingressing through it.

For instance, bootup of linecard. Workaround is to reduce the ingress traffic on that FAP. (645398)

- When a VLAN translation rule is configured on a trunk interface and that interface has `no switchport mac address learning` configured, a packet may have its source MAC address learned if it was matched by the VLAN translation rule. (647338)
- When config-replace of configuration having VPLS commands is performed back-to-back in quick succession, the SandMact agent can restart unexpectedly. (690059)
- The SandTunnel agent may unexpected restart when changing TCAM profiles on systems with high scales of sub-interfaces and SVIs. (725198)
- P4Runtime server may not send PacketIn messages to the P4Runtime client for trace route TTL expiry packets. (745925)
- When segment-security is configured at a high scale, Firewall agent may restart unexpectedly when "shut" or "no shutdown" is issued under "router segment-security", or when a VRF goes down and then comes up. (760155)
- Altering MEP attributes, like changing direction, changing profile etc, while CCMs are already running may result in unexpected RDI being reported in the network. A workaround is to restart the SandOam agent on the affected node, i.e. the node sending the RDI after such a churn has been made to the MEP attributes. (766982)
- With EVPN OISM, IPv6 ND link local packets received from the VXLAN tunnel are not flooded to all ports in the VLAN. (772926)
- When removing a port-channel member from a routed port-channel with segment-security configured, traffic may leak through the other port-channel members. (773786)
- CFM Loss Measurement sessions running on MEPs configured on the same interface at different Maintenance Domain Levels may report incorrect losses. (793764)
- The system may reboot unexpectedly when certain types of transceivers are inserted (796309)
 SKUs Affected: DCS-7280CR3A-48D6-F, DCS-7280CR3AK-48D6-F and DCS-7280CR3AM-48D6-F
- If the device ingress experiences congestion due to a large number of small packets at line-rate, packets which are assigned the receive interface's default traffic-class may be dropped before packets with lower traffic-class. (798612)
 SKUs Affected: 7500R3-48Y4D-LC, 7800R3K-72Y-LC, DCS-7130LBR-48S6QD-F,

DCS-7130LBR-48S6QD-R, DCS-7280CR3-36S-F, DCS-7280CR3-36S-R,
DCS-7280SR3-40YC6-F, DCS-7280SR3-40YC6-R, DCS-7280SR3-48YC8-F,
DCS-7280SR3-48YC8-R, DCS-7280SR3E-40YC6-F, DCS-7280SR3E-40YC6-R,
DCS-7280SR3K-48YC8-F, DCS-7280SR3K-48YC8-R, DCS-7280SR3M-48YC8-F and
DCS-7280SR3M-48YC8-R

- In the presence of a large number of VXLAN MAC addresses, using configure replace to change VXLAN dual stack underlay to IPv4 underlay and also changing the TCAM profile at the same time can lead to permanent traffic loss toward some remote hosts. Workaround is to shut and un-shut the SandMact agent to recover from this state. (802920)
Series Affected: DCS-7280R3, DCS-7500R3 and DCS-7800R3 Series
- If both "ip hardware fib optimize" and resilient ECMP are together configured for VXLAN routes, then SandL3Unicast agent might unexpectedly restart. (813004)
- Configurations to disable MAC learning applied to trunk Port-Channels may stop working on VLANs with ACL or policy-maps applied after LAG membership changes. An MLAG peer-link should never learn MACs. If that peer-link is a Port-Channel, then it is susceptible to this issue, and may start to incorrectly learn MACs.

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

- Hardware accelerated sFlow of Vxlan decapped packet miss first 12B from header. (829460)
- On a switch having an interface configured with traffic-policy unshared policer, unconfiguring the traffic-policy while the SandAegis agent is restarting may cause the SandAegis agent to unexpectedly restart one more time. (836424)
- Executing a config replace between two scaled traffic-policy configurations with unshared policer actions may fail. Additionally, the subsequent config rollback may also fail causing the switch to unapply traffic policies configured on an interface. (838906)
- Configuring more than the maximum number of ports and doing a stateful switchover causes counters to not get updated. (855905)
SKUs Affected: DCS-7816-CH
- If "next-hop fallback decapsulation vrf" config is present and routes with decapsulation/fallback nexthops are churned while route scale is high and

resources are full, then the SandL3Unicast agent might unexpectedly restart. (857899)

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

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

- In a VXLAN routing setup using IPv6 underlay, if the overlay SVI is configured with "ipv6 address virtual" then the virtual VTEP functionality will not work correctly. Specifically, neighbor advertisement responses destined to this SVI will fail to sync across to other L3 VTEPs and hence neighbor resolutions may fail. Also, the ND proxy functionality on Aggregation L3 VTEPs in campus deployments will not work. (868642)
- After reloading the switch or inserting a linecard, multicast traffic ingressing on a GRE tunnel may not be forwarded correctly.
A workaround is to restart the SandMcast agent. (880336)
- In the EVPN VXLAN OISM setup, when the number of VTEPs or that of the configured multicast decapsulated groups goes beyond the limit, restarting SandTunnel or doing a removal and insertion of a linecard modular would lead to losses of some traffic flows.

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

- If only egress traffic-policy is configured in the TCAM profile, but ingress traffic-policy is not, running 'show traffic-policy interface counters' for an egress traffic-policy will cause the ConfigAgent to restart. (908122)
- While in Tap Aggregation mode, linerate traffic on tool ports that have "switchport tool allowed vlan" configured may have corrupted FCS. (910965)
- Inserting a rate-adapting module into a QSFP100 port may cause the adjacent QSFP100 port to link down. (934196)
SKUs Affected: DCS-7280CR3A-24D12-F, DCS-7280CR3A-24D12-R, DCS-7280CR3AK-24D12-F, DCS-7280CR3AK-24D12-R, DCS-7280CR3AK-24D12-S-F, DCS-7280CR3AK-24D12-S-R, DCS-7280CR3AM-24D12-F and DCS-7280CR3AM-24D12-R
- When PFC is configured on a device non-PFC traffic drops can happen for packets going to un-congested queues. (935006)
- When performing linecard OIR and the linecard being removed have an egress policy applied to an interface on the linecard and at least 1 other

interface on another linecard has an egress policy applied, the SandAegis agent may restart. (943175)

- When packets are dropped at a high rate in the egress pipeline, then a small rate of packet corruption might be seen, which show up as FCS errors in the neighboring devices. Egress packets drops could happen during MTU enforcement or during split horizon related packet filtering with features such as VXLAN and L2 DCI. There is a higher chance of the issue happening at larger packet sizes. (944373)
- When a parity error happens during a switchover while a monitor session is configured with an IPv4 GRE tunnel destination, the agent SandL3Unicast may restart. (955988)
- If 'next-hop decapsulation vrf' configuration is active and the TCAM profile is changed, TCAM resources may become unprogrammed.

Remove the 'next-hop decapsulation vrf' configuration and restart the Layer 3 Forwarding agent to allow TCAM resources to be programmed. Once the agent is warm, add the 'next-hop decapsulation vrf' configuration again and restart the Layer 3 Forwarding agent a second time. (958505)

- The system may reboot unexpectedly with an under-voltage fault on the POS1V2_HBM_C_JE0 VOUT_UV power rail. (963157)
SKUs Affected: DCS-7280CR3-96-F and DCS-7280CR3K-96-F
- Interfaces may experience intermittent link flaps with certain optics due to noise on the optics power rail. See Field Notice 94 for more information. (997423)
SKUs Affected: DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R, DCS-7280SR3K-48YC8-F, DCS-7280SR3K-48YC8-R, DCS-7280SR3K-48YC8A-F, DCS-7280SR3K-48YC8A-R, DCS-7280SR3M-48YC8-F, DCS-7280SR3M-48YC8-R, DCS-7280SR3MK-48YC8A-S-F and DCS-7280SR3MK-48YC8A-S-R
- Upon configuring trunk ports on ports associated with incompatible speed groups, SandFapNi agent restart, or linecard removal/insertion, the SandTunnel agent may unexpectedly restart. The agent is expect to stablize and recover. (998216)
- Routed interfaces, untagged sub-interfaces, and interfaces with PVLAN configuration are not supported as member ports of a Port-Channel with "port-channel lacp fallback individual" configured. (1021913)
- In certain in-place or out-of-band action update scenarios where TCAM banks allocated to the PBR feature diverge between forwarding chips, PBR counters may stop recording the correct values. (1030398)

- A traffic-policy applied to an interface that never comes up after a linecard hot-swap may cause traffic-policy to complete stop processing configuration updates. (1044368)
- In a modular switch with both Jericho2 and Jericho2C+ line cards, if the Jericho2 line card comes up after the Jericho2C+ linecard and the route scale is high, then the SandL3Unicast agent might unexpectedly restart.

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

- Insertion or power up of a line card can cause some traffic loss on other cards (1072539)

Series Affected: DCS-7800 Series

- The SandACL agent may restart unexpectedly if a Jericho2C+ linecard is replaced with a Jericho2 linecard. (1072612) (1)
- When traffic-policy 'policer interface port-channel distribution proportional' feature is enabled with 2-color counters enabled, at high counter scale, the forwarding agent may restart due to counter resource exhaustion error.

To fix this issue, use 1-color counter engine or make sure the configured traffic-policy counters are within the allowed scale supported. (1089374)

- When packet time stamping is enabled on a switch in Tap Aggregation mode, the switch may report very high PTP offset from Master. Moreover, the PTP hardware-sync Servo State may report 'measuring clock skew' for an extended period of time. This may cause the packet timestamps to deviate significantly from the PTP time, and may cause PTP to lose sync from its master.

Reset the PTP hardware-sync to recover the switch from this state. (1098998)

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

- Applying traffic-policies to a device in tap-agg mode may lead to SandAegis restarting unexpectedly if the TCAM profile does not contain 'src-ip-label', 'dst-ip-label', 'inner-src-ip-label', and 'inner-dst-ip-label' in the 'feature traffic-policy port ipv4' and 'src-ipv6-label', 'dst-ipv6-label', 'inner-src-ip-label', and 'inner-dst-ip-label' in 'feature traffic-policy port ipv6'. (1175391)
- After EOS upgrade, sFlow hardware acceleration may be disrupted on a forwarding chip. Either one of the following workarounds may be used to avoid the issue.
Workaround 1 : Disable sFlow before upgrade event and enable sFlow after forwarding agents are up.
Workaroun 2 : Proceed with the upgrade without disabling sFlow

acceleration. If eventor port gets stuck after upgrade event, disable sFlow and restart the SandFapNi agent for the corresponding linecard. enable sFlow after forwarding agents are up. (1179823)

- When PVLAN is enabled on a trunk interface and that interface has `no switchport mac address learning` configured, a packet may have its source MAC address learned if it was matched by the PVLAN rule. (1196029)
- If "mpls lookup label count 2" configuration is added/removed, then MPLS-over-GRE and MPLS-over-UDP packets that undergo GRE/UDP tunnel decapsulation followed by MPLS label pop might not be forwarded as expected.

Restarting the SandAcl agent will restore normal packet forwarding operation. (1196860)

- TCP MSS clamping may not work on GRE tunnels in non default VRF. (1203845)
- The SandAegis forwarding agent will restart continuously if TCAM profile contains a mix of *-lpm-label and *-label key fields across traffic-policy features. (1212620)
- The forwarding agent may restart unexpectedly following a sequence of speed or speed-group configuration changes. (1213516)
- On Gen4 modular systems, hot swapping the slot 3 linecard with a different SKU model may cause the forwarding agents (e.g. SandFapNi-Linecard*) to continuously restart.

A workaround is to first power down the modular chassis to perform a slot 3 linecard swap if there is an intentional SKU model change. (1215660)

- Fabric link weight calculation on the FAP is incorrect when the number of TX links is less than the number of RX links, this can result in fabric back pressure to the FAP and cause discards on the FAP. (1220497)
- In Tap Aggregation mode, an update to a traffic-policy that includes new match rules with replicate actions may not get fully programmed in hardware.

As a workaround, instead of incrementally updating the traffic-policy, remove it and add back the updated version. (1222857)

- SandAcl agent may restart repeatedly with failed subinterface reflectors. Workaround is to remove all failed subinterface reflectors. (1232750)
- Configuring a traffic-policy with a redirect action to an unresolved or undefined nexthop-group, instead of going to the FIB route, packets hitting that action would drop. In EOS-4.34.2F and later 4.34 releases, the "dropVoqInAcl" drop counter will increment for these dropped packets. In EOS-4.35.0F and later releases, the "dropVoqInNullRoute" drop counter will increment for these dropped packets. (1254076)

- SandTcam agent may restart continually when the operational TCAM profile has "packet ipv4 mpls ipv6 forwarding routed" in either "storm-control subintf" or "sflow subintf" TCAM features.

The workaround is to remove "packet ipv4 mpls ipv6 forwarding routed" from "storm-control subintf" and "sflow subintf" TCAM features. (1260910)

- TCP SYN packets, where TCP MSS is not the first option are dropped if software MSS clamping is enabled and egress interface is VXLAN tunnel.

Workaround : Disable software MSS clamping by removing 'feature tcp-mss-ceiling ip' from the system TCAM profile (1263193)

- In certain in-place or out-of-band action update scenarios where other features free TCAM banks, PBR counters may stop recording the correct values. (1263498)
- In an EVPN VXLAN A/A multi-homing OISM setup, traffic sourced from a singly connected receiver in a VLAN that is present only on one of the multi-homed switches, might result in duplicate delivery of the traffic to a receiver that is reachable via an A/A Ethernet segment.

Workaround: configure the same set of VLANs on both multi-homed VTEPs. (1270442)

- When both "feature traffic-policy port (ipv4|ipv6)" and "feature qos (ip|ipv6|mac)" have "action set-unshared-policer" in the operational TCAM profile, traffic that matches both an interface traffic-policy rule with a policer action and a QoS service-policy class-map with a count action will not be policed. (1274475)
- If both 'feature traffic-policy vlan-interface' and 'feature traffic-policy port' are enabled in the operational TCAM profile, 'show traffic-policy interface' or 'show traffic-policy interface detail' may cause ConfigAgent to unexpectedly restart.

The workaround is to restart the SandAegis agent. (1279590)

- If a port-channel is configured as a mirroring destination, forwarded traffic egressing this port-channel may be corrupted after it is unconfigured as a mirroring destination.

A workaround is to restart the SandDanz agent. (1280794)

- SandL3Unicast agent may restart unexpectedly when a GRE or GUE tunnel is configured as backup path for fast reroute(FRR).

Workaround: Remove any configurations that makes a GRE tunnel or GUE tunnel a backup path for fast reroute. (1298629)

7300X3, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7060X6, DCS-7260X3 and DCS-7300X Series

- On interfaces with speeds 2.5 Gbps or lower, incoming packets may be incorrectly dropped after a link flap.
The problem can be recognized by observing continual repeating messages of the following form in the system logs:
`%FLOWCONTROL-4-PAUSE_FRAME_RECEIVED: Flow Control receive is disabled on Interface Ethernet<NN>.Interface received <NNN> pause frames`
A workaround is to toggle the flowcontrol config on the interface:
`"flowcontrol receive on" followed by "flowcontrol receive off". (137623)`
SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DP-24S, CCS-720DP-48S, CCS-720DT-24S, CCS-720XPM-48TH-6SY, CCS-722XPM-48Y4-F and DCS-7010TX-48
- MAC ACLs and PACLs cannot co-exist on the same interface if QoS policy ACLs exist in the switch (186949)
- The StrataVlanTopo agent may unexpectedly restart after many linecard online insertion and removal operations. (193787)
Series Affected: DCS-7300X Series
- The Strata fabric or linecard agent may restart due to fabric link flaps leading to traffic outage. (225302)
Series Affected: DCS-7300X Series
- DirectFlow flow in table type vfp will not match malformed Arp requests. (233922)
- When traffic is redirected using PBR and when the FIB lookup based on destination IP address indicates ARP miss, the PBR redirected traffic may be rate-policed.

The workaround is to use count action for the ARP-needed class. (238084)
- When the utilization of the MAC address table grows close to the full hardware capacity, the StrataL2 agent may encounter an out of memory (OOM) condition in the system, causing the agent to get restarted. (238115)
Series Affected: DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7260X3 and DCS-7300X 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-7060X Series

- During an SSO failover, Strata Fabric agent may not detect hardware, resulting in a fatal error reboot with increased traffic outage. (250129)
Series Affected: DCS-7300X Series
- In an L3 EVPN setup, traffic that is sourced by a VTEP that is a pure VXLAN router (i.e., it performs no VXLAN bridging and has no hosts attached to it) might incur some loss at the receiving VTEP, where the packets are decapsulated and forwarded.

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

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

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

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

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

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

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

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

- Egress ACL on SVI for BUM traffic will not match ACL rules for packets being sent across fabric interfaces. (284543)
SKUs Affected: DCS-7304 and DCS-7308
- Applying a large configuration containing many interface changes may result in Strata forwarding agent restart, which could cause traffic loss during the reconfiguration. (330840)
- When VXLAN is configured and when a route or ARP points to an ECMP group with both local and remote nexthops, then a packet destined to this ECMP

group may get looped between the linecards and fabric cards resulting in link congestion and packet drops. The work around is to disable the source port based hashing by removing "ingress-interface" in "ip load-sharing trident fields" command. (368073)

Series Affected: DCS-7300X Series

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

Series Affected: 7300X3 and DCS-7050X3 Series

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

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

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

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

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

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

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

- Changing PFC Watchdog configuration while traffic is flowing can result in sub-second traffic outage when PFC Watchdog is in non-disruptive mode and multi-second traffic outage when PFC Watchdog is in disruptive mode. (403790)
- DirectFlow flows with mirroring action may not be reprogrammed after a supervisor switchover or StrataMirror agent restart.

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

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

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

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

disabling the idle timeout.

Series Affected: 7300X3 (411851)

Series Affected: 7300X3 and DCS-7300X Series

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

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

- Ethernet49-54 may not link up with peers over long copper cables.
Workaround is to enable autoneg on both ends of the link with 'speed auto <speed>' command. (421121)

SKUs Affected: DCS-7060SX2-48YC6-F and DCS-7060SX2-48YC6-R

- After a speed change, flows with mirroring actions on affected interfaces may not be mirrored as expected. Strata Slice agent can be restarted to restore expected behavior. (422221)
Series Affected: CCS-720XP, DCS-7050SX3 and DCS-7260CX3 Series
- The forwarding agent may continuously restart if the speed is changed from 100G to 50G when a 50G QSFP transceiver is inserted.
This only affects 2-lane 50G transceivers, and does not affect 100G transceivers operating in 50G mode.
A workaround is to configure the speed before inserting the transceiver.
The forwarding agent restarts can be stopped by temporarily configuring the affected interfaces in 25G mode. (424990)
- With scale policy-map configuration, Strata agent may restart continuously with SIGQUIT. (430654)
- During Leaf-SSU or SSO, interfaces with 100GBASE-SRBD transceivers may experience a single link flap. (440451)
- With PFC pause stream for lossless priorities on egress interface and data traffic destined to the egress interface for lossy and lossless priorities, few seconds of loss may be seen on lossy priority traffic on toggling PFC

globally.

Workaround: Stopping PFC pause frames before toggling PFC globally. (477941)

- Aggregate storm control rate limit is not honored for broadcast and multicast traffic. (490511)

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

- MAC loopback interfaces may not come up after hitless restart of the Strata forwarding agent with "phy link detection aggressive" configured.

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

- MLAG connectivity between the peers might go down after linecard OIR. Flapping MLAG peer-link interfaces will bring the MLAG connectivity back up (500162)

- With "switchport default phone vlan <vlan-id>" configuration, DHCP offer packets tagged with the default phone vlan tag might be sent out untagged on trunk ports (501749)

- It is possible that storm-control configuration will still be applied in hardware after removing the configuration. Workaround is to restart the Strata agent. (504901)

- Enabling per VLAN routing with an L3 interface that only has an IPV6 address will result in no routing of packets received on this L3 interface. Workaround is to flap the interface. (528490)

- Changing the speed of a QSFP or SFP interface from 10G to 1G may cause a Strata forwarding agent restart under some circumstances. This will cause an interruption in traffic on all ports. Workaround is to change from 10G to another speed before changing to 1G. (550269)

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

- When "show platform trident vxlan multihoming groups non-peering" is executed, StrataL3 agent may restart unexpectedly. (561420)

- Fabric interface on a linecard can sometimes end up in a state where it can discard all packets at egress.

When this condition happens, the "show platform trident fabric counters queue detail" command will show that the drops increment while there is no increase in packets or bytes transmitted on the problematic fabric interface. To recover from the problem, it is necessary to restart the linecard forwarding plane agent by running "platform trident <linecard> reset". The recovery process can cause traffic disruption across several ports in the switch. (562860)

SKUs Affected: DCS-7304 and DCS-7308

- Committing security access lists through CVP may restart forwarding plane slice agent. Post restart forwarding plane functionality will continue to work as is.
Workaround is to use "hardware access-list update default-result permit" CLI. (571289)

- IFA packets do not take same path as original traffic when routed using Policy Based Routing (573778)
- IFA packets corresponding to traffic dropped by QoS policy map are not dropped (575953)
- StrataTcam process may restart when we the primary Switchcard is hotswapped. System should be unaffected from the restart of StrataTcam process.

There is no workaround. (584258)

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

- When many VLANs and/or LAGs are configured, port channels may flap during SSU. (592390)
- PIM over GRE interface does not support over 1K multicast routes (605368)
- During system restart or MLAG split brain, StrataL3 forwarding agent may restart unexpectedly. (616991)
Series Affected: CCS-720XP, CCS-750, DCS-7050CX3 and DCS-7050TX3 Series
- SSU and SSO might fails when MSS-G enabled and there are large number of segments (623639)
- After an interface flap, multicast traffic may not forward correctly in data plane.
Workaround is to remove and re-apply multicast configuration. (627318)
Series Affected: CCS-710P and DCS-7010TX Series
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- StrataL3 forwarding agent may restart unexpectedly when many L3 interfaces are removed and then added again. (636201)
- Migrating from static VXLAN to EVPN VXLAN configuration may result in stale remote MAC entries in hardware leading to black holing of traffic destined to those MAC addresses.

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

- If StrataL3 is restarted when there are insufficient hardware resources to add a virtual port for a VXLAN remote VTEP as indicated by the STRATA-3-

VIRTUALPORT_RESOURCE_FULL syslog, StrataL3 agent may not become warm resulting continuous traffic drops.

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

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

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

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

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

- Removing an SFP BASE-T transceiver from Ethernet33 or Ethernet34 may cause the forwarding agent to restart and a brief traffic outage to occur. (715516)
Series Affected: DCS-7050X4 and DCS-7060X4 Series

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

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

Series Affected: CCS-750 Series

- If MLAG configuration is applied when there is congestion, MLAG stays in inactive state and StrataVlanTopo might restart. When the congestion is stopped MLAG becomes active (727116)
- StrataLag agent will continuously restart unexpectedly if there are more LAGs configured (including LAGs with no members) than the hardware can support.

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

Workaround:

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

- An interface configured to be a member of a LAG may fail to be programmed in the hardware.
The CLI 'show port-channel detailed' will show the status of the interface as "Waiting for hardware to program port to RX...".

Workaround:

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

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

Workaround is to flap the VXLAN interface. (748461)

- MAC initialization failure may cause Strata forwarding plane agent to unexpectedly restart during stateful switchover, and this may cause stateful switchover to fail, and reboot the switch. (776821)

Series Affected: 7300X3 Series

- IP PACL configured on interface receiving VXLAN traffic encapsulated to VARP VTEP IP may result in drops.
Workaround is to configure port range rule in the ACL or apply the following command: `platform trident tcam port-acl vfp disabled`. (784384)

- Looping or double delivery of L2 protocol packets received on the peer-link may be observed in an active-active MLAG setup when L2 Protocol Forwarding is configured on the peer-link. (784857)

- When an interface having "switchport trunk private-vlan secondary" configuration is removed from a port-channel we may observe incorrect tagging on packets egressing the interface.
A workaround is to remove and add the "switchport trunk private-vlan secondary" configuration. (786324)

- IPv6 Underlay is not supported. Receiving EVPN routes from IPV6 VTEPs may result in continuous restart of L3 forwarding agent. (794373)

Series Affected: 7388, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series

- When Reducing number of members in ECMP groups while running traffic with DLB enabled can cause some packets drop.

Workaround is to restart the traffic or disable/enable DLB. (797519)

SKUs Affected: 7358X4-SC, DCS-7050DX4-32S, DCS-7050DX4-32S-F, DCS-7050DX4M-32S, DCS-7050DX4M-32S-F, DCS-7050PX4-32S and DCS-7050PX4-32S-F

- After a Leaf Smart Software Upgrade (Leaf SSU), Strata agent might restart unexpectedly, resulting in a fatal error reboot. (809408)
SKUs Affected: DCS-7060DX5-64S
- VLAN translation may stop working on a port when:
 - the port was removed from a LAG
 - that LAG has the same VLAN translation configured as the port

Workaround:

- restart StrataL2 agent (809637)

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

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

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

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

Series Affected: DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series

- Pause Flow control might not work after Autoneg resolution.

Workaround :

Reconfigure Pause Flow control.

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

Series Affected: 7388 Series

SKUs Affected: 7358X4-SC, 7388X5-SC, DCS-7050CX4-24D8, DCS-7050SDX4-48D8, DCS-7050SPX4-48D8 and DCS-7060DX5-64

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

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

- If the system reboots due to a CPU machine check error, the forwarding plane may fail to initialize after the reboot. Another reboot is required to restore the system. (832968)
SKUs Affected: DCS-7010TX-48 and DCS-7010TX-48-DC
- If many continuous MAC moves occur in a short period of time, the device may not be accessible via SSH while the MAC moves continue. (845283)
- When performing configure-replace with traffic-policy which exhausts TCAM, the Strata agent may restart. Workaround would be to lower the scale of the traffic-policy such that TCAM is not exhausted. (848910)
Series Affected: 7300X3, CCS-720DF, CCS-720DP, CCS-720XP, CCS-722XPM, DCS-7010TX and DCS-7050CX3 Series
SKUs Affected: 7358X4-SC, DCS-7304 and DCS-7308
- A chip reset may occur as a result of a fatal parity error reported in the MMU egress buffer block if the system is configured to exceed the hardware replication limit of 1024 replications per packet. (850455)
Series Affected: 7358, 7368, 7388, DCS-7050X4, DCS-7060X4 and DCS-7060X5 Series
- VLAN source mirroring may not work on multi chip systems (854645)
- If multiple devices behind a phone/hub/switch connected to an Arista switch are authenticated via Dot1x MAC Based Authentication and assigned different VLANs, some of these devices may not be able to move to other interfaces that have Dot1x MAC Based Authentication enabled.

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

- 100BASE-FX modules may fail to link up on ports Et1-Et40 after insertion.

Workaround is to run `shutdown` followed by a `no shutdown` on the affected interfaces. (878027)

SKUs Affected: CCS-720DF-48Y

- Configuring and removing static MAC address entries for MAC addresses that have already been learnt may result in high CPU usage and unexpected restart of StrataL2 forwarding agent. (884326)
- When a packet is received with IP Destination address 0.0.0.0 and multicast MAC address 01:00:5e:00:00:00, the packet is either flooded to the VLAN or sent to the multicast router ports in the VLAN.
Workaround is to configure an ACL to drop such packets. (885556)
- ARP packets sent to VRRP MAC don't get copied to control plane if MACsec is enabled on the interface.
Workaround is to disable MACsec on the interface by running "no mac security

profile <profile>". (921378)

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

- Enabling queue-monitor may cause a high CPU load when queue usage is below the high threshold. (949177) (1)

Series Affected: 7368, 7388, DCS-7060CX, DCS-7060CX2, DCS-7060DX4, DCS-7060PX4, DCS-7060SX2 and DCS-7060X6 Series

SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4M-32S, DCS-7050PX4-32S, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060PX5-64E-F and DCS-7260CX3-64

- During a Leaf Smart Software Upgrade (Leaf SSU), Strata agent may restart unexpectedly, resulting in a fatal error reboot. (959348)
- The switch may unexpectedly reboot during Smart System Upgrade (SSU). (959790)

SKUs Affected: CCS-720XP-48TXH-2C-S-F

- A netdevice refcount that is held on an unresponsive Linecard may prevent Strata from shutting down, resulting in a message similar to: 'unregister_netdevice: waiting for et3_49_1 to become free. Usage count = <n>'. The affected supervisor should be rebooted at the operator's earliest convenience. (967727)
- LAG kernel interfaces may flap when Strata agent(s) are restarted. (975146)
- When uRPF exceptions are configured, forwarding agent may restart unexpectedly during switch card switchover or port events. (980046)
- A switch reboot may cause the system to become unresponsive and hang.

A workaround is to physically unplug and replug the power cable to recover. (989010)

Series Affected: CCS-720DF, CCS-720DP and CCS-720DT Series

- If mac/phy loopback is enabled on a port with autoneg configured, the Strata agent can restart unexpectedly.

Workaround is to disable loopback configuration on affected ports. (992820) (1)

SKUs Affected: DCS-7050DX4-32S

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

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

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

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

- In an MLAG configuration, VXLAN flows may see small traffic loss before SSU boot stages begin, along with some traffic loss on MLAG ports. (1004876)
SKUs Affected: CCS-720XP-48TXH-2C-S, CCS-720XP-96ZC2 and CCS-722XPM-48ZY8-F
- After performing a soft reset on Linecard on a MLAG peer device, MACs learned on orphan port might not be get learned.
A workaround is to restart the StrataL2 agent to get out of the issue state. (1018442)
- The following logs may be emitted occasionally.

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

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

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

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

- Specific groups of four adjacent ports share a common hardware resource: Et1-4, Et25-28, Et29-32, Et53-56.

When configured with 40G speed, configuring `no shutdown` on an even numbered port in these ranges may cause the preceding odd numbered port to flap. e.g. Et2 will cause Et1 to flap. (1029987)

SKUs Affected: DCS-7060CX5-56D8

- In an MLAG config, ARP for some dot1x supplicants may not be resolved post SSU. (1039376)
SKUs Affected: CCS-720XP-48TXH-2C-S, CCS-720XP-96ZC2 and CCS-722XPM-48ZY8-F
- When a pair of non-MLAG switches load balance 8K VXLAN Tunnels between them by terminating 4K with each of the switches through a port-channel interface in an STP enabled topology, a hitful reload of one of the switches can result in a traffic loss for around 25 seconds on the reloaded switch as it comes up.
If STP is disabled the traffic loss reduces to around 3 seconds (1054467) (1)
Series Affected: CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, DCS-7010TX, DCS-7050CX3 and DCS-7050SX3 Series

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

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

- If an SFP slot contains an optic or SFP BASE-T transceiver which is down, or if an SFP slot previously contained an optic or SFP BASE-T transceiver which was down prior to its removal, and the command "reload fast-boot" is performed, the interface may fail to link up after reload.

To work around this problem, hitful restart of the forwarding agent is required, leading to the flap of all interfaces. (1089844) (1)

Series Affected: DCS-7260X3 Series

- Port Security shutdown mode reload persistence is not supported on LAG VXLAN interfaces with normal reboot. A workaround is to use reload fast-boot if supported. (1095212)
- SSU may fail and cold reboot instead if there are continuous MAC moves between interfaces with port-security protect configured. (1104157)
- Port Security protect mode reload persistence is not supported on LAG VXLAN interfaces with normal reboot. A workaround is to use reload fast-boot if supported. (1104344)
- If "switchport port-security violation protect" is configured on an interface, and the number of MACs on the port is continuously moving above and below the port-security MAC-address maximum, a large number of ACLs will be created and removed, which may result in the Strata agent exiting with the SOCKET_CLOSE_REMOTE syslog.

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

- The Strata agent may have high CPU utilization for an extended period of time when subject to several continuous MAC moves between port-security protect interfaces. (1111717)
- Doing a "shutdown" configuration on all the interfaces on which a QoS Policy Map is attached to can cause the Strata agent to restart unexpectedly and this might impact traffic. (1167772)

- Configuring MMU queue profiles which reserve more memory than is available may cause a forwarding agent restart with MMU queue profiles not be consistently applied. (1184733)
- When we have a scaled config with Policy Maps applied on sub-interfaces and there is a big subIntf churn(from say, a config replace), we may see the forwarding plane restart. The agent restart recovers this state. (1187873)
- Dot1x MAC Based Authentication will not authenticate the source MAC address of gPTP packets and P2P PTP packets, thus prevent PTP clock from synchronization if 'ptp mode gptp' or 'ptp mode boundary' with 'ptp delay-mechanism p2p' is used. (1196630)
- Security ACL replacement on a Ethernet or LAG port, from an ACL that is programmed in VFP TCAM, to an ACL that is only programmed in IFP TCAM, may cause security access controlling of the port to be performed by another ACL in VFP TCAM.
To workaround the issue, remove and re-apply the ACL applied to the impacted port. (1205638)
- On trunk ports, PTP packets with the native VLAN will not be dropped when the native VLAN is not in the allow VLAN list. (1226897) (1)
- If the Strata slice agent restarts while RACL programming is underway, CPU traffic leaving RACL-applied VLANs may be continuously dropped even after agent restart. Workaround is to remove and then re-apply the RACL on the affected VLANs. (1249719)
- Setting too small a key size limit for configured key fields in TCAM profile, with slice reservation, will result in the Strata agent continuously restarting. (1253028)
- In active-partial MLAG, unicast packets are looped back transiently when the destination MAC is aged out. (1268059)

CCS-750 Series

- When a switchover occurs, the Strata agent may restart unexpectedly (965328)
SKUs Affected: CCS-755-CH and CCS-758-CH
- Control traffic can be impacted when receiving 100% linerate traffic on an uplink port.
This can be avoided by reducing the incoming rate to 99.98% or less.
(1064421)
SKUs Affected: CCS-750-SUP100 and CCS-750-SUP25
- IPV6 multicast packets will be forwarded in the data plane and not reach PTP agent if "ptp forward unicast" is configured with "ptp mode boundary" or "ptp mode boundary one-step". (1086869)

- A hitless slice agent restart may cause some interfaces to blackhole traffic due to a speed mismatch on the fabric interfaces.

A workaround to bring up the affected interfaces is to hitfully restart the Strata slice agent using the CLI "platform trident LinecardX restart" where X is the linecard number associated with the impacted interfaces. (1268542) (1)

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

7300X3 and DCS-7060X Series

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

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

- Packet buffer memory corruption can result in packets getting discarded without visibility in counters.

When this condition occurs, the buffer usage reported by the "show platform trident mmu queue status" command exceeds the total number of buffers reported by the "show platform trident mmu buffers" command.

A workaround is to run the "platform trident reset" command, which results in forwarding agent restart causing the links of the system to flap momentarily. (553299)

- TCAM resource exhaustion can lead to VXLAN forwarding to fail.

Workaround: Reduce TCAM resource utilization to allow VXLAN entries in TCAM by unconfiguring other features that occupy TCAM. (561231) (1)

- StrataL2 agent might assert when MAC deletion and Strata Slice agent restart happens simultaneously.

System will recover automatically. (897491) (1)

- Changing a static MAC to a drop MAC may not take effect in the hardware. The workaround is to remove the existing MAC first, then apply the drop MAC configuration. (938220) (1)

- Enabling queue-monitor may cause a high CPU load when queue usage is below the high threshold. (949177) (1)

Series Affected: 7368, 7388, DCS-7060CX, DCS-7060CX2, DCS-7060DX4, DCS-7060PX4, DCS-7060SX2 and DCS-7060X6 Series

SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4M-32S, DCS-7050PX4-32S, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060PX5-64E-F and DCS-7260CX3-64

- When enabling fast-polling, the StrataCountersLite agent leaks memory continuously at a rate of approximately 10MB per hour. A workaround is to regularly kill the agent. (982326) (1)

- If a large amount of flows are continuously being re-learned, Strata slice agent may restart. (993211) (1)
- Continuous UFT config changes can result in MAC addresses not being learned. Workaround is to restart StrataL2 (1004102) (1)
- Traffic to a multihomed MAC destination may see packet drops during VTEP reachability change when putting spine devices into maintenance mode.

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

- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- VXLAN traffic that are decapsulated and routed with inner payload frame size greater than 9228 bytes will be dropped. (1182137) (1)
- When StrataL3 agent is restarted immediately after disabling uRPF feature, hardware multicast routing may become disabled from that time on. Workaround is to reload the device. (1213671) (1)
- VXLAN BUM traffic may see over 1second traffic loss during SSU. (1260806) (1)
- Some static MAC addresses may not get programmed into hardware after removal and addition of a VLAN, which can lead to flooding of packets to all members of the VLAN.
A workaround is to remove and add back the static MAC address. (1270937) (1)

DCS-7060X2 Series

- TCAM resource exhaustion can lead to VXLAN forwarding to fail.
Workaround: Reduce TCAM resource utilization to allow VXLAN entries in TCAM by unconfiguring other features that occupy TCAM. (561231) (1)
- StrataL2 agent might assert when MAC deletion and Strata Slice agent restart happens simultaneously.

System will recover automatically. (897491) (1)

- Changing a static MAC to a drop MAC may not take effect in the hardware. The workaround is to remove the existing MAC first, then apply the drop MAC configuration. (938220) (1)
- Enabling queue-monitor may cause a high CPU load when queue usage is below the high threshold. (949177) (1)
Series Affected: 7368, 7388, DCS-7060CX, DCS-7060CX2, DCS-7060DX4,

DCS-7060PX4, DCS-7060SX2 and DCS-7060X6 Series

SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4M-32S, DCS-7050PX4-32S, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060PX5-64E-F and DCS-7260CX3-64

- When enabling fast-polling, the StrataCountersLite agent leaks memory continuously at a rate of approximately 10MB per hour. A workaround is to regularly kill the agent. (982326) (1)
- If a large amount of flows are continuously being re-learned, Strata slice agent may restart. (993211) (1)
- Configuring "traffic-loopback source device mac" and "shutdown" on UnconnectedEthernet interfaces does not prevent these interfaces from recirculating traffic. The workaround is to remove the loopback configuration on shutdown interfaces using "no traffic-loopback". (997940)
- Continuous UFT config changes can result in MAC addresses not being learned. Workaround is to restart StrataL2 (1004102) (1)
- Traffic to a multihomed MAC destination may see packet drops during VTEP reachability change when putting spine devices into maintenance mode.

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

- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- VXLAN traffic that are decapsulated and routed with inner payload frame size greater than 9228 bytes will be dropped. (1182137) (1)
- When StrataL3 agent is restarted immediately after disabling uRPF feature, hardware multicast routing may become disabled from that time on. Workaround is to reload the device. (1213671) (1)
- VXLAN BUM traffic may see over 1second traffic loss during SSU. (1260806) (1)
- Some static MAC addresses may not get programmed into hardware after removal and addition of a VLAN, which can lead to flooding of packets to all members of the VLAN.
A workaround is to remove and add back the static MAC address. (1270937) (1)

DCS-7260X3 Series

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

- After a reload, the forwarding agents may fail to initialize and the Linux dmesg will be filled with "link down" log messages.

Workaround is to reload the switch. (526172)

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

- TCAM resource exhaustion can lead to VXLAN forwarding to fail.
Workaround: Reduce TCAM resource utilization to allow VXLAN entries in TCAM by unconfiguring other features that occupy TCAM. (561231) (1)
- StrataL2 agent might assert when MAC deletion and Strata Slice agent restart happens simultaneously.

System will recover automatically. (897491) (1)

- Changing a static MAC to a drop MAC may not take effect in the hardware. The workaround is to remove the existing MAC first, then apply the drop MAC configuration. (938220) (1)
- Enabling queue-monitor may cause a high CPU load when queue usage is below the high threshold. (949177) (1)
Series Affected: 7368, 7388, DCS-7060CX, DCS-7060CX2, DCS-7060DX4, DCS-7060PX4, DCS-7060SX2 and DCS-7060X6 Series
SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4M-32S, DCS-7050PX4-32S, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060PX5-64E-F and DCS-7260CX3-64
- Symptom: On restarting StrataL3 agent, a traffic outage of around 5-10 seconds is observed.

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

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

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

- When enabling fast-polling, the StrataCountersLite agent leaks memory continuously at a rate of approximately 10MB per hour. A workaround is to regularly kill the agent. (982326) (1)
- If a large amount of flows are continuously being re-learned, Strata slice agent may restart. (993211) (1)
- Continuous UFT config changes can result in MAC addresses not being learned. Workaround is to restart StrataL2 (1004102) (1)

- SSU upgrade sequences starting from EOS versions:
 - 4.20.5 and later 4.20.X
 - 4.21.X earlier than the affected release
 to affected releases when using 25G/50G/100G speeds on optical transceivers may result in these interfaces to fail to link up.

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

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

- When using config replace to update from configuration containing unprogrammed CLB flows due to resource exhaustion to a configuration that should not lead to resource exhaustion, CLB flows may remain in unprogrammed state.

The workaround is to toggle the 'load-balance cluster' configuration to allow CLB flows to be programmed. (1028508) (1)

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

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

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

- In a specific scenario of TCAM features configured such that the features are:
 1. Sharing TCAM hardware slices and,
 2. TCAM slices are vertically chained (a.k.a auto expanded) to increase the number of TCAM entries for the features.

Removing these slice shared feature's configuration in such a specific scenario might lead to the TCAM slices not getting freed up. As a result, any subsequent feature configuration requiring TCAM resources might not get programmed in hardware even though the TCAM has sufficient hardware resources available.

The work around for the problem is to restart Strata forwarding plane agent(s). (1118125) (1)

- Executing "show platform trident l3 software pivots" command when the device is running in ALPM mode ("forwarding table partition mode 4") with IP routing disabled may result in unexpected restart of the forwarding plane agent. (1148631) (1)
- VXLAN traffic that are decapsulated and routed with inner payload frame size greater than 9228 bytes will be dropped. (1182137) (1)
- If user tries to start a snapshot while StrataL3 agent is down, and then tries to issue the `clear snapshot counters ecmp` command, StrataCounters agent will crash. (1185321) (1)
- When StrataL3 agent is restarted immediately after disabling uRPF feature, hardware multicast routing may become disabled from that time on. Workaround is to reload the device. (1213671) (1)
- VXLAN BUM traffic may see over 1second traffic loss during SSU. (1260806) (1)
- Some static MAC addresses may not get programmed into hardware after removal and addition of a VLAN, which can lead to flooding of packets to all members of the VLAN.
A workaround is to remove and add back the static MAC address. (1270937) (1)

7368 and DCS-7060X4 Series

- Transmission of 802.3X pause frames is not supported (348758)
- If speed of a member interface of a port channel belonging to a VLAN is changed, then it can cause packet drops on that interface. Workaround is to unconfigure and reconfigure that port channel or kill L3 agent. (380019)
Series Affected: 7368 Series
SKUs Affected: 7368-16C, 7368-16S, 7368-4D and 7368-4P
- MPLS traffic will be dropped if the nexthop MAC address ages out or not known.

The workaround is to set the MAC address aging-time to a larger value or simply turn off MAC aging if possible. (389425) (1)
- Speed change on a port-channel member can cause a flap in BFD and BGP sessions for the entire port-channel (520638)
- StrataMmu agent has high CPU usage when PFC watchdog is configured with a low polling interval, such as 0.010. (636394) (1)
- If an interface is shutdown, static MACs on that interface may not be counted in number of allowed addresses for port-security (862079) (1)
- StrataL2 agent might assert when MAC deletion and Strata Slice agent restart happens simultaneously.

System will recover automatically. (897491) (1)

- Changing a static MAC to a drop MAC may not take effect in the hardware. The workaround is to remove the existing MAC first, then apply the drop MAC configuration. (938220) (1)
- Enabling queue-monitor may cause a high CPU load when queue usage is below the high threshold. (949177) (1)
Series Affected: 7368, 7388, DCS-7060CX, DCS-7060CX2, DCS-7060DX4, DCS-7060PX4, DCS-7060SX2 and DCS-7060X6 Series
SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4M-32S, DCS-7050PX4-32S, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060PX5-64E-F and DCS-7260CX3-64
- When the next hop for the GRE Tunnel's destination is a SVI and when the ARP is resolved but the MAC is not learnt, the packet is not flooded on the underlay VLAN.
The workaround is to avoid MAC entries from aging out by configuring the ARP aging time lower than the MAC aging time. (955421) (1)
- When enabling fast-polling, the StrataCountersLite agent leaks memory continuously at a rate of approximately 10MB per hour. A workaround is to regularly kill the agent. (982326) (1)
- If a large amount of flows are continuously being re-learnt, Strata slice agent may restart. (993211) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- The default switch chip interrupt pacing configuration may result in missed interrupts, causing a temporary pause in sending and receiving packets. This may cause flaps in time sensitive protocols such as BFD.

As a workaround, the pacing delay may be configured manually. Contact support for more details. (1066763)

- In a specific scenario of TCAM features configured such that the features are:
 1. Sharing TCAM hardware slices and,
 2. TCAM slices are vertically chained (a.k.a auto expanded) to increase the number of TCAM entries for the features.Removing these slice shared feature's configuration in such a specific scenario might lead to the TCAM slices not getting freed up. As a result, any subsequent feature configuration requiring TCAM resources might not get programmed in hardware even though the TCAM has sufficient hardware resources

available.

The work around for the problem is to restart Strata forwarding plane agent(s). (1118125) (1)

- IPinIP tunnel packets with inner destination as self-IP are not decapped via decap-group and sent as encapsulated packet to CPU (1220062) (1)
- An upgrade or downgrade may introduce different default modes to the secondary interfaces on a qsfppd port. These interfaces therefore may become errdisabled or operable according to the default modes.

The workaround is to configure forced speeds on these interfaces. (1284935) (1)

SKUs Affected: DCS-7060DX4-32-F, DCS-7060DX4-32C-RV3-F and DCS-7060PX4-32-F

DCS-7060X5 Series

- Forwarding agent "Strata-FixedSystem" may restart when a large number of ports are shutdown or unshutdown using interface range CLI and if interfaces have one or more TCAM based features applied to them like security ACL, policy QoS, storm-control, etc.
Agent recovers after restart. (821616)
- Traffic loss will be observed upon restarting the StrataL3 agent but recovers after some time. (1032223) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- SSU is not supported with L3-Subinterface configuration. (1038654) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- When the switch is in configured cut-through mode and traffic is sent between a pair of 25Gbps interfaces, some TX errors may be generated.
A workaround is to configure the device in `store-and-forward` mode. (1039054)
- If user tries to start a snapshot while StrataL3 agent is down, and then tries to issue the `clear snapshot counters ecmp` command, StrataCounters agent will crash. (1185321) (1)

7388 and DCS-7060X5 Series

- Traffic loss will be observed upon restarting the StrataL3 agent but recovers after some time. (1032223) (1)
Series Affected: 7388 Series

SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R

- SSU is not supported with L3-Subinterface configuration. (1038654) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- If user tries to start a snapshot while StrataL3 agent is down, and then tries to issue the `clear snapshot counters ecmp` command, StrataCounters agent will crash. (1185321) (1)

DCS-7060CX5 and DCS-7060DX5 Series

- Traffic loss will be observed upon restarting the StrataL3 agent but recovers after some time. (1032223) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- SSU is not supported with L3-Subinterface configuration. (1038654) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- If user tries to start a snapshot while StrataL3 agent is down, and then tries to issue the `clear snapshot counters ecmp` command, StrataCounters agent will crash. (1185321) (1)

DCS-7060X6 Series

- "Show hardware resource route" or "Show hardware resource route-v6" when run as a part of multiple CLI session running concurrently, StrataL3Unicast agent might terminate. As a workaround run above mentioned command as part of only one CLI session. (1014726) (1)
Series Affected: DCS-7060X6 Series

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

- Scaled IP Locking configuration may result in an unexpected restart of Strata forwarding agent. (678926)
SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DP-24S, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DT-24S, CCS-720DT-24S-2F and CCS-720DT-24S-F
- On the relevant platforms with a large scale of VXLAN remote ARP entries, when SSU is performed, there could be up to 30 seconds of traffic loss for bidirectional VXLAN routed flows between local host and remote VXLAN hosts. (1002874) (1)

- Port Security and Dot1x authenticated MAC address entries may forward traffic incorrectly after all interfaces in a VLAN are flapped or when MSTP is flapped. Workaround is to remove the MAC addresses and allow them to get relearned and/or get authenticated. (1036032) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- VXLAN traffic may be dropped when changing between IPv4 and IPv6 encapsulation if remote VTEP reachability shares adjacencies for both IPv4 and IPv6 VTEPs (Use of IPv4 over IPv6 nexthops or enabling adjacency sharing).

Workaround: Flap the VXLAN interface. (1068640) (1)

- In a specific scenario of TCAM features configured such that the features are:
 1. Sharing TCAM hardware slices and,
 2. TCAM slices are vertically chained (a.k.a auto expanded) to increase the number of TCAM entries for the features.
 Removing these slice shared feature's configuration in such a specific scenario might lead to the TCAM slices not getting freed up. As a result, any subsequent feature configuration requiring TCAM resources might not get programmed in hardware even though the TCAM has sufficient hardware resources available.

The work around for the problem is to restart Strata forwarding plane agent(s). (1118125) (1)

- Ingress discards are observed when "flowcontrol send" is enabled on interface with speeds 2.5 Gbps or lower. (1249620) (1)

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

- The forwarding slice agent may restart unexpectedly when a SEU is encountered in the BSC_AG_AGE_TABLE during initialization. This may result in brief traffic outage. Switch is expected to start normal operation after the agent restart. (417588) (1)
- Egress mirroring does not encrypt mirrored packets when MACsec is enabled on the destination port. (769793)
 SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F

- Interfaces configured with MACsec may experience outage during Smart System Upgrade (799195)
SKUs Affected: CCS-722XPM-48Y4 and CCS-722XPM-48Y4-F
- SSU may experience a data-plane outage longer than 1 second. (801186)
- Following a watchdog timeout, links may flap prior to the system restarting. (829364)
SKUs Affected: CCS-722XPM-48Y4
- StrataFlowTracker forwarding agent (StftAgent) may restart unexpectedly, when "flow tracker hardware" is turned on and off in a short interval interval. (922896) (1)
- After an SSU reload, the system may encounter an error which prevents SSU from completing. This can lead to extended traffic loss. A cold reboot is required to recover the system. (954577)
SKUs Affected: DCS-7010TX-48, DCS-7010TX-48-DC, DCS-7010TX-48-DC-F, DCS-7010TX-48-DC-R, DCS-7010TX-48-F and DCS-7010TX-48-R
- On the relevant platforms with a large scale of VXLAN remote ARP entries, when SSU is performed, there could be up to 30 seconds of traffic loss for bidirectional VXLAN routed flows between local host and remote VXLAN hosts. (1002874) (1)
- When MACsec is configured on a VXLAN core interface, broadcast, unknown-unicast and multicast packets may get dropped after forwarding plane agent restart.
The workaround is to toggle MACsec on the VXLAN core interface. (1007386)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- Port Security and Dot1x authenticated MAC address entries may forward traffic incorrectly after all interfaces in a VLAN are flapped or when MSTP is flapped. Workaround is to remove the MAC addresses and allow them to get relearned and/or get authenticated. (1036032) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- VXLAN traffic may be dropped when changing between IPv4 and IPv6 encapsulation if remote VTEP reachability shares adjacencies for both IPv4 and IPv6 VTEPs (Use of IPv4 over IPv6 nexthops or enabling adjacency sharing).

Workaround: Flap the VXLAN interface. (1068640) (1)

- When the reachability of the underlay next hop of a VTEP changes from VLAN A to VLAN B, the packets flooded to the VTEP are still sent to a member of VLAN A instead of VLAN B, resulting in traffic loss for the hosts connected behind that VTEP.

Workarounds is to restart the strataL3 forwarding agent or to shut/unshut the underlay SVI interface. (1068901)

- In a specific scenario of TCAM features configured such that the features are:
 1. Sharing TCAM hardware slices and,
 2. TCAM slices are vertically chained (a.k.a auto expanded) to increase the number of TCAM entries for the features.

Removing these slice shared feature's configuration in such a specific scenario might lead to the TCAM slices not getting freed up. As a result, any subsequent feature configuration requiring TCAM resources might not get programmed in hardware even though the TCAM has sufficient hardware resources available.

The work around for the problem is to restart Strata forwarding plane agent(s). (1118125) (1)

- In a traffic policy deployment, the mirror action for a rule may fail even if the target mirroring session is up. As a result, traffic matching the rule will not be mirrored. (1139348) (1)
- SSU is only supported from 4.30.1 and onwards to newer releases. (1183988)
SKUs Affected: DCS-7010TX-48-F and DCS-7010TX-48-R
- Ingress discards are observed when "flowcontrol send" is enabled on interface with speeds 2.5 Gbps or lower. (1249620) (1)

CCS-720DF, CCS-720DP and CCS-720XP Series

- The forwarding slice agent may restart unexpectedly when a SEU is encountered in the BSC_AG_AGE_TABLE during initialization. This may result in brief traffic outage. Switch is expected to start normal operation after the agent restart. (417588) (1)
- Setup with only Storm-control and scaled IP Locking configuration may cause slice agent to restart continuously. (599935)
- StrataFlowTracker forwarding agent (StftAgent) may restart unexpectedly, when "flow tracker hardware" is turned on and off in a short interval interval. (922896) (1)
- A hitful restart of the PhyIsland agent in the presence of ingressing traffic on ports Ethernet 1-40 may cause an unexpected restart of the Strata

forwarding agent. This leads to a flap of all ports and a brief traffic outage. (951318)

Series Affected: CCS-720DF Series

- On the relevant platforms with a large scale of VXLAN remote ARP entries, when SSU is performed, there could be up to 30 seconds of traffic loss for bidirectional VXLAN routed flows between local host and remote VXLAN hosts. (1002874) (1)
- Exact match table lookup is not supported with IPV6 underlay multicast enabled. As a result we see duplicate packets when we enabled platform trident forwarding-table partition flexible exact-match configuration. (1018543) (1)
- When IGMP snooping is enabled on the VXLAN VLAN and when the core interface is switched from routed port to SVI or vice versa, VTEPs behind that interface may not receive the multicast data traffic.
Workaround is to clear the IP mroutes by issuing the "clear ip mroute *" CLI command. (1057345) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- VXLAN traffic may be dropped when changing between IPv4 and IPv6 encapsulation if remote VTEP reachability shares adjacencies for both IPv4 and IPv6 VTEPs (Use of IPv4 over IPv6 nexthops or enabling adjacency sharing).

Workaround: Flap the VXLAN interface. (1068640) (1)

- In a specific scenario of TCAM features configured such that the features are:
 1. Sharing TCAM hardware slices and,
 2. TCAM slices are vertically chained (a.k.a auto expanded) to increase the number of TCAM entries for the features.Removing these slice shared feature's configuration in such a specific scenario might lead to the TCAM slices not getting freed up. As a result, any subsequent feature configuration requiring TCAM resources might not get programmed in hardware even though the TCAM has sufficient hardware resources available.

The work around for the problem is to restart Strata forwarding plane agent(s). (1118125) (1)

- In a traffic policy deployment, the mirror action for a rule may fail even if the target mirroring session is up. As a result, traffic matching the rule will not be mirrored. (1139348) (1)
- When multicast routing enabled on a VLAN with VXLAN aware IGMP Snooping configured, VTEPs of that VLAN may experience traffic loss.
Workaround: disable L3 multicast routing, or add an orphan port to the VLAN. (1144821) (1)
- An SEU in the EXACT_MATCH_ECC table may not be corrected. This can cause traffic to be misforwarded. (1212381) (1)
- Ingress discards are observed when "flowcontrol send" is enabled on interface with speeds 2.5 Gbps or lower. (1249620) (1)
- Enabling Traffic Trap to CPU on an inactive multi-lane interface may cause the `StrataMirror` agent to unexpectedly restart continuously until the configuration is removed. This condition may occur if traffic trap is enabled on a subordinate interface of a multilane module, followed by a speed or mode change that causes these interfaces to disappear. Similarly, this may also occur if the traffic trap mode is forcibly configured on an inactive interface using either a configuration replace command, or by exposing these interfaces with the `service interface inactive expose` command.

To work around this issue, use the `service interface inactive expose` to expose the inactive interfaces, and then remove traffic trap mode on them with the `no forwarding action trap` command. (1261276) (1)

SKUs Affected: CCS-720XP-48TXH-2C-S-F, CCS-720XP-48ZC2-F, CCS-720XP-96ZC2-F and CCS-720XP-96ZC2-M-S-F

- With VXLAN-aware IGMP snooping enabled, traffic may not egress to snooped VTEPs.

As a workaround, enable "ip multicast-routing" (1292834)

CCS-750 Series

- The forwarding slice agent may restart unexpectedly when a SEU is encountered in the BSC_AG_AGE_TABLE during initialization. This may result in brief traffic outage. Switch is expected to start normal operation after the agent restart. (417588) (1)
- When report flooding is disabled in IGMP Snooping configuration, report packets may still be flooded on port-channel member interfaces.

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

- Linecard slice agents may restart after supervisor switchover. The system recovers normally. There is no workaround. (623381)

- SSO with L3 VXLAN may lead to extended outage greater than 1 second. (637602)
- Inserting secondary switch card may result in unexpected restart of StrataL3 forwarding agent. (639631)
- When packets are flooded on a VLAN and the ingress interface is a downlink, the ingress interface will have lower egress bandwidth for other traffic flows even though the flooded traffic does not egress the ingress interface. (649424)
- Hotswapping the active supervisor, and then issuing "redundancy manual switchover" will cause stateful switchover to fail. (718863)
Series Affected: CCS-750 Series
- Performing SSO may result in extended traffic loss. (748533)
Series Affected: CCS-750 Series
- Feature agent StrataL2 would continuously restart when the number of unique multicast replication list goes higher than 16k. Workaround is to reduce the multicast scale (784949)
- In some rare cases, after a stateful switchover, the switch could start discarding all packets. To recover from the problem state, "platform trident Switchcard1/0 reset" and "platform trident Switchcard2/0 reset" must be executed." (787513)
- 'vxlan multicast headend-replication' command is not supported on this platform. If this command is present in the startup-config StrataL3 process may restart continuously.
To workaround this issue, remove 'vxlan multicast headend-replication' command from startup-config and replace the config using 'config replace flash:startup-config' command or reboot the switch. (798374)
- When the core VLAN of a VXLAN VLAN has a LAG and needs to flood edge-to-core traffic to all of its ports, the addition of another non-LAG port to that core VLAN results in traffic not egressing the LAG. The removal of the LAG and the addition of another port to the core VLAN results in traffic not egressing this new port.

Restart the slice agent to recover from this condition. (828814)
- When IGMP Snooping is enabled, SwitchCard or LineCard restarts or resets may result in unexpected restart of StrataL3 agent. (830962)
- Known unicast traffic may not flood to all ports of a core VLAN when the underlay MAC of the nexthop to the remote VTEP hasn't been learnt.

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

- Upon events causing restart of the forwarding plane agents, L3 traffic that needs to egress VXLAN VLANs might not do so for a few minutes.

The issue resolves itself in a few minutes. To hasten it, run clear the IP routes for the VRF in question. (840366)

- StrataL3 agent restart may result in extended traffic loss of more than 1 seconds on VXLAN VLANs. (855597)
- StrataL3 forwarding agent may restart unexpectedly and recover when StrataCes agent restarts continuously. The StrataL3 agent restarts once and recovers. (870320)
- StrataFlowTracker forwarding agent (StftAgent) may restart unexpectedly, when "flow tracker hardware" is turned on and off in a short interval interval. (922896) (1)
- On restart of all Strata-SwitchcardCes agents, IPv4 multicast traffic might be black-holed. (981827)
- Performing a physical supervisor pull with SSO enabled, with an ECMP group containing interfaces from both supervisors and passing L3 multicast traffic, may result in extended traffic outage for the switchover. (1002219)
SKUs Affected: CCS-755-CH and CCS-758-CH
- If "switchport trunk group peerlink" is configured on an interface and the linecard for that interface is removed, Strata-SwitchcardCes forwarding plane agent will unexpectedly restart. (1021064)
SKUs Affected: CCS-758-X3-SC
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- VXLAN traffic may be dropped when changing between IPv4 and IPv6 encapsulation if remote VTEP reachability shares adjacencies for both IPv4 and IPv6 VTEPs (Use of IPv4 over IPv6 nexthops or enabling adjacency sharing).

Workaround: Flap the VXLAN interface. (1068640) (1)

- when next hop MAC is not learnt for a core VLAN and core VLAN has LAG and ethernet ports, the VXLAN BUM encapsulated packet is not getting flooded to the LAG but only to ethernet port of core VLAN. (1068904)
- An SEU in the MPLS_ENTRY, EGR_VLAN_XLATE_1, EGR_VLAN_XLATE_2, VLAN_XLATE_1, or VLAN_XLATE_2 tables may cause the StrataL2 or StrataL3 forwarding agent

to continually restart unexpectedly. This only occurs after the affected agent is restarted. This can lead to traffic being misforwarded.

A workaround is to reset the affected switchcard after the switch has entered the continual restarts of the StrataL2 or StrataL3 forwarding agents (1111488)

- In a specific scenario of TCAM features configured such that the features are:
 1. Sharing TCAM hardware slices and,
 2. TCAM slices are vertically chained (a.k.a auto expanded) to increase the number of TCAM entries for the features.

Removing these slice shared feature's configuration in such a specific scenario might lead to the TCAM slices not getting freed up. As a result, any subsequent feature configuration requiring TCAM resources might not get programmed in hardware even though the TCAM has sufficient hardware resources available.

The work around for the problem is to restart Strata forwarding plane agent(s). (1118125) (1)

- An SEU in the MPLS_ENTRY, EGR_VLAN_XLATE_1, EGR_VLAN_XLATE_2, or VLAN_XLATE_1 tables that is triggered by an SSU, SSO, or SSS may cause the system to restart. (1130835)
- In some cases, with BGP being used for route learning, when the active Supervisor is pulled from chassis, with ECMP groups(s) having members on both active and standby Supervisor passing traffic, there may be higher than expected traffic loss during the Stateful Switchover (SSO) procedure until the system converges its state to having only one Supervisor. The maximum expected traffic loss is 5 seconds. (1182524)

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

- Configured "ip verify source" on any interface will result in the Strata forwarding agent continuously restarting and lead to traffic blackholing.

To workaround this, remove the "ip verify source" from all interface configurations. (1190135)

- DHCP packets on a trunk port from a MAC-Based VLAN Assignment phone have their VLAN tag removed. (1195390)
- An SEU in the EXACT_MATCH_ECC table may not be corrected. This can cause traffic to be misforwarded. (1212381) (1)
- A hitless slice agent restart may cause some interfaces to blackhole traffic due to a speed mismatch on the fabric interfaces.

A workaround to bring up the affected interfaces is to hitfully restart the Strata slice agent using the CLI "platform trident LinecardX restart" where

X is the linecard number associated with the impacted interfaces. (1268542)
(1)

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

DCS-7050X3 Series

- Changing interface speed can cause all subsequent link up and link down notifications to be delayed by approximately one second on all ports. The workaround is to restart the forwarding agent after the speed change using the CLI command "platform trident reset". This causes all ports to flap and a brief traffic outage to occur. (810073)
- When enabling fast-polling, the StrataCountersLite agent leaks memory continuously at a rate of approximately 10MB per hour. A workaround is to regularly kill the agent. (982326) (1)
- Exact match table lookup is not supported with IPV6 underlay multicast enabled. As a result we see duplicate packets when we enabled platform trident forwarding-table partition flexible exact-match configuration. (1018543) (1)
- When a pair of non-MLAG switches load balance 8K VXLAN Tunnels between them by terminating 4K with each of the switches through a port-channel interface in an STP enabled topology, a hitful reload of one of the switches can result in a traffic loss for around 25 seconds on the reloaded switch as it comes up.
If STP is disabled the traffic loss reduces to around 3 seconds (1054467) (1)
Series Affected: CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, DCS-7010TX, DCS-7050CX3 and DCS-7050SX3 Series
- When IGMP snooping is enabled on the VXLAN VLAN and when the core interface is switched from routed port to SVI or vice versa, VTEPs behind that interface may not receive the multicast data traffic.
Workaround is to clear the IP mroutes by issuing the "clear ip mroute *" CLI command. (1057345) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- VXLAN traffic may be dropped when changing between IPv4 and IPv6 encapsulation if remote VTEP reachability shares adjacencies for both IPv4 and IPv6 VTEPs (Use of IPv4 over IPv6 nexthops or enabling adjacency sharing).

Workaround: Flap the VXLAN interface. (1068640) (1)

- In a specific scenario of TCAM features configured such that the features are:
 1. Sharing TCAM hardware slices and,
 2. TCAM slices are vertically chained (a.k.a auto expanded) to increase the number of TCAM entries for the features.

Removing these slice shared feature's configuration in such a specific scenario might lead to the TCAM slices not getting freed up. As a result, any subsequent feature configuration requiring TCAM resources might not get programmed in hardware even though the TCAM has sufficient hardware resources available.

The work around for the problem is to restart Strata forwarding plane agent(s). (1118125) (1)

- In a traffic policy deployment, the mirror action for a rule may fail even if the target mirroring session is up. As a result, traffic matching the rule will not be mirrored. (1139348) (1)
- When multicast routing enabled on a VLAN with VXLAN aware IGMP Snooping configured, VTEPs of that VLAN may experience traffic loss.
Workaround: disable L3 multicast routing, or add an orphan port to the VLAN. (1144821) (1)
- An SEU in the EXACT_MATCH_ECC table may not be corrected. This can cause traffic to be misforwarded. (1212381) (1)

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

- Interface speed change or a transceiver insertion may cause one of the following:
 - One or more interface to discard all packets.
To recover: restart the forwarding agent using the CLI command "platform trident reset". This will cause all interfaces to flap and a brief traffic outage to occur.
 - Forwarding agent to restart. (527484)
- IFA packets corresponding to traffic dropped by ingress security ACLs are not dropped (561922)
Series Affected: DCS-7050CX3 and DCS-7050SX3 Series

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

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

- Momentary MAC Rx local faults and remote faults may not be reflected in MAC Rx fault counters and may not cause the interface state to flap. (880612)
SKUs Affected: 7300X3-32C-LC, 7300X3-48TC4-LC and 7300X3-48YC4-LC

- SSU from a release < EOS-4.24.0F may fail with a fatal error if interfaces on a single QSFP port are configured with speeds 50G and 10G, 25G or default. Such interfaces are identifiable by the speed-misconfigured errdisabled state. The workaround is to configure matching speeds on all interfaces in a QSFP port. (891375)
- When enabling fast-polling, the StrataCountersLite agent leaks memory continuously at a rate of approximately 10MB per hour. A workaround is to regularly kill the agent. (982326) (1)
- Exact match table lookup is not supported with IPV6 underlay multicast enabled. As a result we see duplicate packets when we enabled platform trident forwarding-table partition flexible exact-match configuration. (1018543) (1)
- When IGMP snooping is enabled on the VXLAN VLAN and when the core interface is switched from routed port to SVI or vice versa, VTEPs behind that interface may not receive the multicast data traffic.
Workaround is to clear the IP mroutes by issuing the "clear ip mroute *" CLI command. (1057345) (1)
- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- In a specific scenario of TCAM features configured such that the features are:
 1. Sharing TCAM hardware slices and,
 2. TCAM slices are vertically chained (a.k.a auto expanded) to increase the number of TCAM entries for the features.
 Removing these slice shared feature's configuration in such a specific scenario might lead to the TCAM slices not getting freed up. As a result, any subsequent feature configuration requiring TCAM resources might not get programmed in hardware even though the TCAM has sufficient hardware resources available.

The work around for the problem is to restart Strata forwarding plane agent(s). (1118125) (1)
- In a traffic policy deployment, the mirror action for a rule may fail even if the target mirroring session is up. As a result, traffic matching the rule will not be mirrored. (1139348) (1)
- When multicast routing enabled on a VLAN with VXLAN aware IGMP Snooping configured, VTEPs of that VLAN may experience traffic loss.
Workaround: disable L3 multicast routing, or add an orphan port to the VLAN. (1144821) (1)

- An SEU in the EXACT_MATCH_ECC table may not be corrected. This can cause traffic to be misforwarded. (1212381) (1)

7358 and DCS-7050X4 Series

- When more than 500 access list rules are configured and applied to range of interfaces, slice agent may encounter SIGQUIT and restart when shut and no shut on the range of interfaces.

Workaround is to reduce number of access list rules. (679292)

Series Affected: 7368 Series

- When security ACL is being updated, port blocking rules are applied which drop the packets to avoid mis forwarding. This will also block traffic from reaching the control plane. This will cause the protocol like BGP, BFD etc to not work temporarily which requires control plane for processing. As a workaround, use CLI `hardware access-list update default-result permit` to avoid port blocking rules taking affect. With this configuration packets are not dropped due to port blocking rules when ACL is being updated. (703753)

- StrataL3Unicast agent restart may result in the flapping of Layer3 sub interfaces. (732310)

- When a route adjacency transitions from a non-ECMP configuration to to an ECMP configuration, some interfaces might see a traffic loss of 300-500ms. (743031)

- After parity error occurs in LTS_TCAM memories, traffic forwarding may not work as expected.

Workaround is to restart the Strata-FixedSystem agent after the parity error. (756038)

- Configuring "ip hardware fib next-hop sharing" while VXLAN routes are present may lead to continuous restart of StrataL3Unicast forwarding agent.

The workaround is to disable next-hop sharing. (796010)

- Enabling queue-monitor may cause a high CPU load when queue usage is below the high threshold. (949177) (1)

Series Affected: 7368, 7388, DCS-7060CX, DCS-7060CX2, DCS-7060DX4, DCS-7060PX4, DCS-7060SX2 and DCS-7060X6 Series

SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4M-32S, DCS-7050PX4-32S, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060PX5-64E-F and DCS-7260CX3-64

- If mac/phy loopback is enabled on a port with autoneg configured, the Strata agent can restart unexpectedly.

Workaround is to disable loopback configuration on affected ports. (992820)

(1)

SKUs Affected: DCS-7050DX4-32S

- When the device is configured with 128K MAC scale and with "forwarding vxlan disabled", a subset of routes in LPM hardware may drop traffic instead of forwarding it.

Workaround is to change the forwarding scale configuration or to enable VXLAN with 128K mac mode. (1055002) (1)

- Traffic Policy rules configured with "match fragment" will miss matching the tail fragments.

Use security ACLs to avoid this. (1135594) (1)

SKUs Affected: DCS-7050CX4-24D8-F, DCS-7050CX4-24D8-R, DCS-7050CX4M-48D8-F, DCS-7050CX4M-48D8-R, DCS-7050DX4-32S-F, DCS-7050DX4M-32S-F, DCS-7050PX4-32S-F, DCS-7050SDX4-48D8-F, DCS-7050SDX4-48D8-R, DCS-7050SPX4-48D8-F and DCS-7050SPX4-48D8-R

- A Traffic Policy configured with a match fragment rule will not output a SYSLOG indicating that tail fragments will not be matched. (1257531) (1)

SKUs Affected: DCS-7050CX4-24D8-F, DCS-7050CX4-24D8-R, DCS-7050CX4M-48D8-F, DCS-7050CX4M-48D8-R, DCS-7050DX4-32S-F, DCS-7050DX4M-32S-F, DCS-7050PX4-32S-F, DCS-7050SDX4-48D8-F, DCS-7050SDX4-48D8-R, DCS-7050SPX4-48D8-F and DCS-7050SPX4-48D8-R

DCS-7050X4 Series

- Repeated addition followed by removal or updation of ingress port access lists may result in forwarding plane slice agent restarting. The forwarding plane slice agent post restart will continue to work as expected.

Currently there is no workaround for this issue. (811140)

- When the device is configured with 128K MAC scale and with "forwarding vxlan disabled", a subset of routes in LPM hardware may drop traffic instead of forwarding it.

Workaround is to change the forwarding scale configuration or to enable VXLAN with 128K mac mode. (1055002) (1)

- Traffic Policy rules configured with "match fragment" will miss matching the tail fragments.

Use security ACLs to avoid this. (1135594) (1)

SKUs Affected: DCS-7050CX4-24D8-F, DCS-7050CX4-24D8-R, DCS-7050CX4M-48D8-F, DCS-7050CX4M-48D8-R, DCS-7050DX4-32S-F, DCS-7050DX4M-32S-F, DCS-7050PX4-32S-F, DCS-7050SDX4-48D8-F, DCS-7050SDX4-48D8-R, DCS-7050SPX4-48D8-F and DCS-7050SPX4-48D8-R

- A Traffic Policy configured with a match fragment rule will not output a SYSLOG indicating that tail fragments will not be matched. (1257531) (1)

SKUs Affected: DCS-7050CX4-24D8-F, DCS-7050CX4-24D8-R, DCS-7050CX4M-48D8-F, DCS-7050CX4M-48D8-R, DCS-7050DX4-32S-F, DCS-7050DX4M-32S-F,

DCS-7050PX4-32S-F, DCS-7050SDX4-48D8-F, DCS-7050SDX4-48D8-R,
DCS-7050SPX4-48D8-F and DCS-7050SPX4-48D8-R

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

7388 and DCS-7060X5 Series

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

The workaround is to set the MAC address aging-time to a larger value or simply turn off MAC aging if possible. (389425) (1)

- StrataMnu agent has high CPU usage when PFC watchdog is configured with a low polling interval, such as 0.010. (636394) (1)
- When StrataL2 agent restarts, some static MAC addresses pointing to port-channel may not get programmed in hardware table resulting in flooding. Workaround is to removing and add the static MAC address configuration. (787583)
- With the scale of VXLAN VLANs and flood VTEPs, churn in flood VTEP list may result in exhaustion of replication group used for flooding. Workaround is to flap the VXLAN interface. (824256)
- If an interface is shutdown, static MACs on that interface may not be counted in number of allowed addresses for port-security (862079) (1)
- StrataL2 agent might assert when MAC deletion and Strata Slice agent restart happens simultaneously.

System will recover automatically. (897491) (1)

- VXLAN VTEP(s) and HER nexthop entries may not be cleared/deleted from the system ("show vxlan vtep command") , even after corresponding EVPN routes are deleted, when the system fails to allocate new multicast groups due to scaled config. Workaround is to flap the VXLAN interface. (908722)
- Configuring auto-negotiation on an interface (e.g., "speed auto 400g-8") may cause an unexpected restart of the forwarding agent. This leads to a flap of all interfaces and a brief traffic outage. (928889)
SKUs Affected: DCS-7050SDX4-48D8, DCS-7050SPX4-48D8, DCS-7060DX5-32, DCS-7060DX5-64, DCS-7060DX5-64E and DCS-7060DX5-64S
- Enabling queue-monitor may cause a high CPU load when queue usage is below the high threshold. (949177) (1)
Series Affected: 7368, 7388, DCS-7060CX, DCS-7060CX2, DCS-7060DX4, DCS-7060PX4, DCS-7060SX2 and DCS-7060X6 Series

SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4M-32S, DCS-7050PX4-32S, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060PX5-64E-F and DCS-7260CX3-64

- When the next hop for the GRE Tunnel's destination is a SVI and when the ARP is resolved but the MAC is not learnt, the packet is not flooded on the underlay VLAN.
The workaround is to avoid MAC entries from aging out by configuring the ARP aging time lower than the MAC aging time. (955421) (1)
- When "platform trident forwarding-table partition flexible ip-multicast <>" is configured and when the unicast route scale is high enough to fill the entire unicast forwarding table partition, an unexpected restart of StrataL3 agent may be seen. Work around is to reduce unicast route scale. (965291)
- With `storm-control <type> cpu ...`, policed CPU-bound packets may end up on the `Other` CPU queue.

The workaround is to configure both `storm-control <type>` and `storm-control <type> cpu`. (967978) (1)

- When using config replace to update from configuration containing unprogrammed CLB flows due to resource exhaustion to a configuration that should not lead to resource exhaustion, CLB flows may remain in unprogrammed state.

The workaround is to toggle the 'load-balance cluster' configuration to allow CLB flows to be programmed. (1028508) (1)

- When many storm control rules are applied, the Strata agent may unexpectedly restart after an interface speed change. (1029139) (1)
- Traffic loss will be observed upon restarting the StrataL3 agent but recovers after some time. (1032223) (1)

Series Affected: 7388 Series

SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R

- SSU is not supported with L3-Subinterface configuration. (1038654) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R

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

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

- Policy-map action used to overwrite VLAN header fields does not apply for VXLAN Decapped packet. (1064504) (1)

Series Affected: 7388 and DCS-7060X Series

- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1066274) (1)
- CLB may take significant time to re converge upon remote link failure when configured without the ExactMatch table. (1143879) (1)
- Unexpected traffic loss may be observed if StrataL3 agent restarts when Bridged VXLAN CLB flows are programmed without the Exact-Match table.

A workaround to reduce traffic outage is to clear existing flows to be relearned. (1150563) (1)

- Hardware Multicast routing and features that use Exact-Match are not supported together. (1164266)
- The issue occurs in the following cases:
 1. L3 Multicast
 2. Vxlan HER

In the case of L3 Multicast, if there is an L2 receiver and an L3 receiver, and both copies need to go out on the same egress port, only the L3 copy is transmitted - the L2 copy is not sent.

In the case of Vxlan HER, there are situations where the chip must send out two copies on the same egress port - one naked copy for the overlay VLAN and one encapsulated copy in the underlay VLAN. In such cases, only the encapsulated copy is transmitted in the egress packet - the naked copy is not sent in the overlay VLAN. This situation typically occurs during Mlag Core Port Down scenarios, where both copies need to be sent over the peer link. Another scenario involves the egress port being a member of both the overlay VLAN and the core VLAN.

In the case of Mlag Core Port Down scenario, the switch should resolve the issue after the core port comes back up. (1178900) (1)

- Starting an AGM snapshot when there are no ECMP groups configured may result in continuous restart of StrataCounters forwarding agent.

Workaround is to stop the AGM snapshot using "clear snapshot counters ecmp". (1188343)

- Interfaces may fail to link up or may experience link flaps. (1196465)
SKUs Affected: DCS-7060DX5-64S
- Configuring speed-group configuration on a set of affect interfaces can result in VXLAN decapsulation traffic egressing out the affected interfaces without a VLAN tag.

Workaround: Reconfigure the interface to 'switchport mode access` and back to 'switchport mode trunk' to recover. (1206433) (1)

- IPinIP tunnel packets with inner destination as self-IP are not decapped via decap-group and sent as encapsulated packet to CPU (1220062) (1)
- Link flaps on interfaces with auto-negotiation configuration present may cause the forwarding agent to restart and all links to flap. This would primarily be an issue with links which are prone to flapping. (1249552) (1)
- VXLAN BUM traffic may see over 1second traffic loss during SSU. (1260806) (1)
- When an ACL is applied via configure-replace or an interface-range command to a large number(40+) of SVIs or sub-interfaces, or when many ACLs are applied across a large number of interfaces via configure-replace, the Strata-FixedSystem forwarding agent may restart due to a SIGQUIT signal.

Workarounds:

Option 1: Configure "hardware access-list update default-result permit" command. This command will permit traffic during ACL updates instead of denying all traffic.

Option 2: Reduce the number of ACL applications per configuration session. (1262317) (1)

- On certain devices, the 'show forwarding destination [detail]' command may continuously fail after an initial failure until the device is rebooted. (1263648) (1)

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

DCS-7060X6 Series

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

The workaround is to set the MAC address aging-time to a larger value or simply turn off MAC aging if possible. (389425) (1)

- Enabling queue-monitor may cause a high CPU load when queue usage is below the high threshold. (949177) (1)

Series Affected: 7368, 7388, DCS-7060CX, DCS-7060CX2, DCS-7060DX4, DCS-7060PX4, DCS-7060SX2 and DCS-7060X6 Series

SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4M-32S, DCS-7050PX4-32S, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060PX5-64E-F and DCS-7260CX3-64

- With `storm-control <type> cpu ...`, policed CPU-bound packets may end up on the `Other` CPU queue.

The workaround is to configure both `storm-control <type>` and `storm-control <type> cpu`. (967978) (1)

- With PFC Watchdog feature enabled, sometimes a port-channel link flap in the presence of continuous PFC pause storm can cause the Strata agent to restart. (1005078)
- "Show hardware resource route" or "Show hardware resource route-v6" when run as a part of multiple CLI session running concurrently, StrataL3Unicast agent might terminate. As a workaround run above mentioned command as part of only one CLI session. (1014726) (1)

Series Affected: DCS-7060X6 Series

- In a nexthop-group configured setup with close to limit ECMP groups configured, network churn could cause ECMP group exhaustion. Workaround is to restart StrataL3Unicast agent; the restart will be hitful (1026122)
- When using config replace to update from configuration containing unprogrammed CLB flows due to resource exhaustion to a configuration that should not lead to resource exhaustion, CLB flows may remain in unprogrammed state.

The workaround is to toggle the 'load-balance cluster' configuration to allow CLB flows to be programmed. (1028508) (1)

- When many storm control rules are applied, the Strata agent may unexpectedly restart after an interface speed change. (1029139) (1)
- The coprocessor hosting bcmaccel may get stuck in reset due to an incomplete reset sequence. Once in this state the only way to exit the state is to reset the switch. If the behavior occurs during SSU the upgrade will revert to a normal hitful upgrade. To work around the issue a switch reset may be initiated by terminating the slice agent associated with the switch. (1056442)

SKUs Affected: DCS-7060X6-64PE-F

- Policy-map action used to overwrite VLAN header fields does not apply for VXLAN Decapped packet. (1064504) (1)

Series Affected: 7388 and DCS-7060X Series

- A SEU in EGR_IP_TUNNEL might result in an uncorrectable parity error and cause dropped traffic for IP, IPv6 and MPLS tunnels.

Restarting the forwarding slice agent will correct the parity error.
(1066137)

- The Strata agent may unexpectedly restart when config replace is performed or a large number of interfaces are flapped. (1108280)
- If dot1x MAC-based authentication is configured on an interface, it may take around 10 to 15 seconds for an interface to become active. (1122571)

Series Affected: DCS-7060X6 Series

- CLB may take significant time to re converge upon remote link failure when configured without the ExactMatch table. (1143879) (1)
- Unexpected traffic loss may be observed if StrataL3 agent restarts when Bridged VXLAN CLB flows are programmed without the Exact-Match table.

A workaround to reduce traffic outage is to clear existing flows to be relearned. (1150563) (1)

- The issue occurs in the following cases:
 1. L3 Multicast
 2. Vxlan HER

In the case of L3 Multicast, if there is an L2 receiver and an L3 receiver, and both copies need to go out on the same egress port, only the L3 copy is transmitted – the L2 copy is not sent.

In the case of Vxlan HER, there are situations where the chip must send out two copies on the same egress port – one naked copy for the overlay VLAN and one encapsulated copy in the underlay VLAN. In such cases, only the encapsulated copy is transmitted in the egress packet – the naked copy is not sent in the overlay VLAN. This situation typically occurs during Mlag Core Port Down scenarios, where both copies need to be sent over the peer link. Another scenario involves the egress port being a member of both the overlay VLAN and the core VLAN.

In the case of Mlag Core Port Down scenario, the switch should resolve the issue after the core port comes back up. (1178900) (1)

- A multiple bit SEU in the EGR_IP_TUNNEL_IPV6 or EGR_IP_TUNNEL_MPLS tables may not be corrected, leading to traffic misforwarding or drops.
Restarting the forwarding slice agent will correct the parity error.
(1189048)
- Configuring speed-group configuration on a set of affect interfaces can result in VXLAN decapsulation traffic egressing out the affected interfaces without a VLAN tag.

Workaround: Reconfigure the interface to 'switchport mode access' and back to 'switchport mode trunk' to recover. (1206433) (1)

- With VTEP(s) reachable over Port-Channel interfaces, configuration replace from existing VXLAN configuration to non-VXLAN configuration or vice-versa may result in unexpected restart of StrataL3 forwarding agent. (1226650)
- A coprocessor may get stuck in reset due to an incomplete reset sequence. Once in this state the only way to exit the state is to reset the switch. (1226902)
- Restart of the Srv6 agent could intermittently trigger an unexpected restart of the StrataL3Unicast agent. (1229438)
- Link flaps on interfaces with auto-negotiation configuration present may cause the forwarding agent to restart and all links to flap. This would primarily be an issue with links which are prone to flapping. (1249552) (1)
- A coprocessor may get stuck in reset due to an incomplete reset sequence. Once in this state the only way to exit the state is to reset the switch. To work around the issue the a switch reset may be initiated by terminating the slice agent associated with the switch. (1249888)
- In a configuration session, BGP's neighbor interface command is evaluated against an interface's current state instead of the target state. As a result, if an interface was subsumed prior to the config session commit, it's BGP neighbor related commands will be ignored even if the configuration session contains the commands required to breakout the port and activate the interface.

The workaround is to configure replace one more time with the same config file. (1252169)

- When an ACL is applied via configure-replace or an interface-range command to a large number(40+) of SVIs or sub-interfaces, or when many ACLs are applied across a large number of interfaces via configure-replace, the Strata-FixedSystem forwarding agent may restart due to a SIGQUIT signal.

Workarounds:

Option 1: Configure "hardware access-list update default-result permit" command. This command will permit traffic during ACL updates instead of denying all traffic.

Option 2: Reduce the number of ACL applications per configuration session. (1262317) (1)

- On certain devices, the 'show forwarding destination [detail]' command may continuously fail after an initial failure until the device is rebooted. (1263648) (1)

- Configuring EVPN L3 routing on unsupported releases may result in continuous restart of StrataL3Unicast forwarding agent. Workaround is to remove the EVPN L3 configuration. (1264099)

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

Transceiver Series

- When performing an accelerated switch upgrade with the command `reload fast-boot`, some transceivers may experience a link interruption. This can occur on dual-speed QSFP100 optical modules that support both 100G and 40G, specifically when the module is configured to operate at 2x50G. Examples include the 100GBASE-PSM4 and the 100GBASE-CWDM4. (343684)
- 100GBASE-SR4 modules may not link up or experience errors when inserted into a QSFP-DD slot or an OSFP slot using an OSFP to QSFP Adapter (OQA) on some platforms.

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

SKUs Affected: 100GBASE-SR4

- SFP-10G-MRA-T 10GBASE-T ports configured with `'speed auto 100full'` or `'speed 100full'` (100M port speed) may experience extended traffic loss during the Leaf Smart Software Upgrade (SSU) boot up sequence. (866682)
- Traffic may be briefly lost when performing SSU with QSFP100 optics. There is no workaround. (957040)
SKUs Affected: DCS-7060DX5-64S
- Upon an XcvrAgent restart, 800G-2XDR4 and 800G-2PLR4 interfaces may flap briefly. (1205694)
SKUs Affected: OSFP-800G-2PLR4, OSFP-800G-2XDR4, QDD-800G-2PLR4 and QDD-800G-2XDR4

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

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

- Interfaces with certain types of optics do not come up after doing a administrative no-shut following a hitless reload operation if those interfaces were kept in admin down state prior to hitless reload. (417223)
- Some NAT rules may not work correctly if the StrataL2 forwarding agent happens to restart. Packets may not get translated as expected
The impacted rules need to be un-configured and then re-configured. (483821)
Series Affected: DCS-7300X Series
- Having any ACLs with UDF qualifiers as well as any of the following LAG/ECMP hashing features present in the configuration: RRoce hashing or TTL/HopLimit hashing, will result in incorrect operation of the LAG/ECMP hashing where the switch incorrectly hashes on IP payload. To avoid this problem the conflicting features should not be configured at the same time. (497016)
Series Affected: DCS-7050CX3, DCS-7050SX3 and DCS-7300X Series
- If NAT configuration is modified or if a NAT-enabled interface is shutdown and enabled while traffic is flowing, then some of the NAT connections might not be setup properly leading to packet loss or untranslated traffic.

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

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

Workaround is to reload the switch. (539951)

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

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

- With GRE tunnels or IPinIP tunnels or Nexthop-group tunnels configuration along with subinterface in the tunnel core interface, tunnel encapsulated packet will not be tagged with dot1q tag as configured in the subinterfaces, instead packet will go out tagged with internal vlanId assigned for the subinterface. There is no workaround for this issue other than removing subinterface configuration in the core interface. (548327)
- When VXLAN encapsulated packet is received with an inner payload tag, the switch may decapsulate and send the packet with inner tag stripped when the egress chip is not same as the ingress chip. There is no workaround to this issue. (548492)
SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F
- Some NAT translations may stop working after rebooting a switch.

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

- When VXLAN IPv6 underlay is configured and when core interface is a port channel on a SVI, continuous port channel interface flaps may result in persistent loss of VXLAN traffic to remote VTEPs that are reachable through that port channel. (567455)
- VXLAN encapsulated packets received with inner payload having more than one dot1q tag will result in outer dot1q tag being stripped when egressing out on a port. (567587)
Series Affected: 7300X3 and CCS-750 Series
SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F
- Replacing new scaled configuration with old scaled config may cause StrataL3 agent to restart. Feature will function normally after agent restart. (591050)
- When a linecard with LAG members configured statically is removed, subsequently moving or adding the LAG to a different VLAN may cause BUM traffic on that VLAN to be blackholed.

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

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

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

Workaround is to manually flush all the entries using the command "conntrack -F" in the corresponding network namespace bash mode. If configured on a non-default VRF (vrfName, for example), run the command "ip netns exec ns-vrfName conntrack -F". (615735)

- A single event upset (SEU) in the IFP_LOGICAL_TABLE_SELECT TCAM or IFP_TCAM_WIDE may not be detected. This may cause traffic to be misclassified.

A forwarding agent restart is required to forward traffic normally. (626452)

- Configuring 'switchport source-interface tx [multicast]' while the Strata-FixedSystem/Strata-Linecard agents are shutdown will result in an unexpected StrataL3 agent restart.

Workaround: Ensure the Strata-FixedSystem/Strata-Linecard agents are running before configuring 'switchport source-interface tx [multicast]' (635275)

- A single-event upset (SEU) in the IFP_TCAM can lead to forwarding agent restart, after which the switch forwards traffic normally. (641105)
- With EVPN multicast solution configured, first few underlay multicast <S,G> entry may not get programmed; leading to any traffic arriving with the underlay <S,G> to be punted to CPU.
Workaround is to create an extra VLAN/VRF to underlay multicast group mapping before configuring the rest of the mappings. (643866)
- During config replace with NAT configuration, TCAM churn may result in forwarding agent restart. (676117)
- With IPV6 VXLAN Underlay configuration, when routing is toggled VLAN encapsulated packets sent on routed ports may be sent incorrectly with a VLAN tag resulting in packets getting dropped.
Workaround is to flap the core underlay interface. (679980)
- SSO switchover may result in remote ARP entries getting refreshed resulting in brief traffic loss for traffic routed to VXLAN remote hosts. (724315)
SKUs Affected: CCS-755-CH, CCS-758-CH, DCS-7304 and DCS-7308
- With a large route scale, executing the CLI command "show platform trident l3 software alpm-table routes" may result in the restart of multiple forwarding agents. There is no workaround for this issue. (762247)
- During hardware reset, routed multicast traffic egressing private VLAN tagging enabled port channel may not recover. This behaviour is not always consistent, but can hit once in a while.
To work around this issue, disabling and enabling secondary trunk tagging on the port-channel configuration, recovers the traffic. (790247)
Series Affected: CCS-720DF, CCS-720DT, CCS-720XP, CCS-722XPM, DCS-7010TX, DCS-7050CX3, DCS-7050SX3 and DCS-7050TX3 Series
- Flapping of all underlay links through which a remote VTEP is reachable may lead to StrataL3 agent restart. (804231)

- On devices running EVPN VXLAN multicast OISM, Overlay multicast traffic with TTL1 arriving on VXLAN enabled VLAN in a VRF with `evpn multicast` enabled is not VXLAN encapsulated when there is a receiver in the same VLAN on remote VTEP.

Use TTL>1 or enable `redistribute igmp` under router BGP for the VLAN on the receiver VTEP. (833210)

- Executing "platform trident * reset" continuously may result in unexpected restart of StrataL3 forwarding agent. (854688)
- When MoFrr is configured with IPv4 and IPv6 multicast groups having the same interested receivers, core interface flaps can lead to StrataL3 agent restarting unexpectedly.

Workaround is to disable and re-enable MoFrr. (860221)

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

- If StrataNat agent is killed or shutdown while NAT traffic is flowing, it may cause the StrataNat agent to continuously restart. System needs to be rebooted to recover from it.

Workaround is to stop NAT traffic before shutting down the StrataNat agent. (862801)

Series Affected: CCS-720DP, CCS-720XP, DCS-7050CX3 and DCS-7050SX3 Series

- When the StrataMirror agent restarts while egress port mirror-to-cpu is configured, mirrored traffic may be flooded temporarily on the VLANs of the mirror source port(s). Workaround is to use a recirculation channel with egress mirror-to-cpu. (883792)
- Global routing disable with MSS-G enabled can result in switch local IPs being unreachable.

Workaround: Enable routing globally. (890058)

- StrataL2 agent might assert when MAC deletion and Strata Slice agent restart happens simultaneously.

System will recover automatically. (897491) (1)

- In Active-Active NAT synchronization mode between the two peers, TCP and ICMP connection establishment might fail if the control packets are received in the other peer.

A workaround is to make sure that all the control packets for a specific connection arrive on the same peer. (954033) (1)

- Moving from MPLS configuration to VXLAN configuration and performing SSU or StrataL3 agent restart may result in VXLAN with IPv6 underlay to be dropped.

Workaround: Reload the affected device. (965290)

- Cos value in the egress packet after VXLAN decapsulation is not set correctly in a MultiChip switch , when ingress and egress ports are in different chips. Cos value is always 0 even if cos rewrite is enabled. (988967)

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

- ECN bit may not be set correctly in VXLAN Encapsulated packet in a MultiChip switch , when ingress and egress ports are in different chips. (992398)

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

- If NAT fragmentation settings are changed where there are active NAT connections/traffic, NAT translations might stop working. Recommendation is to change NAT config when there is no NAT traffic.

Work-around is to restart StrataNat agent. (992522)

- If a large amount of flows are continuously being re-learned, Strata slice agent may restart. (993211) (1)
- VXLAN IPv6 underlay encapsulation flows (tunnel initiation) may experience a few milliseconds traffic loss when the new members are added to the underlay ECMP. (1037644)
- Traffic to a multihomed MAC destination may see packet drops during VTEP reachability change when putting spine devices into maintenance mode.

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

- VXLAN traffic may be dropped when changing between IPv4 and IPv6 encapsulation if remote VTEP reachability shares adjacencies for both IPv4 and IPv6 VTEPs (Use of IPv4 over IPv6 nexthops or enabling adjacency sharing).

Workaround: Flap the VXLAN interface. (1068640) (1)

- With EVPN multicast configured, first few packets that are sent to CPU till the multicast route is learnt in hardware will not be routed. (1077713) (1)
- A single-event upset (SEU) in the IFP_POLICY_TABLE_WIDE can lead to forwarding agent restart, after which the switch forwards traffic normally. (1081241)
- When config replace is done with a large VLAN translation configuration, it is possible some of the VLAN translation entries do not get programmed.

A workaround is to restart the StrataL2 agent. (1081790) (1)

- If the device is restarted when overload mode dynamic network address translation is configured and the NAT traffic is running then the NAT translations may brake and some or all of the flows subject to NAT are being forwarded without translation after the device completes the restart.

Workaround is to manually flush all the NAT entries using the command "conntrack -F" in the corresponding network namespace bash mode. If configured on a non-default VRF (vrfName, for example), run the command "ip netns exec ns-vrfName conntrack -F". (1122152) (1)

- Parity error in the IFP_POLICY_TABLE_WIDE table can cause a fatal error during SSU, leading to a cold reboot and extended traffic loss (1130453)
- Executing "show platform trident l3 software pivots" command when the device is running in ALPM mode ("forwarding table partition mode 4") with IP routing disabled may result in unexpected restart of the forwarding plane agent. (1148631) (1)
- When a NAT translation is being offloaded to hardware, some return packets of that connection might be lost while the nexthop is updated for the NAT entry.

Workaround is to pre-establish a connection to cache the nexthop. Also adding a static route instead of a Null0 route for the NAT IP might work in some cases. (1173317)

- A port belonging to a VXLAN VLAN supporting IPv6 encapsulation may tag packets belonging to its native VLAN. (1185201)
- When StrataL3 agent is restarted immediately after disabling uRPF feature, hardware multicast routing may become disabled from that time on. Workaround is to reload the device. (1213671) (1)
- After a StrataL3 restart, EVPN Multi-Home MAC entries may become stale and remain programmed in hardware with outdated values. This can occur when a MAC entry, after going through an aging interval with no traffic, receives traffic again and transitions to a remote-learned state. If StrataL3 is restarted any time after this transition, the stale entry may persist. Restarting the StrataL2 agent clears the stale entries and resolves the issue. (1225739)
- Performing SSU with VXLAN MLAG configuration where the MLAG Port-Channel is used for VXLAN VTEP underlay reachability may result in additional traffic outage. (1231988)

- Changing the switchport mode or native VLAN on an inactive port can result in packet drops on active ports.

A workaround is to restart the StrataL2 agent. (1255142)

- VXLAN BUM traffic may see over 1second traffic loss during SSU. (1260806) (1)
- The StrataL2 agent may crash repeatedly if per-MAC ACLs are enabled through both MAC traffic policy CLI configurations and Dot1x. This is unsupported so the recommendation is to remove either configuration. (1264185)
- In the case of EVPN Bud Role traffic, when an underlay multicast encapsulated packet is received, if the inner multicast packet is an L3 packet and results in a multicast miss, such packets are dropped, causing traffic loss.

The workaround is to issue the command "fab internal_lpbk evpn_bud 65 131". (1272387) (1)

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

7358 and DCS-7050X4 Series

- StrataMmu agent has high CPU usage when PFC watchdog is configured with a low polling interval, such as 0.010. (636394) (1)
- Traffic loss may be seen during StrataL2 restart (745553)
- MAC Moves may happen for very small duration when forwarding agent restarts on any of the peers of a MLAG. In case of VXLAN BGP EVPN, this may result in the peer learned MACs being locally learned and being advertised back to remote VTEPs. (749358)
- Ipv4 microBfd packet sent to non-routed port with interface destination mac will be treated like regular L2 traffic. (769736)
- When BFD is configured but the session is down, BUM traffic can cause loops in a LAG. (778369)
- CLI commands prefixed with `platform trident adaptation-table` are unsupported and will cause repeated restarts of the StrataLag agent. (791141)
- IGMP Snooping does not work when ingress vlan translation is enabled. (814426)
- Scaling down static MAC address configuration while traffic with the same source MAC address as one of the removed MAC addresses is running might result in the MAC address not being learnt.
A workaround is to execute the "clear mac-address-table dynamic" command to prompt relearning of such MAC addresses. (852039)
- Scaling down static MAC address configuration while traffic with the same source MAC address as one of the removed MAC addresses is running might

result in the newly learnt dynamic MAC address still being considered a static MAC address.

A workaround is to execute the "clear mac-address-table dynamic" command to prompt relearning of such MAC addresses. (852203)

- If an interface is shutdown, static MACs on that interface may not be counted in number of allowed addresses for port-security (862079) (1)
- When routes are updated while nexthop utilization is exhausted, an unexpected agent restart of StrataL3Unicast may occur. (881945)
- Multicast MAC routes will persist in hardware after config-replace although the configuration has been removed. "show mac address-table multicast" will display that the software table is empty, but "show platform trident mac-address-table" will display that the routes are still present in hardware and consuming resources. To clear the stale MAC table entries and free up resources, a switch reload is required. (907142)
- At very high ALPM scales, routes that are not programmed due to hardware resource exhaustion may incorrectly indicate they were successfully programmed. The switch will not syslog that IP resources are full due to these failures in this scenario. Subsequent route insertions will correctly indicate they are unprogrammed and will trigger a resource full syslog, but the existing failed routes will still indicate they are programmed.

Workaround is to reduce the route scale. (928775)

- Encapped VXLAN packets can be dropped after multiple StrataL2 and StrataL3Unicast restart at the same time. The workaround is to restart StrataL3Unicast alone (949366)
- StrataL3Unicast may unexpectedly restart when route churn occurs at high ALPM route scale. (971310)
- When the dna chip-profile is active. Greenspan sessions that have a nexthop destination that is the same as a mirror source can cause an infinite packet loop.

The workaround is to either change the GRE nexthop to another port, or change the mirror source. (1048962)

- When IGMP snooping is enabled on a VLAN, under some conditions - (1) VLAN is also a VXLAN VLAN or (2) MLD Snooping is also enabled in the VLAN - non-IP unknown multicast traffic may not be flooded to all member ports in the VLAN. Workaround is to configure "unresolved packet flood" under "router multicast". Workaround requires "pim ipv4 sparse-mode" to be configured under the SVI. (1072368)

- When config replace is done with a large VLAN translation configuration, it is that possible some of the VLAN translation entries do not to get programmed.

A workaround is to restart the StrataL2 agent. (1081790) (1)

- VLAN mapping may stop working after StrataL2 agent restart.

Workaround:

- execute the CLI:
platform trident reset (1126331)

- During device reload with multiple ACLs configured to full TCAM scale, "Strata-FixedSystem" agent may restart multiple times post device reload.

To workaround this issue, remove the ACLs and re-apply them. (1153186)

- L2 forwarded traffic is unexpectedly dropped when the packet length exceeds the programmed MTU, even though MTU enforcement should not apply on the given ingress interface. (1165583)
- Link flaps on interfaces with auto-negotiation configuration present may cause the forwarding agent to restart and all links to flap. This would primarily be an issue with links which are prone to flapping. (1249552) (1)
- IPv6 multicast packets with destination IP address in the FF03::/16 to FF0E::/16 range are not routed. (1258179)
- When an ACL is applied via configure-replace or an interface-range command to a large number(40+) of SVIs or sub-interfaces, or when many ACLs are applied across a large number of interfaces via configure-replace, the Strata-FixedSystem forwarding agent may restart due to a SIGQUIT signal.

Workarounds:

Option 1: Configure "hardware access-list update default-result permit" command. This command will permit traffic during ACL updates instead of denying all traffic.

Option 2: Reduce the number of ACL applications per configuration session. (1262317) (1)

- Configuring a monitor session to GRE tunnel with an IPv6 destination address residing on an interface in trunk mode with a GRE key configured may cause the `StrataMirror` agent to endlessly restart.

Workaround is to use IPv4 or remove the GRE key. (1272638)

Limitations and Restrictions in 4.34.3.1M

Accelerated System Update

- During a hitless reload upgrading from a release before 4.21.3, more than 3 LACP packets could be transmitted within a one second interval. This should not cause problems. (338048)
SKUs Affected: DCS-7260CX3-64, DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F and DCS-7260CX3-64E-R
- Aborting 'reload fast-boot' or 'reload fast-boot now' command with Ctrl-C is unsupported. It may lead to the incorrect reload cause in `show reload cause` in the subsequent reload. (344574)
- SSU of VRRP Master and Backup routers at the same time is unsupported and may cause extended traffic loss. Successfully SSU one router then SSU the other router. (345657)
- On systems with 4GB of /mnt/flash space, there may be insufficient space to store both an old swi image and a new swi image to perform an SSU upgrade. The workaround is to delete the old swi image from /mnt/flash prior to performing the upgrade. (385364)
- Performing SSU to boot into an image with this limitation may require that the currently running image is deleted to ensure that there is sufficient space available to perform the upgrade. (525870)
- The "show interfaces transceiver hardware" CLI command may fail to run temporarily after reload. (854489) (1)
- Performing SSU reload with a transceiver that is unsupported in the source image but supported in the destination image may cause the port not to link up after SSU reload completes.

The workaround is to hotswap the affected ports after SSU reload. (955515)

BGP EVPN L2/L3 VXLAN/MPLS

- On a BGP EVPN device, if a VLAN-aware bundle is configured to include VLANs from multiple IP VRFs and if there are any IP addresses reused across multiple VRFs, then only one of the corresponding MAC/VLAN ARP bindings is distributed over EVPN. As a workaround, configure VLAN-aware bundles such that all VLANs in the bundle are in the same IP VRF. (514980)

- VRRP+VXLAN is not supported in deployment with EVPN VXLAN. (549525)
- In an EVPN/VXLAN environment with VLAN-Based MAC-VRF, type 2 and type 3 routes are incorrectly imported when the route-targets match even if the VNIs are different.

The workaround is to use different route-targets for the VLAN-Based MAC-VRF's to prevent the routes from being incorrectly imported. (631589)

- IP ARP proxy is not compatible with EVPN VXLAN enabled VLANs.

Disable IP ARP proxy on VLANs with EVPN VXLAN enabled. (710344)

- In a EVPN DCI deployment, Symmetric Integrated Routing and Bridging is not supported in the default VRF. (733143)
- In a EVPN VXLAN deployment with IP prefix routes with a gateway IP, address resolution may be unevenly distributed across various VXLAN endpoints, VNIs and router MAC addresses, each destination (VTEP/VNI/router MAC) as identical VIAs are not considered when forming ECMP. (810595)
- In networks using EVPN Active/Active multihoming with the modulus designated forwarder election algorithm, EOS versions prior to 4.29.0 consider all of the VLANs configured in a vlan-aware-bundle MACVRF when calculating the DF. EOS versions starting with 4.29.0 only consider VLANs configured in both the MACVRF and on the device itself. This can lead to disagreements on the elected DF if the range of VLANs configured in a bundle MACVRF is wider than the range of VLANs actually configured on the device, and the devices in the ethernet segment are running disparate EOS releases on either side of version 4.29.0. If different devices in the same ES elect different DFs, traffic loss can ensue.

As a workaround, either make sure all devices participating in an EVPN active/active multihoming ethernet segment are running the same version of EOS, or alternatively, before upgrading any device, ensure that the lowest VLAN configured in a vlan-aware-bundle is also configured on the device itself. This will ensure that the VLAN considered as the lowest VLAN in the MACVRF is the same pre and post-upgrade. (868750)

- When an ethernet segment comes up on a multi-homing peer there can be an outage for several seconds of traffic destined to IP/IPv6 addresses whose ARP/ND entries age out on a multi-homing peer that is already forwarding the traffic.

If the interface is being brought up administratively (e.g. bring interface out of maintenance mode) then a workaround can be used to reduce the outage. The workaround is to force an immediate refresh of the ARP/ND cache on multi-homing peer not under maintenance, before initiating the end of maintenance cycle. The immediate refresh should be triggered using 'refresh

arp-cache vrf all' for IPv4 and using 'refresh ipv6 neighbors vrf all' for IPv6. The refresh should be followed by waiting for a sufficient amount of time for the it to complete, which can be calculated as follows:
(ARP or ND scale) x (3 refresh attempts) / (refresh rate of about 500 addresses per second)

Note that the ARP timeout must be sufficiently large (30 min or more) during the operation for the workaround to be effective. (1051719)

- In EVPN setups, the MAC duplication detection logic helps prevent control plane churn by adding MACs that move over a certain number of times over a certain rolling period of time. This is typically caused by an L2 loop. While control plane churn is mitigated, the data plane traffic remains affected from the L2 loop.

An event-handler such as one attached to the bug can be deployed as a workaround. The event-handler is triggered when a syslog is thrown in response to a host being denylisted, and it removes the VLAN corresponding to the host from the offending interface. Since the syslog does not contain the offending interface, the interface lookup is done in a separate call. If the host is removed before the interface could be obtained, the event-handler would not take any action (1117596)

Containerized EOS

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

vEOS Router / CloudEOS

- The /mnt/flash/cloud_ha_config.json based configuration for the CloudHa agent is not supported from this release. If you are upgrading from an older image (< 4.20.5) please reconfigure HA using EOS CLI. Please see information on converting json file based config to CLI commands in the vEOS Router Configuration Guide. (264660)
- While vEOS instances in an AWS cloud can be created using either a Bring Your Own License (BYOL) or Pay As You Go (PAYG) format, there currently is no show command in vEOS to easily help the user or support determine which mode this instance uses. (272649)
- Tx/Rx ring parameters of an "Elastic Network Adapter (ENA)" interface cannot be set through config.json in vEOS. (276382)
- vEOS supports up a maximum of 8 physical cores. With hyperthreading enabled it is 16 vCPUs (278286)
- In dpdk based vEOS(MODE=sfe in /mnt/flash/veos-config), even a packet matches a data plane ACL rule with log action, the log message might not show up if

there is a similar rule with more exact ttl/dscp value match before that rule. (293641)

- In dpdk based vEOS(MODE=sfe in /mnt/flash/veos-config), SR-IOV mode doesn't work with Intel x520/82599 on ESXi. As a workaround use mixed mode instead of the SRIOV only mode in ESXi (296718)
- Hot detaching network interface instances is unsupported. Reboot the instance after detaching a network interface. (306132)
- CloudEOS does not support more than 10k /30 routes. The SFE agent will restart if more than 10k /30 routes are added to the system. The workaround is to reduce the number of /30 prefixes. (540204)
- NIC i40en driver version 2.2.4.0 is required on ESXi 7.0 host for SRIOV virtual function trust mode on/off functionality. (680174)
- NAT interface configuration mode is not supported in SFE. This needs to be kept in mind when migrating to SFE mode from Sfa.

The NAT interface configuration will need to be changed to profile based configuration in SFE mode. (682248)

- The command "show flow tracking hardware flow-table" may trigger OOM events when there are over 1 million flows present in the system. This could occur as a result of many short lived flows and the need to process many adds or deletes.

The flow-tracking information can be viewed correctly on Cloud Vision As A Service (721460)

- The system does not support the interface tx-queue shape rate commands, however the CLI still permits the configuration. Entering these commands will have no effect. (751249) (1)

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

CVX

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

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

- Switches and CVX instances participating in a CVX setup must either all run 4.15.4F or later, or all run images from before 4.15.4F. (146664)

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

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

- Intercept hosts on trunk ports with native vlan are only supported on 7050X2 platforms (257580)
- The Macro-Segmentation Service (MSS), requires MLAG fast redirection be enabled on the service MLAG TOR pair, when configured to work with a L2 transparent FW. This configuration minimizes traffic loss when any individual interface in the MLAG port channel goes down. (268291)
- The MssClient agent, can restart unexpectedly, when the MacroSegmentation Service, deployed with a L2-transparent firewall and is configured to intercept more than 10,000 end-points on a single switch (a number that's far greater than the total number of intercepts possible on the switch). (283221)
- The MacroSegmentation Service (MSS) allows users to configure multiple tags on CVX that can be used in firewall policy definition, to identify policies MSS should work on. It is however not possible to delete one of the multiple tags configured.

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

- In Macro-Segmentation Service (MSS) for L3 Firewalls , Virtual SVI IP addresses cannot be treated as intercepted hosts and thereby cannot be added to any redirect policy. However, these IPs can still be part of offload policies, even though implicit redirect will not work for them. (372504)
- If the Macro-Segmentation Service (MSS), working with a L3 firewall, is configured to intercept traffic from a host generating multicast traffic, the traffic can be black-holed.

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

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

- MLAG devices in an environment with the Macro-Segmentation service (MSS) configured (with L3 firewalls), should have a high-priority DirectFlow rule configured to ignore traffic over the MLAG peer link. This can be configured

using the following flow definition on each switch (assuming Po1 is the MLAG peer link):

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

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

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

- The MssPolicyMonitor agent might unexpectedly start running and eventually exit when the CVX functionality is disabled on the master CVX instance. This does not have any side-effect on the operation of the Macro-Segmentation Service (MSS). (378077)
- When the Macro-Segmentation Service (MSS) is enabled and works with a L3 firewall, there can be a momentary traffic outage on reloading a MLAG peer if the CVX to VTEP (MLAG switch) connection is not re-established before the MLAG reload delay expires at the reloaded peer. (418560)
- If MSS policy enforcement rules configuration is changed from verbatim to group verbatim, the SandAcl agent can restart unexpectedly (433592)
- When Macro-Segmentation Service (MSS) is enabled on the CVX in an environment where no switch is configured as CVX client, multiple "excessive warmup delay" syslog messages for MssPolicyMonitor agent might be noticed. (497132)
- In a Macro-Segmentation Service (MSS) deployment with a PaloAlto firewall, enforcement policy configured with application using dynamic port assignment is not supported. (515749)
- Macro-Segmentation Service (MSS) running in consistent enforcement mode may emit "CVX-3-MSS_HOST_UNPROGRAM_ERROR" syslog messages for the hosts

attached to only one MLAG peer when that peer is reloaded.

These messages can be ignored. "CVX-3-MSS_HOST_RECOVER" messages will be observed as soon as these hosts are learned again at the reloaded MLAG peer afterwards. (524705)

- For Macro-Segmentation Service (MSS) deployment running in consistent enforcement mode, "CVX-3-MSS_HOST_UNPROGRAM_ERROR" or "CVX-3-MSS_RULE_UNPROGRAM_ERROR" syslog messages may be seen when a switch is reloaded (planned reload).

Use "clear cvx connections client <switch-id>" command to clear the switch from CVX registration. (524717)

- In a Macro-Segmentation Service (MSS) deployment in a multi-VRF environment with EVPN as control plane, when an MLAG peer is reloaded, CVX-3-MSS_HOST_PROGRAM_ERROR message is logged for the intercept hosts followed by CVX-3-MSS_HOST_RECOVER message. (525851)
- In a Macro-Segmentation Service (MSS) deployment with a Palo-Alto Firewall, when using an aggregator interface in a virtual system, all its member interface must be in the same virtual system. (528511)
- In a Macro-Segmentation Service (MSS) deployment with a Palo-Alto firewall, a policy with multiple source or destination zones will be enforced over only one source zone and/or one destination zone.

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

- In a Macro-Segmentation Service deployment with Layer 3 firewall (MSS-FW), when a policy matching a subnet with a mask less than /16 is tagged as a group-based redirect or a regular offload policy, the service may become non-operational. However, any subnet mask length is supported for all variants of verbatim policies. (716283)
- The Media Control Service (MCS) with CVX high availability (CVX-HA) will be hitlessly upgraded if the deployment is upgraded to a version of EOS greater than or equal to the version specified in the "Version Introduced" field of this defect and less than that of the "Version Fixed".

As a result of this defect, multicast routes configured by the MCS will be removed. (907855)

Enterprise routing

- When large number of subinterface (~4k) are configured, restarting SFE/ NacIes agents may cause SFE agent to continuously restart unexpectedly

(856104)

SKUs Affected: AWE-5310 and AWE-5510

- Any config change in DSCP value configured via AVT policy/Application profile is not updated for existing flows. New flows created after the config change inherit the updated DSCP value. (873846)
- If SfeAgent is administratively disabled or unexpectedly restarting, transceiver diagnostic data may not be available. Given that SfeAgent is unavailable for approximately 10 minutes, the a message may be logged for TRANSCEIVER_SMBUS_COMMUNICATION_FAILURE. Once SfeAgent is stable, the workaround is to virtually remove and reinsert the transceiver with the interface config level command 'transceiver diag simulate removed' followed by 'no transceiver diag simulate removed'. This will temporarily bring down the links associated with the transceiver. (907440)

SKUs Affected: AWE-7220RP-5TH-2S-F

- Pluggable BASE-T SFP transceivers are not supported on AWE-7220RP-5TH-2S-F. This is a hardware limitation. There is no workaround. (909692)

SKUs Affected: AWE-7220RP-5TH-2S-F

- Some pluggable SFP transceivers may be unsupported. Unsupported optical transceivers will have their transmitters disabled. The Ethernet Interfaces associated with the unsupported transceiver will be disabled with reason "xcvr-unsupported". A system log for TRANSCEIVER-DISABLED may be seen.

The list of supported transceiver types are as follows. Some transceiver SKUs which fall under this list may still be unsupported by this platform.

- SFP+ active copper cables
- SFP+ active optical cables (AOCs)
- SFP+ 10G SR/LR optical modules
- 1000BASE-SX/LX optical modules
- SFP+ direct-attach twin-ax cables which comply to SFF-8431 v4.1 and SFF-8472 v10.4

There is no workaround as this is a hardware limitation. (940078)

SKUs Affected: AWE-7220RP-5TH-2S-F

SDWAN

- In a High Availability configuration, packets may get dropped if the MTU of the LAN DPS path between the HA peers is not at least 4 bytes greater than the MTU of the WAN DPS paths.

Configure an MTU for the LAN DPS path between the HA peers which is at least 4 bytes greater than the MTU of the WAN DPS paths to avoid packet drops.

(1104868)

Series Affected: AWE-5000, AWE-7000 and CloudEOS VM Series

- In a High Availability configuration with internet exits present on both the HA peers and default route advertised from both the peers as an ECMP to the remote sites, all packets of a given internet bound flow from a remote site may not be routed to the same HA peer causing different NAT IPs to be used for different packets of the same flow depending on which peer the packet lands on resulting in application traffic disruption.

Enable NAT sync on both the HA peers so that the same NAT IP is used for the flow irrespective of which peer forwards the flow to the internet.

(1135268)

Series Affected: AWE-5000, AWE-7000 and CloudEOS VM Series

- In a High Availability configuration, DPI applications recognition may be incorrect if multiple LAN DPS paths are configured between the HA peers.

Use a port channel if multiple HA peer links are required. Otherwise, DPI applications recognition may be incorrect. (1161876)

Series Affected: AWE-5000, AWE-7000 and CloudEOS VM Series

EOS SDK

- When an empty argument is provided to EosSdkRpc BgpLinkWeightMgrService service's bulk_link_weight_set() RPC, the error code returned is INVALID_ARGUMENT (3) instead of OUT_OF_RANGE (12) (1268485) (1)

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

General

- CLI does not allow access to files whose names contain spaces. (539)
- 'speed forced 1000full' is not supported on interfaces with the "Management" prefix such as mal. Use speed autonegotiation instead. (1506)
- Ports configured as mirroring destinations will be shown as active even when they are shutdown. (65221)
- When VmTracer is used in a network that includes Cisco devices, ESX hosts connected to the Arista switch will see CDP frames from the Arista switch as well as from the other Cisco devices. VmTracer will display the VM info when the ESX host reports Arista's CDP info, but then will delete it when the ESX host reports other CDP info.

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

- EOS SDK-based applications must be run via the 'daemon' command configuration. They cannot be run directly from bash. (158431)
- EOS extension scripts and agents which do not use EOS-SDK will need to be modified to work in 4.16.6M and beyond. (158467)
- EOS swix extensions containing RPMs which were based out of or procured from Fedora distributions prior to Fedora 18 may fail to install due to dependency issues. Suggested workaround is to recreate the swix files with the corresponding RPMs from Fedora 18 distribution. (162325)
- Packets with size greater than MTU of egress interface will not honor directflow flows and may not be forwarded (180927)
- SSO is hitful with subinterfaces. Protocols running on subinterfaces may experience session flaps during SSO. (188070)
- ACL drop counters cannot be cleared. (202905)
- when ACLs are configured with stats over port-channels, these stats are kept at the physical port level not at the port-channel level. (203131)
- With international and HW-REV-01 images and the introduction of "ip access-group <access-list name>" other configuration commands beginning with 'ip' that are configurable in the global config mode are not accessible in "router bgp" mode.

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

- Changes made to running-config after a config session is created might be reverted when the config session is committed which contains changes in similar or related areas. Use 'show session-config diff' inside the config session before committing to verify what changes will be made. (226850)
- The uptime in the output of "show module" may not match the uptime as shown by "show version" or "show uptime". The uptime in "show module" represents the time the module has been up since the last change in the "Status" column in the output of "show module", which differs from the meaning of uptime in "show version" and "show uptime". (248230)
- To avoid interoperability problems when using strong SNMPv3 encryption algorithms, follow the following minimums:
 - When using AES192 for privacy, use a minimum of SHA224 for authentication.
 - When using AES256 for privacy, use a minimum of SHA256 for authentication. (268290)
- Mirroring to an MLAG port channel is not supported. (276973)

- Copy commands with sftp: or scp: URL do not work via eAPI by passing password through the "input" parameter. SSH keys have to be setup to run the commands without a password. (283716)
- EOS support alphanumeric VLAN names, however ODL gives an error when reading the leaf /interfaces/interface/routed-vlan/config/vlan if the name contains alphabetic characters.

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

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

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

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

- If an EOS CLI command is used to create a reference to a VRF that does not exist, a subsequent OpenConfig YANG operation may cause the VRF to be created in EOS. (411908)
- Large amounts of monitoring traffic processed at the control plane via aggressive sflow sampling or mirroring to CPU could affect other control plane functions due to CPU bandwidth contention. (415992)
- IP Locking may send out additional queries when STP changes state (even possible in portfast config which is common for the IPLocking use-case). (417806)
- When hardware or software flow tracking is enabled and tracking flow is a VXLAN flow, the exported IPFIX packet may contain incorrect inner VLAN ID. (421437)
Series Affected: CCS-720XP Series
- When "logging format hostname ipv4" is configured, log messages may not display the correct IP address if the IP address is updated. (422666)
- Egress MAC ACL logging is not supported (433169)
- In the 'show interfaces interactions' CLI, the list of speeds an interface is limited to does not account for the limitations that result from speed-group configuration. (434363)
- SSH may fail when sFlow and Management SSH have the same QoS DSCP value configured, and sample rate for sFlow is set to dangerous. (434763)
- Mirror-On-Drop does not monitor IP packets with an invalid IP version. (443839)
- When we first configure a port-channel lag_type, we must set values for both aggregation/config/lag-type and ethernet/config/aggregate-id (i.e. member configuration) in the same transaction. The lag-type is populated in EOS once the port-channel has one member port and remains set. (445421)
- "show mac/ip/ipv6 access-lists summary" now excludes remarks from "total rules configured". (452368)
- When eAPI is configured with invalid SSL profile, nginx will stop running and all http/s requests will be refused (474422)
- Currently EOS does not support parsing and separating multiple ACL rules in the same NAS-Filter-Rule AVP using zero byte delimiter. (475210)
- EOS-2GB.swi is no longer available. (475680)
- Installed licenses without corresponding licensing features configured might not be visible from the CLI (504793)

- CVE-2019-16905: The version of openssh shipped in EOS does not enable support for XMSS, so EOS is not affected by CVE-2019-16905. (511376)
- CVE-2020-1968: EOS is susceptible to CVE-2020-1968. The workaround is to only use ephemeral (DHE) ciphersuites. (519024)
- If Octa, OpenConfig or TerminAttr are started under low memory conditions, they can cause a spike in CPU usage before unexpectedly restarting. (534841)
- Configuring NTP to serve clients in more than 500 VRFs fails with a syslog message "ntpd: Too many sockets in use". (543116)
- When the "tunnel destination" address is configured as a broadcast address for static tunnel interfaces then there could be packets drops during decapsulation at the tunnel endpoint. (544032)
- For a given interface, multiple simultaneous MAC address flush requests on different VLANs may result in clearing all MAC addresses on that interface. This is an optimization to support a higher scale of VLANs and interfaces. To disable this optimization, use the "mac address-table flushing interface vlan aggregation disabled" CLI configuration. (544712)
- The "dot1x re-authenticate" command is not supported for failed supplicants. Workaround is to flap the interface. (567028)
- Removing "vlan" key-field from custom TCAM profile for ingress security ACL feature causes ACL application on SVIs to fail.

To workaround this issue, add "vlan" key-field to custom TCAM profile for ingress security ACL features. (570680)

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

- In SSO redundancy mode, the switch may take longer (more than 20 minutes) to be declared switchover ready. This is to accommodate for all agents to become ready for switchover. (572615)
- When unsupported transceiver support is unlocked using a key, the ports that are in disabled state before unlocking will continue to be disabled.

The workaround is to remove such transceivers and re-insert them after unlocking the unsupported transceiver support. (574146)

- The 'logging level all' CLI expands to individual commands in the configuration, preventing dynamic updates of the logging list as part of EOS upgrades. (579047)
- TACACS+ authentication and authorization through SSH does not work with Pulse Secure server because EOS sends different port names in authentication ("ssh") and authorization (actual vty) requests which are rejected by the server. (583678)

- Sampled flow tracking exports IPFIX flow records with egress interface as 'discard' for flows on pseudowire patched interfaces. (603916)
- A CPU internal error (BIT30 SyncFlood BIT11) may cause an unexpected reboot (623392)

SKUs Affected: DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7050SX3-48C8-F, DCS-7050SX3-48C8-R, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050SX3-96YC8-F, DCS-7050SX3-96YC8-R, DCS-7050TX3-48C8-F, DCS-7050TX3-48C8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-36S-F, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-36S-F, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F, DCS-7280PR3K-24-F, DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R, DCS-7280SR3K-48YC8-F and DCS-7280SR3K-48YC8-R
- IP Locking does not support supervisor stateful switchover (SSO). (631563)
- Dot1x agent incorrectly accepts larger than advertised values for idle timeout and session timeout RADIUS response attributes (639283)
- /mnt/flash cannot be used as the log archiver destination on systems formatted as VFAT. To check if a product is formatted as /mnt/flash, from bash, run "mount | grep flash", and look for "type vfat" or "type ext4". If you see "type vfat", then that product's /mnt/flash storage is formatted as VFAT. (642415)
- "mac security profile" is not supported on L2 subinterfaces. (680967)

SKUs Affected: 7800R3-48CQM2-LC, 7800R3K-36DM-LC, DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- When Path MTU discovery is enabled under static tunnel interfaces using "tunnel path-mtu-discovery", the discovered MTU is only used for software forwarding in the slow path. When MTU is configured under static tunnel interfaces using "mtu" config, the configured MTU value is used for hardware forwarding. (681101)
- OpenConfig does not support BFD configuration for VLAN interfaces. BFD configuration under VLANs must be removed for any gNMI set operation to succeed. (682920)

- If RFC2544 test is started on a sub-interface and the parent interface has "traffic-loopback source system device mac" configured, the test packets will get into infinite loop.

Unconfigure traffic-loopback on the parent interface using "no traffic-loopback source system device mac" to stop the traffic loop. (683129)

- If the AVA switch sensor ("monitor security awake") is enabled along with the TerminAttr (CloudVision), the TerminAttr agent will also process IPFIX packets sent to the CPU, and the "Traffic flows" feature on the CloudVision portal is automatically enabled and can cause additional CPU usage. If "Traffic flows" feature is not required on CloudVision, disable TerminAttr IPFIX processing by adding "-ipfix=false" argument to "exec /usr/bin/TerminAttr" under "daemon TerminAttr" configuration. (688791)

Series Affected: CCS-720XP Series

- EOS images are optimized during installation on selected platforms. The optimized EOS image saved on non-volatile storage may be incompatible with other platforms. An attempt to install an incompatible SWI will result in the following warning: "Next image does not support required optimization". Only install full EOS images or compatible optimized EOS images. (691810)
- From EOS release 4.29.2F, the 64-bit image of EOS supports "monitor security awake" or AwakeSensor.swix/NDRSensor.swix on platforms with at least 8GB memory. EOS releases prior to 4.29.2F, the 64-bit image of EOS does not support AwakeSensor.swix on any of the platforms. Use the 32-bit EOS version instead. (719835)
- System will not boot if a cable connected to an ethernet switch is connected to the console interface of the system. All system LEDs may remain red and the console cable will subsequently be unresponsive, until the cable connected to an ethernet switch is removed from the console interface of the system. (724181)

SKUs Affected: DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7050SX3-48C8-F, DCS-7050SX3-48C8-R, DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050SX3-96YC8-F, DCS-7050SX3-96YC8-R, DCS-7050TX3-48C8-F, DCS-7050TX3-48C8-R, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-36S-F, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-36S-F, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F, DCS-7280PR3K-24-F, DCS-7280SR3-40YC6-F, DCS-7280SR3-40YC6-R, DCS-7280SR3-48YC8-F, DCS-7280SR3-48YC8-R,

DCS-7280SR3E-40YC6-F, DCS-7280SR3E-40YC6-R, DCS-7280SR3K-48YC8-F,
DCS-7280SR3K-48YC8-R, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R

- Inbound packets destined for local agents on the switch (e.g., LACP, STP) will not be decoded properly when the "tcpdump" command listens on the "any" interface.

The workaround is to direct the "tcpdump" command to listen on the interface on which the packets are being received. (730292)

- PTP traffic cannot be captured using a monitor session on PTP enabled ports for 750X devices configured as a PTP Boundary Clock (738554)
- Smbus card agent may contain assertion errors when restarting due to card removal. There is no effect on service. (740093)

Series Affected: 7368 Series

- Removing one or more rule(s) with sequence numbers which causes a hole in ACL sequence number, i.e., removing rules from middle of an ACL and, re-adding ACL rules (with auto-generated sequence numbers) can cause sequence numbers to be re-used.

This series of ACL operations will cause an "Duplicate sequence number" error to be emitted in CLI.

To workaround this issue, if a rule in middle of ACL needs to be removed then,

1. Copy the rules which appear after the rule of interest,
2. Remove all rules from the rule of interest till the end of ACL,
3. Re-add rules copied in step 1. (743473)

- * network-instances/network-instance/policy-forwarding/policies/
policy[policy-id=POLICY-ID]/rules/rule[sequence-id=SEQUENCE-ID]/ipv4/config/
source-address
* network-instances/network-instance/policy-forwarding/policies/
policy[policy-id=POLICY-ID]/rules/rule[sequence-id=SEQUENCE-ID]/ipv4/state/
source-address
configures/displays at most 1 source-address prefix per rule, despite EOS CLI supporting >1 source-address prefix per rule.

As a workaround, either use OpenConfig exclusively for configuration/query/telemetry on IPv4 source-address prefixes, OR do not use OpenConfig for configuration/query/telemetry on IPv4 source-address prefixes if >1 IPv4 source address prefixes per rule have already been configured via EOS CLI. (747080)

- * system/logging/remote-servers/remote-server/selectors/selector/facility
* system/logging/remote-servers/remote-server/selectors/selector/config/
facility
do not filter syslogs sent to the remote-server based on RFC5424 facility.
Instead, these two paths configure the RFC5424 facility for all syslogs,

regardless of original facility, sent to the remote-server (filtered by having severity >= configured severity). (747096)

- AAA RADIUS servers configured with shared-secret profiles can only be used to authenticate 802.1X supplicants. They can't be used to authenticate control-panel login or CLI enable access. (750835)
- On dual supervisor products, supervisor card 1 is not always the active supervisor when executing "reload all" or powering off/on the chassis. (768242)
- JSON does not specify the precision of a Number in the specification. eAPI can generate numbers up to 6-bit, which can be truncated by some JSON clients with lower limit (e.g., some implementations have a max value of $2^{53}-1$). The only workaround is to use a different JSON deserializer (769194)
- Workaround for the restrictions on the list
 - * /system/logging/remote-servers/remote-server[host]/selectors/selector[facility severity]
is to update/replace/merge all /system/logging/remote-servers/remote-server[host]/selectors/selector[facility severity] list items at once, such that the end configuration is that all /system/logging/remote-servers/remote-server[host]/selectors/selector[facility severity] list have (a) exactly 1 list item with identical facility and severity key values, or (b) no list items. (774251)
- * system/logging/console/selectors/selector/facility
* system/logging/console/selectors/selector/config/facility
do not filter syslogs sent to the console based on RFC5424 facility. Instead, only the "ALL" facility can be configured and is a no-op. Console syslogs are only filtered via severity. (774310)
- Workaround for the restrictions on the list
 - * /system/logging/console/selectors/selector[facility severity]
is to configure at most 1 list item and the "facility" key value must be configured as "ALL" for the single allowed list item. (774317)
- When a fifth non-unique TTL value is configured for Static Interface Tunnels on Hardware Forwarded Platforms using the "tunnel ttl" config command, the Static Interface Tunnel pipeline programming may not be completed due to hardware limitations and the traffic may not flow through the Tunnel. Only 4 unique TTL value are supported across all the Static Interface Tunnels on Hardware Forwarded Platforms. If TTL value is not configured for Static Interface Tunnels using the "tunnel ttl" config command, then it implies that the TTL value is zero. This zero TTL value is also a unique TTL value. For Static Interface Tunnels to be operational, configure 4 unique tunnel TTL values across all the Static Interface Tunnels on a Hardware Forwarded Platform (782720)

- EOS with version 4.29.0 or later generate a different user id from the same username. As a result, files on flash may show random user id instead of the expected user name. For example, a file created by the "admin" user under an older version may show that it is now owned by user "106904808" after upgrading to a newer version. The workaround is to run the following command from bash:

```
find /mnt/flash -uid 106904808 -exec sudo chown admin {} \; (798231)
```

- Across all flow tracking features, when the number of configured exporters for a given tracker exceeds the maximum supported for the feature on the platform, the active exporters may differ after a reload or shutdown/no shutdown of the flow tracking feature.
Reduce the number of configured exporters to maximum supported for the flow tracking feature on the platform. (806831)

- ZeroTouch provisioning bootstrap scripts using 'import jsonrpclib' and 'import sleekxmpp' in python2 will fail.

Use python3 instead. (823133)

- A switch configured with both MACsec L2 protocol forwarding and configurable EAPOL DMAC will send the EAPOL packets with the configured DMAC to CPU instead of transparently forwarding it in the dataplane. (824887)
- Sampled flow tracking doesn't export proper nexthop in IPFIX flow records for flows using ECMP BGP routes via IS-IS SR tunnels.
There is no workaround for this issue. (825053)
Series Affected: DCS-7500R Series

- When using "flow tracking mirror-on-drop" to generate sFlow datagrams containing dropped packets, the "drops" field is used to denote the number of dropped packets, instead of the number of drops by the policer or sFlow agent as is intended for that field. (834493)
- When PTP boundary clock is configured, incoming PTP packets on a PTP-enabled access switchport that contain a VLAN other than the one configured on the access port will be sent to CPU. (837200)
Series Affected: CCS-750 Series

- Octa/OpenConfig has incorporated implementation guidance at <https://github.com/openconfig/public/blob/master/doc/implementation-guidance/leafrefs-in-openconfig.md> and the agent no longer validates leafrefs on a YANG transaction. (861998)
- In a MLAG VXLAN deployment with anycast gateways, ARP synchronization between the MLAG peers may result in packet drops in the CoPP "copp-system-

selfip" queue and adversely affect any application traffic over that queue.
(881798)

- SSH login may fail with broken pipe error when "authentication protocol password keyboard-interactive" is configured and empty password login is allowed at the same time.
Disallowing empty password with either "no aaa authentication policy local allow-nopassword-remote-login" from the config mode or "authentication empty-passwords deny" from the management ssh submode will fix this issue. (899698)
- L3 routes resolved over multiple next-hops (ECMP/UCMP) may remove otherwise reachable next-hops from their via set when more than one next-hop in the group is resolved over the same L3 interface/are in the same subnet, but only one peer in the subnet/in the VLAN has a BFD session with the router. Workarounds include removing the BFD session with that singular peer, or adding an additional BFD session with one or more other peers on that L3 interface/subnet. (901915)
- Arista Networks Product Bulletin for EOS Release 4.29.1F incorrectly included DCS-7280DR3A-36-F, DCS-7280DR3AM-36-F, DCS-7280DR3AK-36-F in the New Platform Capabilities section. DCS-7280DR3A-36-F, DCS-7280DR3AM-36-F, DCS-7280DR3AK-36-F were introduced in 4.30.0F. The Arista Networks Product Bulletin for EOS Release 4.29.1F was revised on January 15, 2024 to correct the documentation. (902473)
Series Affected: DCS-7280R3 Series
- If a VRF selection policy and a match rule is created using CLI first, and the rule is updated via gNMI, the resulting VRF selection policy configuration may have unexpected additional rule.
Workaround is to use only either CLI or gNMI instead of both. (910953)
- SNMPv2/v3 requests for data in VRFs as a context is limited to VRFs with maximum length 32.

The workaround is to limit such VRF names to a maximum length of 32 characters. (924461)
- In a VXLAN routing setup, the VxlanSwFwd agent restart time may be high (1 to 2 minutes). During this time, new ARP or IPv6 neighbor entries may not be learned over VXLAN. (924494)
- In an EVPN VXLAN deployment using IPv6 overlay routing, if a anycast gateway is configured with 'ipv6 address virtual' and a host generates a DAD solicitation for the link local corresponding to the switch virtual MAC address configured for the anycast gateway, the switch won't respond to defend the link local address. This may result that the link local address corresponding to the switch virtual MAC to get advertised by the host. (953497)

- On a 7368 system with a 7368X4-SC switchcard, a combination of 4 plugged power supplies and 2 or fewer linecards, can lead to the system not powering back on after reload. It might require removal/insertion of the power supplies for it to power on.

To avoid running into this, the recommendation is to use a maximum of 2 PSUs on any 7368 chassis running just with 1 or 2 LCs. (957422)

Series Affected: 7358 and 7368 Series

SKUs Affected: 7289

- If a forced reboot is triggered via 'reboot -f', the reload cause may not be accurate. (974023)
- OpenConfig gNMI subscriptions to paths that contain wildcards and module-prefixed identityref list keys will return no data for those paths. The workaround is to not module-prefix identityref keys in subscription paths that contain wildcards. (975123)
- gNOI killProcess RPCs require both a process name and PID to be specified in the request. (991502)
- "flow tracking cpu-queue" feature only supports a single active tracker. The show command may incorrectly show more than one active tracker. (1001487)
- When looking at the status of NTP, the DNS refresh setting is not available in any of the NTP "show" commands.

Workaround: when enabled, the setting is available in the configuration output. (1025742)

- The switch may not respond to ARP requests over LAG interfaces that are configured as forwarding-unviable. (1037955)
- The M2Crypto library is deprecated in this version of EOS. If you are using custom scripts that made use of M2Crypto please migrate to the PyCA/cryptography library for continued support and compatibility. This should not impact operators using standard EOS features. (1039341)
- BGP extension may be missing with egress sFlow if the ingress interface is a Port-channel (1041816)
- In VXLAN routing setup, if the CLI commands "no icmp helper" "icmp helper" are issued in rapid succession, e.g. tens of microseconds apart, the command can get out of sequence.

The workaround is to redo the CLI commands with sufficient gap between them (milliseconds or more). (1069967)

- In a MLAG VXLAN deployment, BFD sessions between the MLAG peers may not get established between virtual IPv6 VLAN interfaces.

A workaround is to use `ipv6 virtual-router address`.
Alternative workaround configure a static neighbor entry for each interface
`ipv6 neighbor vrf <vrfName> <m1ag peers link local address> Vlan<N>
<system mac of m1ag peer>` (1089341)

- In a VXLAN routing setup using SVIs configured with "ipv6 address virtual", a host ping to the IPv6 Link-local IP corresponding to the virtual-router MAC address won't respond. This does not affect routing function and does not affect the neighbor resolution for the Link-local IP address. The workaround is to create the route manually e.g. for link-local IP fe80::200:80ff:fe00:0 on vlan10:
#bash sudo ip -6 route add local fe80::200:80ff:fe00:0/128 table local dev vlan10
The above workaround would have to be repeated on each switch restart. (1095686)
- Default STP mode is MSTP. If OpenConfig is used to configure STP, enabled-protocol should be specified if the default mode is required to be changed due to the lack of support for default of a leaf-list in YANG 1.0 (1135020)
- The SSH config option ObscureKeyStrokeTiming is disabled in EOS even though it's enabled by default in OpenSSH. (1195054)
- As of OpenSSH 9.6 most shell metacharacters have been banned when specifying a username or hostname via the command line. Please see the OpenSSH 9.6 release notes for more information.

As a workaround, users and hostnames containing these metacharacters can be specified using the ssh_config file instead. However, manual modifications to the EOS ssh_config file can be overwritten by the system. (1195136)

- The LoadAverage attribute is not available in Sysdb. To workaround this use the "uptime" command or get necessary information from the device's /proc/loadavg file. (1205981) (1)
- Starting in OpenSsh 9.8 validation behavior changed for a client connecting to a server using SSH Host Certificates. This can cause clients which only use SSH Host Certificates and not public key types to fail to connect. This will cause lines similar to "No matching CA found. Retry with plain key" and "host key <plain key type> not permitted by HostkeyAlgorithms" to be emitted when the client is run in `-vvv` mode.

To fix this either change the client setting for HostKeyAlgorithms to allow public key types or ensure the CA that signed the server's SSH Host Certificate is included in the known hosts file. (1238932)

- An NTP server can be configured with a minpoll value greater than its maxpoll value.

In this case, the effective value of minpoll is capped at the value of maxpoll. (1246325)

- If trap hosts are unreachable and traps are generated at high scale, snmpd may show high memory usage.

A workaround is to disable traps, and/or ensure trap host reachability. (1257541)

- During a system reboot, Linux kernel message "Failed to unmount /mnt/flash: Device or resource busy" may be displayed on the console. This message is not an error and does not indicate a problem with the device or the reboot process. It is a standard and expected diagnostic log from the kernel, a normal byproduct of how EOS boots from the flash memory. The reboot will proceed and complete successfully. (1262123)
- The ssh-keygen executable fails when the FIPS flag is set and no key type is specified.

As a workaround explicitly specify the key type when using ssh-keygen with the FIPS flag. (1271751)

- When running SSH in FIPS mode it is possible to login with an authorized 1024 bit RSA key.

As a workaround don't configure any 1024 bit RSA keys for SSH login when running SSH in FIPS mode. (1274136)

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

IPSEC

- We do not support NAT on Ipsec autoVPN in this release. (564431)
- When `ethernet untagged allowed` is configured under `ip security` config, TCP MSS clamping may not apply to packets encapsulated over a VXLANsec tunnel. (685484)
Series Affected: DCS-7280CR3 Series
- When a tunnel interface is shut and IPsec profile is removed from the tunnel interface, ipsec profile can't be deleted from the config.
The workaround is to delete the tunnel and delete the ipsec profile. (994463)
- For IKEv1 connections, the public key algorithm used to generate public and private keys must be the same as that used to generate the Certificate Authority's certificate. (1140300)
- NAT traversal and flow parallelization encapsulation udp are not supported with Policy Based VPN. (1156652)

Layer 1

- The MACSec outPktsEncrypted and outOctetsEncrypted counters are not supported with the MACSEC Proxy feature (281715)
- Clause 37 (BASE-X) auto-negotiation is not supported on SFP25 ports. (492271)
SKUs Affected: DCS-7010TX-48-F and DCS-7010TX-48-R
- Ports configured at PAM4 speeds (50G-1, 100G-2, 200G-4, 400G-8) do not support link-debounce intervals shorter than 2 seconds. (500674)
Series Affected: 7368 and DCS-7060X4 Series
SKUs Affected: DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32P4-F and DCS-7280CR3MK-32P4-R
- InPktsNoSCI counter increments for a MACsec frame with AN that is not in use by the receiver. InPktsSAnotInUse should have incremented instead. (532869)
SKUs Affected: 7800R3-48CQM-LC
- Interfaces with BCM82391 PHY do not support auto-negotiation. (537530)
SKUs Affected: DCS-7050CX3M-32S, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- QSFP to SFP adapters are not supported on front panel interfaces ethernet25/1-36/1. To work around this, use interfaces ethernet1/1-24/1 instead. (625046)
SKUs Affected: DCS-7280CR3-36S-F and DCS-7280CR3K-36S-F
- MACsec frames may be dropped when the sci setting in the "mac security profile" does not match the peer's sci setting. (664596)
SKUs Affected: 7388-8D, 7800R3-48CQM2-LC, 7800R3K-36DM-LC, DCS-7050CX4M-48D8, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R
- MACsec frames with SL less than 5-bytes are silently dropped on ingress (714497)
SKUs Affected: 7800R3-48CQM2-LC, 7800R3-48CQMS-LC, 7800R3K-36DM-LC, DCS-7170B-64C-F, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F,

DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R

- MACsec frames with SL less than 24-bytes are silently dropped on ingress (714597)
SKUs Affected: DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- Cable diagnostics test (TDR) on a BASE-T port may report inconclusive result if link partner is configured at forced 100M speed. (716208)
Series Affected: CCS-710P, CCS-720XP, CCS-750 and DCS-7050TX3 Series
- The MACsec inPktsNoTag counter includes the non-standard MACsec inPktsCtrl counts. inPktsCtrl (non-standard) counts packets that are parsed as control packets and hence forwarded to the uncontrolled port. For instance, EAPOL (MKPDU) and possibly LLDP or Pause/PFC if bypassed. (735450)
SKUs Affected: 7388-8D, 7800R3-48CQM2-LC, 7800R3-48CQMS-LC, 7800R3K-36DM-LC, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S, DCS-7280CR3MK-32P4S-F and DCS-7280CR3MK-32P4S-R
- The MACsec inPktsNoTag counter includes the non-standard MACsec inPktsCtrl counts. inPktsCtrl (non-standard) counts packets that are parsed as control packets and hence forwarded to the uncontrolled port. For instance, EAPOL (MKPDU) and possibly LLDP or Pause/PFC if bypassed. (735455)
SKUs Affected: DCS-7050CX3M-32S, DCS-7050CX3M-32S-F, DCS-7050CX3M-32S-R, DCS-7280SR3M-48YC8, DCS-7280SR3M-48YC8-F and DCS-7280SR3M-48YC8-R
- 50G-2 autonegotiation is not supported on interfaces with the CRT50216 PHY. (755659)
SKUs Affected: 7358-16C, 7368-16S, 7368-4D, 7368-4P, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3K-48CQ-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4A, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4A and DCS-7280CR3K-96
- On AWE-5510 ports 1-8 with SFP-10-T, the port may still link up when one side has auto-negotiation enabled and the other side does not. (793496)
SKUs Affected: AWE-5510
- 100G-4 autonegotiation with FEC disabled is not supported on interfaces with the CRT50216 PHY. (796531)
SKUs Affected: 7368-16C, 7368-16S, 7368-4D, 7368-4P, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3K-48CQ-LC, DCS-7060DX5-64S-F, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R,

DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R,
DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R and DCS-7280CR3K-96-F

- The "show interfaces transceiver hardware" CLI command may fail to run temporarily after reload. (854489) (1)
- On port 6 and 7 with 1000BASE-SX/1000BASE-LX optics inserted, "speed auto lgfull" is not supported. (942257)
SKUs Affected: AWE-7220RP-5TH-2S-F
- BASE-T management ports do not support operation in non-negotiated 100M/10M mode, despite accepting the associated configuration commands. (972448)
- The link may flap once before coming up through errdisable recovery. (998012)
SKUs Affected: AWE-7220RP-5TH-2S-F, CCS-720XP-24ZY4-F, CCS-720XP-96ZC2-F and DCS-7010TX-48
- SSU upgrade sequences starting from 4.21 or 4.22 to affected releases when 10G or 25G is configured on any interface in ports Et1-12, Et21-44, or Et53-64 will result in new interfaces appearing on upgrade. These interfaces will consume additional logical ports which may cause other interfaces passing traffic to go down due to a lack of resources.

Note that SSU upgrade sequences that stop at some other intermediate release(s) prior to upgrading to an affected release will still encounter the issue upon upgrade to the affected release. (1019049)

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

- Front panel interface counters will only count packets rather than octets when configured to retrieve counters from MetaMux, since MetaMux is only able to count packets. (1034409)
SKUs Affected: DCS-7130LBR-48S6QD, DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R
- 1000BASE-LX transceivers may fail auto-negotiation on Ethernet1/1 through Ethernet1/8.

A potential workaround is to use 1000BASE-SX transceivers. (1227433)

SKUs Affected: AWE-5510, AWE-5510-F, AWE-7250R-16S, AWE-7250R-16S-F, ZTX-7250S-16S and ZTX-7250S-16S-F

- 1000BASE-LX transceivers may fail auto-negotiation.

A potential workaround is to use 1000BASE-SX transceivers. (1227523)

SKUs Affected: NIM-4S

- Reading the eye diagram from an Ethernet port might briefly interrupt its link. (1265312)
SKUs Affected: DCS-7130LBR-48S6QD

Layer 2

- LACP PDU fast rate ("lacp rate fast" on an Arista switch) should not be configured on any port in an MLAG pair, or any port connected to an MLAG pair. (13950)
- Multicast VXLAN encapsulated packet is counted as a multicast packet on transmit, even though the inner packet which is sent out on the egress interface (after the VXLAN header is decapsulated) is a unicast packet. (55971)
- Static ARP inspection may not work on modular platforms after SSO switchover. The workaround is to disable and re-enable static ARP inspection after switchover. (244477)
- The command "lacp rate fast" is not supported on modular systems with stateful supervisor switchover (SSO) enabled. Neighboring devices should not have "lacp rate fast" configured on ports connected to the said system. (248121)
- Port Security Protect does not support trunk ports. If a port is secured to source MAC <A> in VLAN <X>, traffic from source MAC <A> in other VLANs of the trunk port will also be allowed. (344773)
- While performing SSO, a device's peers might not receive LLDP packets within default TTL (120 seconds) and those peer devices may not maintain LLDP status information for the device undergoing SSO. After SSO is complete, the peers will again have LLDP state available.

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

- An MLAG switch cannot be used as an EVPN multi-homing provider edge. When a switch has MLAG enabled, MAC addresses from ports with an ethernet segment configured are advertised through BGP as if they were single-homing, and local ethernet segments are not advertised through BGP. (397867)
- Even with fallback to unprotected traffic configured, traffic loss in the order of milliseconds may be seen during initialization (491232)
- LLDP always sends frames using the interface's physical MAC address as the source MAC even when it's used as an L3 interface. (501949)
- Upon programming a large number (~250) of VTEPs in a flood set at once, the L2Rib agent may restart. Workaround is to program them in batches, waiting a few seconds between each batch. (544051)
- 802.1X is supported on LAG member ports using EAPOL authentication. (544556)

- The SyncE wait-to-restore (WTR) timer is armed when administratively shutting / unshutting an Ethernet interface. This is inconsistent with disabling / enabling SyncE on the interface in which case the WTR timer is not armed.

To workaroud the WTR timer, issue 'no sync-e' followed by 'sync-e' for the given interface(s). (648153)

- When an EEC has redundant links traceable to the same reference (e.g. ITU G. 781 section 5.13.2 "'bundles' of ports"), and the quality level of the reference transitions from a higher quality to a lower quality (e.g. QL-PRC to QL-SSU-A), the EEC may briefly enter a holdover condition for a few seconds. (687326)
- Synchronous Ethernet is not supported for 10G BASE-T transceivers. (689493)
- Dot1x session move is not supported if the supplicant changes authentication methods during the move. (705466)
- Loop protection is not intended to be used along with Spanning Tree Protocol (STP) (745522)
- Idle-host idle timeout only applies to supplicants which have been authenticated. (747150)
- In a VXLAN MLAG anycast gateway deployment, if the dynamic ARP cache is cleared using CLI "clear arp-cache", the ARP cache is cleared on the MLAG peer where the CLI is applied and not on both peers.

The CLI needs to be applied to both MLAG peers for ARP cache consistency. (751712)

- LAG interface counters will have constant counts discrepancy compared to the aggregate interface counts of its member interfaces when there is a link flap event on LAG interface. (785073)
- In VXLAN routing deployment, when using EVPN Symmetric IRB to reach a destination host, the remote router IP address in the traceroute address may be arbitrarily set to any interface IP address in the VRF of the destination VTEP. (815530)
- If a custom MAC is configured without an active IP on the interface, even after the IP is configured, no compensatory GARP/GND will be sent. The workaround is manually sending GARP/GND or reconfig MAC. (826750)
- IP Phones doing untagged EAPOL authentication and sending tagged traffic may not work properly when 802.1X MAC based authentication is disabled. (841592)

- An idle MBA supplicant that is in DISCONNECTED-CACHED state may not be cleaned up when caching is disabled and link comes back up.
A workaround for this is to configure "dot1x timeout idle-host TIMEOUT seconds". This config would cleanup the idle supplicant after the configured timeout. (854465)
- LLDP packets from unauthenticated supplicant do not trigger MAC based authentication on 802.1X configured interface with "dot1x mac based authentication" enabled. (860140)
- The Lldp agent is always running, even when the protocol is disabled in the configuration. (875252)
- Disabling EAPOL using "dot1x eapol disable" command doesn't disable the fallback mechanism and its default timeout as configured using "dot1x eapol authentication failure fallback mba" configuration command.
Possible workarounds:
 1. Use the fallback CLI with timeout 0. eg "dot1x eapol authentication failure fallback mba timeout 0"
 2. Use MBA always option, ie "dot1x mac based authentication always", which supersedes "dot1x eapol authentication failure fallback mba" CLI. (918141)
- Traffic for a Dot1x authenticated supplicant received on a non-Dot1x interface will cause the MAC address learnt for that traffic on the Dot1x interface to not get cleared due to idle-timeout until the traffic stops. (938650)
- If a supplicant receives a dynamic VLAN from AAA server and this VLAN is outside the allowed dynamic VLAN range configured on the switch, the switch will proceed to successfully authenticate this supplicant. However, the supplicant traffic will be dropped, since the requested VLAN is not available on the switch. (948547)
- If Port Security is enabled on a 802.1X port, then the existing authenticated supplicants on the port will be logged off and they will be re-authenticated post logoff. (958954)
- If Port Security is enabled on a 802.1X port, then the existing authenticated supplicants on the port will be logged off and they will be re-authenticated post logoff. The delay between the logoff and the next successful re-auth (log in) might be as much as 1 minute. (959675)
- In a IPV6 VXLAN deployment, a system log entry of type IPV6_UNDERLAY_UNSUPPORTED may be emitted after the device is reloaded while having remote IPv6 VTEP configuration.

These messages can be ignored. The configured IPv6 addresses will be learned shortly after the IPV6_UNDERLAY_UNSUPPORTED syslog message appear. (961492)

- When there is a continuous ARP churn on L3 servers, the Dot1x agent will be busy processing all the updates and may send IPs from stale ARP entries on the L3 server in Framed-IP-Address RADIUS attributes (992621)
- If the reachability to L3 server from the L2 switch or the reachability to the L2 switch from the L3 server is over non-default VRFs on these devices, Dot1x agent on the L2 switch may not be able to send Framed IP Address for supplicants. (1008576)
- When a Dot1x supplicant changes to a new VLAN and acquires a new IP address, Dot1x may still send Accounting messages with the old IP in the Framed-IP-Address attribute till the new IP learnt in the new VLAN. (1009134)
- There is no support for MVRP with Rapid-PVST. It is not a supported configuration and may lead to undefined behaviour. (1035620)
- Caching of 802.1X devices credentials is not shared between MLAG peers. 802.1X devices will require credentials on both peers. (1053123)
- Dot1x authenticated MAC addresses do not persist after a port shutdown due to a violation, even though persistent port security is enabled. (1060769)
- When a large number of EAPOL logoff packets are received on a single interface, the switch may drop some packets and not process the logoff for the corresponding supplicants. (1079779)
- LLDP Neighbors are identified by a combination of the Chassis ID and Port ID. When the same Chassis ID is received with a new Port ID, we will add that as an neighbor and allow the old entry to timeout. (1095153)
- If 802.1X authentication feature is used with Arista-Interface-Profile VSA, then it may not work without any switchport configuration. The workaround is to have at least 1 switchport command either statically configured on the port or in the interface profile configuration. (1140078)
- Dot1x is not supported on layer 3 port-channel interfaces (1207542)
- Dot1x multi-host authenticated host mode is not supported on routed ports. (1216888)
- gNMI update/replace through OpenConfig may accept MACsec mka policy name with non alphanumeric characters. This may result in configuration not getting applied in some scenarios. The workaround is to use only alphanumeric characters for mka policy name. (1283873)

Layer 3

- ECMP is not supported for routes learned via RIP. (34773)

- DHCP relay is not supported when running virtual machines on EOS (101754)
- CLI does not detect TCAM exhaustion and automatically revert the change when a PBR policy is applied to an L3 interface in the "shutdown" state, since the actual hardware programming happens when the interface comes up. (122651)
- OSPFv3 support for multiple address families (RFC 5838) in EOS introduces support for IPv4 routing with OSPFv3. OSPFv3 for IPv4 AF in EOS requires configuring neighboring OSPFv3 IPv4 routers on the same link with primary IPv4 addresses in the same subnet. Adjacency is established in OSPFv3 IPv4 AF even if the neighbor's primary IPv4 address is in a different subnet but routes through the neighbor on a different subnet will however not be installed. This limitation may be removed in a future release. (140234)
- When BFD is unconfigured with Fhrp as the only client and/or BFD is administratively shut, both VRRP routers will become Master. This transition is temporary and Master-Backup relationship will be restored. (159647)
- BFD sessions established between Arista switches and Cisco peers over IPv6 link-local addresses may flap when BFD echo mode is enabled. (159803)
- vEOS does not support BFD echo sessions. (160770)
- If "bfd per-link rfc-7130" is configured on port-channel(s) with primary and secondary IPv4 addresses configured in the same subnet, and the "bfd neighbor" is configured with peer switch port-channel's secondary IP address and there is a BFD session using primary IP address, then the port-channel might flap continuously.

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

- BFD echo functionality is not available on L2 Port-Channels with BFD RFC7130 mode enabled. (190418)
- If a port channel is configured with "bfd echo" and "bfd per-link rfc-7130" and is used for OSPFv3 or PIM, BFD sessions will not come up with echo functionality. BFD with SSO for OSPFv3 or PIM is not supported. (190997)
- An LACP port channel configured with BFD per-link and echo may experience BFD session flap on existing member links if new member link(s) are added to the port channel while the BFD peer device undergoes Stateful Switchover (SSO). (255042)
- Eos does not support programming IPv6 ECMP paths to kernel (258387)

- The DSCP value of BGP control plane packets set by the command 'bgp transport qos dscp <value>' will be overwritten when the platform 'qos rewrite dscp' is enabled and the outgoing interface is a SVI based on traffic-class to dscp map present on the platform. (269764)
- In the output of the "show ipv6 dhcp relay counters" or "show ip dhcp relay counters" command, interface-specific DHCP relay counters might not sum up to "All Req" and "All Resp" counters. (278786)
- On interfaces with traffic-engineering enabled, configuring secondary IP addresses may exclude these interfaces from being considered for constrained based traffic engineering RSVP paths. (279803)
- Tunnel interfaces flap when configuring ttl, tos and path mtu discovery on them. (281464)
- IPv6 addresses are not supported as tunnel source and destination. (283912)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289217)
- High MPLS route scale or high MPLS route churn may cause the platform agent to restart unexpectedly. (289224)
- When BFD is configured with per-link mode rfc-7130 for an IPv6 link-local address on a Port-Channel and the Port-Channel transitions to the Down state, the Port-channel may not transition back to the UP state.

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

- On Macro-Segmentation Service (MSS) deployment with L3 firewalls, a flow entry that fail to get programmed due to lack of TCAM space, is not programmed when TCAM space becomes available at a later time. The workaround is to "shut" and "no shut" on the direct flow config. (372232)
Series Affected: DCS-7050X3, DCS-7060X and DCS-7060X2 Series
- The update wait-install feature may not work as expected if there are routes which fail programming due to hardware resource exhaustion. The routes which failed programming may still get advertised out to peers. (372765)
- DHCP packets will not be processed by EOS DHCP Server when DHCP Relay is relaying packet from one SVI where DHCP Relay is configured to another SVI where DHCP Server is configured. (389017)

- DHCP clients will fail to obtain IPv6 leases when "ipv6 dhcp relay destination" is set to IPv6 multicast address. (426796)
- Layer 3 protocols incorrectly see Pseudowire local connectors as layer 3 interfaces. As a result, reachability to interface's attached IP networks may be advertised, depending on the routing protocol configuration, irrespective of pseudowire connector's operational state. (429818)
- Nexthop group tunnel counters does not work, when tunnel counters is enabled using 'hardware counter feature mpls tunnel' command. (474078)
- VXLAN features are not compatible the 'ip hardware fib next-hop sharing' configuration. Configuring VXLAN routes may result in identical FECs not being shared and potential misforwarding after changes to FECs are made. (477471)
- The config 'ip hardware fib next-hop multi-path maximum' does not restrict the size of installed next-hop group ECMP FECs (498484)
- CLI configuration of a static route with nexthop pointing to self ip address is not rejected. (514894)
- NAT dynamic profiles are not supported on multi-chip and modular system. (525213)
Series Affected: 7300X3, CCS-720XP and DCS-7300X Series
- BFD over Vxlan does not work for IPv6 underlay. (526897)
- Rsvp autobandwidth feature does not work for single hop tunnels. The workaround is to configure Rsvp explicit-null on the egress router for single hop tunnels. (571177)
- VARP VTEP IP configuration is not supported on AP aggregation VTEPs

It is recommended that either EVPN VXLAN symmetric or asymmetric IRB be used for distributed VXLAN routing in campus environments. (571517)

- IS-IS Extended administrative group (EAG) sub-TLV is not supported. If a system in the network advertises EAG, Arista router will ignore the sub-TLV. (628807)
- The `show ip hardware fib ecmp resilience` CLI may sometimes incorrectly report an unordered resilient ECMP as ordered. (685318)
- If a preferred IPv6 source address for an IPv6 kernel route is obtained from a local Port-Channel interface that is used in a "software forwarding hardware offload route local interface INTERFACE" command, and the interface has a IPv6 BFD neighbor configured, and the Port-Channel interface is subsequently shut down, the preferred IPv6 source address of the kernel

route is not updated as expected.

The workaround is to remove the "bfd ipv6 neighbor" command from the Port-Channel interface and re-apply the command. (698716)

- If a "software forwarding hardware offload route rcf FUNCTION" command is configured under "router kernel/address-family" sub-mode and FUNCTION returns True for a connected route, then kernel traffic to that prefix is forwarded in hardware as expected. But if the connected route later goes away for some reason (e.g., the interface is reconfigured) and is replaced in the FIB by a route for the same prefix but learned from a routing protocol peer, the kernel route is programmed for software forwarding instead of hardware forwarding.

Restarting the KernelFib agent should cause the affected kernel route to be reprogrammed such that forwarding of kernel traffic to the prefix is offloaded to the hardware. (701746)

- CLI command show ip route may incorrectly display an unprogrammed route as programmed. (731198)
- BFD running over L3 interfaces can reduce the time it takes to shrink an ECMP when one of the ECMP member links fail. If all the ECMP paths use the same L3 interface (such as when using SVI interfaces) then the time to shrink the ECMP may take longer. (746361)
- The `show ip hardware fib ecmp resilience` command may display `VRF <name> not found` when given a non-default VRF. (764757)
- 'show ip hardware fib ecmp resilience' and 'show ip hardware fib ecmp resilience summary' commands may show status of unprogrammed resilient routes that have already been removed from FIB. (798614)
- IPFIX and sFlow BGP nexthop value cannot be accurately determined for ECMP flows going over the same egress interface that are encapsulated with implicit-null labels. There is no workaround for this issue (816528)
- When "ip virtual-router mac-address mlag-peer" is enabled in the MLAG environment, BFD is not supported on the MLAG peer interface. (818124)
- "MIXED" prefix-sets are not supported in OpenConfig in policy constructs. The workaround is to use separate prefix-sets for IPv4 and IPv6 prefixes. (858818)
- Setting accumulated IGP (AIGP) metric of "-0" using the "set aigp-metric" route-map clause is not supported and results in a traceback in the EOS CLI.

Workaround is to use "0" (887974)

- The `show ip fib hardware ecmp resilience` command may continue to display unprogrammed routes with reason "Route not programmed". These are routes that were previously resilient, but were subsequently deleted/unprogrammed. (903779)
- The next hop in the extended router data of an sFlow datagram will not be populated if the destination IP address of the sampled packet is of a different address family than the next hop IP address (939695)
- 'show ip dhcp relay counters' and 'show ipv6 dhcp relay counters' count both IPv4 and IPv6 DHCP packets (952513)
- CLB does not handle flows with the same destination IP and destination queue-pair but different source IP addresses (995904)
- Traffic loss may occur if the OSPF grace period exceeds 600 seconds (1039979)
- An OSPFv3 NSSA ABR may incorrectly translate type 7 NSSA LSAs into external LSAs (1049021)
- When the same prefixes are added from static EVPN as well as from BGP sources, the FIB route count may encounter a double counting issue. (1063801)
- OSPF LDP Sync does not take effect on an interface when auto-cost reference-bandwidth is configured for an OSPF instance and `ip ospf cost ...` is NOT configured for an interface. OSPF applies the auto-cost metric instead of max metric. To workaround this behavior, explicitly configure a cost on the interface with `ip ospf cost`. (1089392)
- If the same static EVPN route is configured through both the EosSdk and the CLI, it can result in a duplicate count in the FIB. There is no available workaround. (1101151)
- Aggressive BFD interval configuration (timers less than 300 ms with a multiplier of 3) are not supported on T3.X3 switches. (1134966)
 SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DF-48Y-2F, CCS-720DF-48Y-DC-2F, CCS-720DF-48Y-DC-F, CCS-720DF-48Y-F, CCS-720DF-48Y-M-S-2F, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-24S-M-S-2F, CCS-720DP-24ZS-2F, CCS-720DP-24ZS-F, CCS-720DP-48S-2F, CCS-720DP-48S-F, CCS-720DP-48S-M-S-2F, CCS-720DP-48ZS-2F, CCS-720DP-48ZS-F, CCS-720DT-24S-2F, CCS-720DT-24S-2R, CCS-720DT-24S-F, CCS-720DT-24S-M-S-2F, CCS-720DT-48S-2F, CCS-720DT-48S-2R, CCS-720XP-24Y6-F, CCS-720XP-24Y6-R, CCS-720XP-24ZY4-F, CCS-720XP-48TXH-2C-S-F, CCS-720XP-48Y6-F, CCS-720XP-48Y6-R, CCS-720XP-48ZC2-F, CCS-722XPM-48Y4-F and CCS-722XPM-48ZY8-F
- BGP-LS may continue to produce IS-IS LSDB information when the IS-IS agent is shutdown using the "agent Isis shutdown" command.

Instead of shutting the IS-IS agent, using 'shutdown' under 'route isis

<instance>' can be used to shutdown IS-IS which also removes the LSDB information for BGP-LS (1177930)

- When there are multiple OSPF instances configured and a routemap is configured for redistribution in more than one instance. If include-leaked is configured in one instance for redistribution, the leaked routes will also be available in the other instance. (1182487)
- Loopback interfaces are the only interface type that should be used as a lending interface. (1266259)

Multi-agent

- If a BGP peer is configured for BFD fallover, the BFD session is not cleaned up even after the BGP peer is unconfigured. (217199)
- Traffic destined to an L3EVPN route whose underlay has ECMP BGP-LU routes, may experience temporary traffic loss when the underlay goes from n way to n-1 way or lower. The duration of loss can vary depending on the scale of BGP-LU routes or L3EVPN routes. (251692)
- When the multi-agent routing protocol model is configured, the maximum number of next hops for an OSPF ECMP route is 128. (274888)
- With very a very low multi-hop BFD timer configured, a BFD session may flap when a path (belonging to an ECMP set), over which the BFD session is established, goes down. (287571)
- Route-map specified as match-map for dynamic prefix-list is evaluated only against BGP routes. (345444)
- The IpRib agent runs out of addressable memory when a large number of routes are leaked to a large number of VRFs. (346575)
- Routes which are imported into a target VRF using a "leak routes source-vrf" policy cannot be matched on the basis of the source-protocol from the source-VRF as part of the "rib fib policy" (403739)
- In multi-agent mode, extended Sflow BGP data is not supported for BGP LU routes. (411762)
- When using 64-bit versions of EOS, Bgp, BgpCliHelper and Bmp agents may report a virtual memory size up to 50 GB larger than actual memory used. The actual memory usage is better represented by the sum of "RES" and "SHR" from "show processes top". (427364)
- When BGP router has advertised a large scale of routes (e.g. several million) and it reloads, there may be traffic drops while it is coming back up. This happens when the reloading router comes back up before the directly connected peering routers (which are receiving routes via an

intermediate Route Reflector) get a chance to process millions of withdrawals and delete all the routes from FIB. (458127)

- An AS_PATH modification operation in an RCF function applied on inbound for an IBGP peer is ignored. (474006)
- IP ACLs are not supported for control plane route filtering for VRF leaking (536329)
- Match dynamic prefix-list clauses in route maps are only supported for BGP inbound and outbound route maps. (564552)
- CVP will not monitor multiple BGP peers that are configured using the same peering IP address. (588954)
- In multi-agent, the aggregate link bandwidth calculation by BGP does not honour the "maximum-paths" configuration. It is calculated as the sum of the link bandwidth of all available best paths for a prefix, rather than the sum of the link bandwidth of the paths actually selected for forwarding. This may result in advertisement of incorrect higher bandwidth to BGP peers and may impact accurate functioning of UCMP. (601570)
- BMP, SNMP and EosSdk will not report multiple BGP peers configured with the same BGP peering address. (602657)
- When default-originate is configured, default route will not be advertised to a neighbor if a prefix-list configured directly for this neighbor denies the default route. (648587)
- When multiple single-hop EBGp sessions are configured using point-to-point physical interface IP addresses between the same pair of switches, some of the BGP neighborship may get established using the other peer's IP address. When any BGP neighborship comes up using such wrong IP address from cross-link, the BGP neighbor configured on the other link will have connection establishment failure with Notification error: 'connection collision resolution'.

To workaroud this issue, either use 'update-source' option for the neighbor or explicitly enable 'no neighbor <addr> route-to-peer' for all the single-hop EBGp peers using point-to-point physical interface IP address. (686390)

- BGP session tracker may not detect quick BGP peer flaps and may not error disable interfaces (687998)
- For the RT membership AFI/SAFI, RT membership routes advertised by a given peer are also reflected back to the same peer. i.e. there is no echo suppression of the RT membership NLRI. (717330)

- IPv4 static route using both the interface and associated IPv6 address as nexthop requires both the IPv4 and the IPv6 address to be present on the the interface before configuring this static route. (778576)
- BGP-LU routes directly redistributed into IS-IS using command 'tunnel source-protocol bgp ipv4 labeled-unicast' under 'router isis' mode, uses MED as metric instead of AIGP.

To workaround this issue, IS-IS metric value can be derived from AIGP metric using RCF policy. Note that the 'aigp.metric' attribute value, used by RCF at redistribution point of application, is same as AIGP metric path attribute value received from the peer and doesn't account for the next-hop's IGP metric cost. (791558)

- At a high route and peer scale, BGP may take a long time to report ROA validation counters. (884609)
- Routes where the destination prefix covers the next hop of the route are not installed in RIB and consequently in the FIB. (898293)
- Debugging RCF policy with the "show bgp debug policy" command may not work with neighbor additional-paths receive. (921294)
- BGP link local peers are not supported through OpenConfig. The interface discriminator that is displayed in the EOS CLI (such as %Et1) cannot be conveyed in the openconfig model in either config or status. Upon querying the state of existing BGP peers through a gNMI GET request, unexpected values will be present in the 3rd and 4th octets of the IPv6 address. (928239)
- If a VPLS-BGP PE advertises a VPLS-BGP route as part of a VPLS pseudowire group, it will continue to advertise the VPLS-BGP route even if the VE block in the route no longer covers the VE-ID of any PE in the VPLS-BGP group that the PE is participating in.
To withdraw the VPLS-BGP route, the advertising VPLS-BGP PE must be deconfigured from the VPLS pseudowire group and then reconfigured. (975517)
- The "ping mpls pseudowire vpls" CLI command can currently only be used to ping VPLS pseudowires signaled using LDP, not VPLS pseudowires signaled using BGP. (976255)
- VPLS pseudowire signaled using BGP will ignore any "pseudowire-type" config when negotiating a VPLS pseudowire and will instead select a pseudowire type (of either 'raw' or 'tagged') based on the VLAN configuration of the VPLS instance the VPLS pseudowire is inside of. (976887)
- When next hop is updated using either outbound route-map or outbound RCF policy for the labeled-unicast peer, the advertised LU label is not updated correctly.

To workaround this issue, configure 'neighbor <peer-address> next-hop-self' for such labeled-unicast peer. (1001447)

- "next-hop resolution route Igp-nexthop-cost protocol bgp" feature is not supported when the BGP next hop is resolved over tunnels. (1003571)
- If a link bandwidth extended community is being introduced outbound, route-map or RCF policy may not match on the link bandwidth value being advertised over the wire.

As a workaround, match on the link bandwidth value using inbound policy on the received device. (1055916)

- In multi-agent mode, the conditional default route originate feature is limited to specific types of routes such as BGP, IS-IS, RIP, static and connected routes. For other sources of routes such kernel routes, DHCP routes etc. the conditional default route originate functionality does not work in multi-agent mode. (1098356)
- "show ip bgp" may not show all the paths as ECMP contributors when the next hops resolve over the same terminal/outgoing interface. (1114712)
- When BGP RT membership is configured locally on a peer (but not configured on the remote peer), default RT membership routes will appear in the 'show bgp neighbor ... received-routes' output even though they were not received from the remote peer.
This is just a cosmetic issue and shouldn't affect any functionality. (1220400)
- Link bandwidth regeneration does not support VPN and VPN route resolution over MPLS GUE tunnel (1232890)
- When an empty argument is provided to EosSdkRpc BgpLinkWeightMgrService service's bulk_link_weight_set() RPC, the error code returned is INVALID_ARGUMENT (3) instead of OUT_OF_RANGE (12) (1268485) (1)

Rib

- PBR recursive resolution for nexthops is not supported for packets routed in software (like packets with IP options) (82077)
- No support for IPv6 labeled Unicast capability when BGP neighbor is established using IPv4 transport. (208402)
- Match dynamic prefix-list clauses in route maps are only supported for BGP inbound and outbound route maps. (261865)
- Dynamic prefix-list match clauses are not supported in BGP neighbor inbound route maps. (269663)

- The BGP configuration commands, "bgp next-hop-unchanged" and "neighbor next-hop-unchanged", are not supported in ribd mode (single-agent mode).

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

- Ribd agent may restart while running with Internet scale routes and with a large number of (50+) Rib-Out groups. Such a high scale is not supported with single-agent mode.

To workaround this issue, either reduce the number of Rib-Out groups by checking the output of 'show bgp update-groups' or upgrade to multi-agent mode which can support very high scale. (558044)

- Static routes with administrative distance of 255 are installed in the RIB and eligible for redistribution into other protocols with "redistribution config". The routes are also installed in the FIB for programming in hardware.

To deny FIB programming of routes, the RIB route control configuration policy in the form of "rib ipv4|ipv6 fib policy" can be used.

To prevent RIB redistribution of routes, the routes can be filtered out as part of the "redistribution" policy into the target protocol. (663662)

- There may be a behavior difference in OSPF between single agent and multi agent mode when both OSPF and BGP are configured in a non-default VRF. In single agent mode, a route that would come from an external type 5 LSA might be dropped where in multi agent mode, the route would not be dropped.

The workaround to prevent loops is to follow RFC4577 (use of the DN-bit) (774087)

MetaMux

- MuxIn Interfaces from non-EB profile, do not get enabled when sourcing directly from different MetaMux instance.

Atleast one another MuxIn Interface in that non-EB profile, must be sourced from either front panel interface or, switch interface from J2C, to make it enabled (982810) (1)

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

MetaWatch

- Disabling the time-sync functionality will not stop MetaWatch from passing traffic or synchronizing time.

Workaround for stopping traffic is to disable MetaWatch, and to stop the synchronization action is to switch the time-sync reference to 'free-running'. (743245)

- If the PTP Grand Master's time fluctuates and/or is not of good quality, the MakoTimesync agent would be unable to synchronize to it, staying in fast converge mode. (782482)
- If the PTP interface is enabled, but the time-sync reference is set to a value other than PTP, the device will continue tracking the upstream PTP network but not sync to it. In particular, the output of `show metawatch ptp` will still report offset from GM (if any).

If using a time-sync reference other than PTP, disable the PTP interface with `no ptp enable`. (791377)

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

- In order to use the 40G-capable profiles, the metawatch-3.5 SWIX (or greater) must be used. (794428)
- If the PTP Grand Master clock is in the UTC timescale, the PTP interface might enter a faulty state when either of its associated internal CPU interface or Switch interface flap.

It is recommended to always use a PTP Grand Master clock in the TAI timescale to avoid such potential instability issues. (796022)

- When SwitchApp is running on one application FPGA and MetaWatch is enabled on another application FPGA, for the time MetaWatch is synchronising SwitchApp will observe fluctuations in its offset from a PTP master within range of +-800ppb. This behaviour is transient and ends once MetaWatch is synchronised, it can be worked around by starting MetaWatch before SwitchApp. (840573)

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

- MetaWatch will not start correctly after a reboot if it is configured with instance ID 1.

Use an instance ID other than 1 for the MetaWatch application. (925244)

Series Affected: DCS-7135 Series

- When using deduplication, any inserted VLAN tag (in the original packet) that is not removed prior to hashing will disable the IPv4 TTL and header masking. (935469)
- Combination of "time-sync reference ptp" and "ptp system clock interface ptpX/Y" commands is not supported. (955934)
Series Affected: DCS-7130 Series
- PTP on MetaWatch may intermittently lose synchronization when using a Gigabit Ethernet transceiver. (984836)
Series Affected: DCS-7132 and DCS-7135 Series
- An Arista FDK application with Tscore cannot be configured alongside MetaWatch, even if the configurations are disabled.

To use an Arista FDK application with Tscore, all MetaWatch instance configurations must be removed (not just disabled).

To use MetaWatch, any Arista FDK applications with Tscore must be disabled. (1089489) (1)

- PTP synchronisation quality will be degraded if a device with an atomic clock module is synchronised via a management interface.

Devices with an atomic clock module should use a more stable time source, such as PPS. (1210702)

SKUs Affected: DCS-7130-48LA, DCS-7130-48LAS, DCS-7130-48LBA, DCS-7130-48LBAS, DCS-7130-96LA, DCS-7130-96LAS, DCS-7130-96LBA and DCS-7130-96LBAS

- It is possible for the follower instance to synchronize early during initial synchronization while the offset from master is still significant. This can result in extended initial synchronization time. (1220024)
SKUs Affected: DCS-7130LBR-48S6QD, DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R

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

MLAG

- Spanning-tree mode backup does not work in an MLAG configuration. (15151)
- The port-channel configured on each MLAG switch with the same MLAG ID must also have the same channel-group ID. (22890)
- Configuring a device connected to an MLAG with a round-robin LAG distribution algorithm is not supported if the device is going to participate in IGMP. (28370)

- On a scaled MLAG setup, Lag+LacpAgent agent may restart unexpectedly if MLAG is reinitialized due to configuration changes. (116125)
- On an MLAG system configured with dual primary detection and with default control plane ACL enforced, MLAG will not form if the management interfaces of the MLAG peers are connected through extra hops (e.g., on different subnets), due to the default control plane ACL dropping MLAG control packets with TTL < 255.

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

- TCAM profile switch in a MLAG switch may result in some packet duplication for broadcast or multicast packets (341353)
- Configuring private VLAN with fast MAC redirection on a MLAG setup may cause traffic loss in the private VLAN. Fast MAC redirection also gets enabled on configuring "platform sand lag hardware-only".
The workaround is to disable fast MAC redirection using the CLI command "inactive mac-address destination peer-link direct" under MLAG configuration. (683982)

Series Affected: DCS-7500R Series

- With MLAG egress filter interlock configured, duplicate BUM packets may be seen on MLAG interfaces along with "%MLAG-4-PEER_ACL_STAGE_TIMEOUT" syslog when the MLAG peer-link state changes.

If duplicates are seen while exiting MLAG dual-primary detection, the workaround is to configure a large enough MLAG dual-primary recovery delay such that PEER_ACL_STAGE_TIMEOUT syslogs are not observed. If duplicates are seen when reloading a MLAG peer, ensure the MLAG reload delay is configured properly. (724152)

- On an MLAG device VRRP sessions configured on VLAN interfaces may briefly transition to the master state after device reboot even though the VLAN is not associated with any active L2 ports. The VRRP session will become Stopped once the MLAG reload delay timer has expired. Beyond unexpected log messages there is no impact of device operation. (799723)
- When MLAG subinterfaces are configured and a MLAG interface flaps, there may be traffic loss. Fast MAC redirection is not supported on MLAG subinterfaces. (1018070)

MPLS

- The MPLS agent may continuously restart on a router with high BGP LU Tunnel scale. (377939)

- BGP LU Tx MPLS routes corresponding to BGP LU advertisements with multiple labels are not supported. (395029)
- The control word configuration does not take effect for one-hop pseudowire tunnels (406396)
- When RSVP is configured with split-tunnels, the `show traffic-engineering rsvp tunnel` command may display a last sampled bandwidth value for the tunnel that is not exactly equal to the sum of the last sampled bandwidth value of all sub-tunnels. This can occur if the command is used while the values are actively changing as it is possible that some values in the output have not yet been updated while others have. (625866)
- Dynamic IGP metric is not supported for non /32 RSVP tunnel destination IP addresses. (785652)
- In the "show mpls rsvp counters" output, received node hello messages will increment the unknown neighbor error count. (894064)

MultiAccess

- When MultiAccess restarts, traffic will temporarily be dropped when each interface has its ACL re-asserted. (983414)

Multicast

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

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

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

on an interface and feature becomes disabled, flap the pim sparse mode configuration on non-DR on that interface to attempt to get out assert winner state. (401672)

- Multicast routes are not created when a source is transmitting only fragmented multicast packets.
Switch to using SFE for multicast software forwarding. (605326)
- In certain EVPN Multicast deployments, the PEG designated router (DR) might send PIM Registers to the RP for sources located in the external PIM domain. Once the PEG DR receives register stops, it ceases to forward registers. This behavior can initially result in a few packets being duplicated. (630762)
- A device with 8GB RAM might be able to support a maximum of 20 VRFs or lesser depending on the memory requirements of other features configured on the device. A workaround is to reduce the VRF scale or use a switch with higher RAM for such a scale requirement. (707709)
- Multicast scale is limited on devices running with 4GB memory. Running `software-forwarding sfe` on these platforms can result in BessMgr agent crashing. (786687)

NAT

- SNAT will not be performed if the original source IP is not fully resolved, and DNAT will not be performed if the translated destination IP address is not fully resolved. For routed ports, a static or dynamic ARP entry must be present for the relevant IP address. For SVIs, both an ARP entry and a MAC table entry for the MAC address associated with the IP address must be present. (756211)
Series Affected: DCS-7130 Series
- Changing the forwarding profile between nat-vxlan and nat may cause the Nat agent to unexpectedly restart. (792534)
Series Affected: DCS-7170 Series
- When an interface configured with static NAT translations is taken out of the shutdown state, untranslated traffic will flow briefly. (805141)
Series Affected: DCS-7130 Series
- If static source NAT is configured on the OUTSIDE interface of an INSIDE/OUTSIDE pair of interfaces, configuring static destination NAT on the INSIDE interface will cause the static source NAT translation to fail. (821248)
SKUs Affected: DCS-7130-16G3-F (HwRev:C1), DCS-7130-16G3-R (HwRev:C1), DCS-7130-48G3-F (HwRev:B1), DCS-7130-48G3-F (HwRev:B2), DCS-7130-48G3-R (HwRev:B1), DCS-7130-48G3-R (HwRev:B2), DCS-7130-48L-F (HwRev:A2), DCS-7130-48L-F (HwRev:A3), DCS-7130-48L-R (HwRev:A2), DCS-7130-48L-R

(HwRev:A3), DCS-7130-48LB-F (HwRev:A3), DCS-7130-48LB-R (HwRev:A2) and DCS-7130-48LB-R (HwRev:A3)

- NAT synchronization is not supported together with NAT overload. (862459)
- It is expected to see some connections exceed over the configured connection limit when connection limit is used with NAT Synchronzation config. With higher scale, the number of connections exceeding the connection limit will be much lower. It is recommend that NAT Synchronization is configured over MLAG. (1002400)
- It is expected to see some connections exceed over the configured connection limit when connection limit is used with NAT Synchronization config. With a higher scale in terms of the number of connections and with larger value of connection limit, the number of connections exceeding the connection limit will be much lower. It is recommended that NAT Synchronization is configured over MLAG. (1002486)
- Dynamic source NAT translations in different VRFs sharing the same ACL is not supported, when the ingressing traffic flows in both VRFs have the same source and destination address and port. (1105562) (1)
Series Affected: DCS-7132 Series
- In a configuration, two or more dynamic NAT translations are defined with ACLs that can match the same packet (due to either using the same ACL rule or the ACL rules overlap). If those dynamic NAT translations are configured and applied in different NAT profiles, it is possible that some new dynamic NAT translations will not be installed in the hardware forwarding engine. (1237023)
- Twice NAT translations will only match and translate ingress traffic on ports that have reachability for the reverse path's destination. (1254435) (1)
- Configuring the ACL on a dynamic source NAT and dynamic twice-NAT rule in the same NAT profile to be the same, or to be ACLs that overlap allowing packets to hit on either rule is not supported. (1266441)

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

Switch Aggregation Group

- The LoadAverage attribute is not available in Sysdb. To workaround this use the "uptime" command or get necessary information from the device's /proc/loadavg file. (1205981) (1)

SwitchApp

- IGMP snooping filters are not supported when SwitchApp over forwards IGMP reports. Reports are over forwarded if the "igmp flooding" CLI is enabled through the FPGA instance CLI, or if the FPGA profile does not support correct forwarding of IGMP reports. (601917)
- When a MLAG with DCS-7130 LB series devices is running SwitchApp application with L3 mid-latency profile, because of the pipeline bandwidth constraints, if the peer-links are in the same pipe as other ports and the receive traffic rate across all ports is high, MLAG control traffic across peers may be dropped causing MLAG relationship to fail.

The recommendation when using this profile is to have the peer-link ports from a dedicated pipeline.

The port to pipeline mapping is available via "show fpga switch mapping" command. (715712)

- Multicast boundaries do not filter unicast PIM encapsulated packets sent to the RP. This may result in unicast PIM traffic continuously reaching CPU in certain network topologies. (718224)
- When running the SwitchApp application as a PTP grandmaster, it will distribute its internal free running counter instead of system time, which corresponds to the time the SwitchApp has been started since 1 January 1970. (736982)
- When using SwitchApp with routing and MLAG, if link failure causes more than 64 ARP entries to be needed on the peer link then traffic may be dropped.
Design the network with care to ensure that this limitation does not occur for the relevant failure scenarios. (765647)
- MakoSharedResources agent may restart hitfully if MakoL3Unicast agent performs a hitful restart due to crashing during a hitless restart. The system will recover to normal operation without intervention. (800466)
- SwitchApp platforms are limited to 64 nexthops per L2 port. This includes the inter MLAG peer link. As a result MLAG systems where more than 64 hosts are reachable via MLAG ports may hit this limitation during failure scenarios where traffic for more than 64 hosts need to cross the peer link resulting in traffic loss. It is therefore recommended to design configurations to ensure that this limitation is not exceeded in the event of failures. (811102)

- When the hardware resources become full, new NAT rules will not be programmed in hardware, and the system will log:

```
`%NAT-3-HW_RESOURCE_FULL: Hardware resources are insufficient to program all NAT rules`
```

Those rules will remain pending, even if hardware resources subsequently become available.

The workaround is to remove any pending rules (``show ip nat translation pending``) and re-add them to use the available hardware resources. (848105)

- SwitchApp platforms may fail to insert ECMP prefixes in hardware (due to running out of ECMP group resources) even though other prefixes with the same set of nexthops are inserted. This occurs because the routes do not get assigned the same FecId due to differences in route attributes in the RIB.

Use `'ip hardware fib next-hop sharing '` or `'neighbor <> next-hop-self'` commands to ensure prefixes share the same FecId. (910657)

- In a case when there are both a dynamically source NAT-ted and a routed VLAN sharing the same trunk interface on egress and the source IP address for the routed traffic matches the ACL for the dynamic NAT rule, the routed packets are being trapped by the CPU and therefore highly rate-limited. The only remedy is to keep the ACL as specific as possible, so the ACL overlapping doesn't happen. (950797)
- Interface link flaps that result in multicast route re-programming may result in software forwarding of multicast packets. Software forwarded packets are not subject to multicast NAT configurations. (958816)
- A router running SwitchApp may emit the following error `"MakoL3Unicast: %ROUTING-3-FEC_RESOURCE_FULL: Not all FECs are programmed in hardware"` followed soon after by `"MakoL3Unicast: %ROUTING-3-FEC_RESOURCE_NORMAL: All FECs are programmed in hardware."` This is expected in scenarios where a number of IP next hops are learned in the routing table before their MAC addresses have been learned in the MAC address-table.

There is no workaround required - the `".. _RESOURCE_NORMAL"` message indicates that the condition has cleared and the router is now forwarding correctly to all destinations. (1009531)

- Routes can fail to be programmed in hardware and `HW_RESOURCE_FULL` errors can occur even though LPM table utilizations are as low as 25% of capacity. The limit will vary with the prefixes in use and may also vary between reloads of the switch.

Workaround: Reducing the number of prefixes at the relevant prefix lengths until the utilization below 25% will resolve the problem. (1049134)

- On an interface where both storm control threshold(s) and ingress mirroring are configured, traffic mirrored to the destination may be rate-limited by storm control.

To remedy this, we advise configuring storm-control only on interfaces that are not configured as mirroring sources. (1092112)

- On SwitchApp with NAT translations configured, if the number of active ports in the ingress or egress VLAN(s) becomes zero, the NAT translation will be marked as in progress. Any hardware resources allocated while the translation was active, will remain active and programmed, although no translation will be possible. This is considered a transient condition, and is intended to support network reconfigurations while reducing any packet loss due to hardware reprogramming.

The resources can be recovered by deleting and re-adding the NAT translation; resource will only be reallocated once the ingress and egress VLAN(s) become populated. (1101466)

- Storm control is not supported on control plane interfaces. (1108395)
- When a transceiver is plugged into an Ethernet interface that does not support any speeds that the transceiver supports, the speed will default to 25G for that interface and become errdisabled. (1232542)
Series Affected: DCS-7132 Series
- Modifying an ACL that includes rules with a subnet that is a sub-network of another ACL's rules could lead to a very brief traffic leak on the second ACL. (1239360)
- Twice NAT translations will only match and translate ingress traffic on ports that have reachability for the reverse path's destination. (1254435) (1)

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

DCS-7130 and DCS-7170 Series

- If an ingress packet contains more than one 802.1Q tag and the egress port is a trunk port, the egressing packet may get corrupted. (180436)
Series Affected: DCS-7170 Series
- For frames that are routed and VXLAN encapsulated, underlay MTU is not used for MTU enforcement. Therefore the packets will be routed, encapsulated and transmitted without fragmentation. (254767)

- Changing port speed from 10G/25G/50G to 40G/100G will cause the forwarding agent (BfnSliceAgent) to restart. This can result in a brief traffic outage in the order of 100ms on all ports. The link state of other ports will not change. (256519)
- 10/25/50 Gbps interfaces only support MTU up to 7Kbytes. (258823)
- After applying or removing shaping configuration on an L2 sub-interface, the L2 sub-interface must be flapped for the configuration to take effect. (281312)
Series Affected: DCS-7170 Series
- After setting CoPP rate = 0, a few packets (<10) might be let into the CPU, but everything will be dropped afterwards. (350676)
- The number of L2 subinterfaces which can be supported on a parent interface is 63.
On a given parent interface, there cannot be more than one L2 subinterface with the same forwarding vlan. (357946)
- Full mroute scale may not be achievable under certain scenarios due to hash collision for the multicast route table. (380545)
- When egress NAT and MLAG are both used, 'show ip mroute' may show NAT'ed addresses in addition to the receiving mroute addresses. (394571)
- firewall profile on DCS-7170 platform does not support IP multicast (401134)
- Accelerated Software Upgrade (ASU) is not supported in combination with the packet-filter profile. The hitless reload will time out and fall back to a normal reload. (473985)
Series Affected: DCS-7170 Series
- On the DCS-7170 series, ASU is not supported on the firewall profile. (475112)
Series Affected: DCS-7170 Series
- When sFlow and monitor session are both enabled in the system, the mirrored traffic may be incorrectly throttled to the sFlow copp policy rate. (604224)
- The 32-bit image of EOS does not support the 7170B series switches. Use the 64-bit version instead (646589)
- IP Multicast packets that are software forwarded are not subjected to egress static NAT. (771707)

DCS-7130 and DCS-7170 Series

- On the 7170 series of switches, when IPv6 is used in VXLAN underlay, the UDP checksum

will be 0 for all the packets going out with an IPv6 VXLAN encapsulation.
(297660)

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

- Configuring Authentication Failure VLAN to same value as 'Internal VLAN' or 'AAA-server assigned successful VLANs', may lead to un-desirable behaviour (376571)
- When hardware flow tracking is enabled, tcpControlBits information element in IPFIX flow records for VXLAN encapsulated TCP flows is incorrect. (416079)
- On receiving COA-REQ for changing VLAN for EAPOL supplicant, the NAS deliberately fails authorization for the first time in the new vlan. (449890)
- The per-interface configuration to ignore reauthorization timeouts from the AAA server ("dot1x timeout reauth-timeout-ignore") is not designed to work if manual reauthorization is requested from the command line. The reason for this is that manual reauthentication is designed to force the supplicant to contact the AAA server to retrieve the latest credentials. (453044)
- CoA request for VLAN change for EAPOL supplicant returns CoA-NAK. The VLAN does get changed when device is reauthenticated and the AAA server sends the new VLAN in an Access-Accept response. (459726)
- Hitless FPGA upgrades are not supported due to hardware limitations. Therefore FPGA images will not be upgraded during SSU. (474978)
 SKUs Affected: CCS-720XP-24Y6-F, CCS-720XP-24Y6-R, CCS-720XP-24ZY4-F, CCS-720XP-48Y6-F, CCS-720XP-48Y6-R and CCS-720XP-48ZC2-F
- Negation operator (!) and 'assigned' keyword are not supported in Radius NAS-Filter-rule(attribute-id 92) Attribute. (475207)
- Ip options qualifiers are not supported in the NAS-FILTER-RULE attribute if present in the Access-Accept or COA packet from the AAA server. If Ip options are present in NAS-FILTER-RULE, it may cause supplicant to fail authorization. (480659)
- WOL feature is currently only available for phone trunk port, but not for regular trunk port. (481468)
- Fallback to MBA authentication feature may not work as expected with Caching Auth feature. (482333)
- Timeout value configured for MBA fallback feature can block MBA Auth for upto 60 seconds more than the configured value since first non-EAPOL packet is received. It is recommended to configure the hold period to be lower than the timeout value for MBA fallback. (483995)

- In Mac-Based-Authentication in Dot1x, once a host's mac address gets authenticated with one vlan, it cannot change its vlan to get authenticated with a new vlan. (490696)
- On phone trunk ports which are not controlled by Dot1x, phones do not honour port-security max-address limit and get installed in the L2 table even after the port-security limit is breached (494095)
- CoA request to change VLAN of multiple supplicants identified with NAS-PORT-ID may not work as expected. (495905)
- MBA for passive device can not be triggered by ARP/ICMP if IP locking is enabled (500439)
- If a switchport loses power, the credentials of the attached phone and PC will remain cached when authentication caching is enabled. However, when power is restored, the EAPOL enabled PC may not be reachable within the default reauthorization limit of 3 attempts and the supplicant will be removed as a result.
The solution in this situation is to increase the reauthorization limit using the per-interface dot1x command: dot1x reauthorization request limit 10 (500488)
- Limitation - 802.1X Web Authentication Implicit ACL feature doesn't work when IP Locking is configured on the interface. The workaround is to use 802.1X ACL based Web Authentication. (557986)
- PFC and pause generation are not supported when MACsec is enabled. (577310)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- An interface that acts as a Dot1x supplicant cannot be paired with a third party device that acts as a Dot1x authenticator. (643209) (1)
Series Affected: CCS-750 Series
SKUs Affected: DCS-7050CX4M-48D8
- Mirrored packets sent to a MACSec enabled destination interface are sent out unencrypted. (649877)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- Command "ptp free-running source clock system" is not supported. Although it may complete without errors, it will not take effect. (708316)
Series Affected: CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-722XPM and DCS-7010TX Series
- PTP will not synchronize over interfaces configured with half-duplex link speeds. (740727)

- Phone sending LLDP packets before getting authenticated (by sending non-LLDP packets like DHCP) may not get classified as phones.

The workaround is to configure "dot1x protocol lldp bypass" (742933)

- "VTEP to VTEP bridging" is not supported. (791265)
SKUs Affected: CCS-720XP-96ZC2 and CCS-720XP-96ZC2-F
- PTP transparent clock mode is not supported on interfaces which have MACsec enabled. The workaround is to use boundary clock mode, or interfaces without MACsec. (821972)
SKUs Affected: CCS-722XPM-48Y4 and CCS-722XPM-48ZY8
- The system will not boot if a cable connected to an ethernet switch is connected to the console interface of the system. To recover the system, remove the cable from the console port and power cycle. (934172)
Series Affected: CCS-720DF, CCS-720DP and CCS-720DT Series
SKUs Affected: CCS-710P-12, CCS-710P-16P, CCS-720DF-48Y-2, CCS-720DF-48Y-2F, CCS-720DF-48Y-DC, CCS-720DF-48Y-DC-2, CCS-720DF-48Y-DC-2F, CCS-720DF-48Y-DC-F, CCS-720DF-48Y-F, CCS-720DP-24S, CCS-720DP-24S-2, CCS-720DP-24S-2F, CCS-720DP-24S-F, CCS-720DP-24ZS, CCS-720DP-24ZS-2, CCS-720DP-24ZS-2F, CCS-720DP-24ZS-F, CCS-720DP-48S, CCS-720DP-48S-2, CCS-720DP-48S-2F, CCS-720DP-48S-F, CCS-720DP-48ZS, CCS-720DP-48ZS-2, CCS-720DP-48ZS-2F, CCS-720DP-48ZS-F, CCS-720DT-24S, CCS-720DT-24S-2, CCS-720DT-24S-2F, CCS-720DT-24S-2R, CCS-720DT-24S-F, CCS-720DT-48S-2, CCS-720DT-48S-2F and CCS-720DT-48S-2R
- Non-unicast traffic egressing on Port-Channel interfaces that are not uplinks do not use hashing to load balance. A single Port-Channel member will be designated to transmit non-unicast traffic for a given replication group (VLAN or multicast group). (1095704)
Series Affected: CCS-750 Series

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

AWE-5000, AWE-7000 and ZTX-7000 Series

- 100M speeds are not supported on the BASE-T interfaces Et1/1-1/4 (720921)
SKUs Affected: AWE-5310-F
- The system does not support the interface tx-queue shape rate commands, however the CLI still permits the configuration. Entering these commands will have no effect. (751249) (1)
- Arista Networks branded AOC-S-S-10G-yM manufactured before 2016 may not link up on interfaces 1-8. There is no workaround.

In more detail, we require the SFP transceiver memory to be compliant with

the SFF-8472. The exact memory requirements for a 10G AOC SFP are as follows. We require bit 4 of page A0h byte 3 to be set to 1. We require a non-zero value in page A0h byte 8. We require bit 2 of page A0h byte 60 to be set to 1. (788102)

SKUs Affected: 10GBASE-AOC and AWE-5510-F

- This platform only supports 64-bit versions of EOS. (794027)
- Arista Networks branded 1000BASE-T SFP transceivers starting with serial number AHT19 or XHT19 are not compatible with the first 8 interfaces Ethernet1/1 through Ethernet1/8. They are not capable of linking up on these specific ports. Workaround is to use Arista Networks branded 1000BASE-T SFP transceivers starting with XTH, ATH, XMD, and AMD serial numbers, or modules with XHT22 serial numbers and above. (810006)

SKUs Affected: 1000BASE-T and AWE-5510-F

- When all of the first 8 ports (Et1/1-Et1/8, xl710 ports) are active, sending and receiving traffic, the aggregate throughput of the 8 ports are unable to reach line rate at packets sizes below 1400 bytes. (815929)

SKUs Affected: AWE-5510-F

- Tunable SFP transceivers are not supported on the first 8 slots (Ethernet1/1 to Ethernet1/8). Therefore the interface configuration commands "transceiver frequency ..." and "transceiver channel ..." are not supported on the first 8 slots (Ethernet1/1 to Ethernet1/8). (823756)

SKUs Affected: 10GBASE-DWDM, 10GBASE-DWDM-T, 10GBASE-DWDM-ZR and AWE-5510-F

- Support of SFP-1G-T transceivers is limited on Ethernet1/1 through Ethernet1/8.

The following third-party SFP-1G-T SKUs are supported:

- Finisar FCLF-8521-3
- Kinnex A XSFP-T-RJ12-0101-DLL
- Avago ABCU-5710RZ
- Finisar FCLF8521P2BTL (described to be functionally equivalent to Finisar FCLF-8521-3)

The following Arista Networks SFP-1G-Ts SKUs are supported:

- Serial number with: "HN" or "PF"
- Serial number with: "TH"
- Serial number with: "MD"
- Serial number with: "HT" (907088)

SKUs Affected: 1000BASE-T and AWE-5510-F

- SFP 10G AOC transceiver page A2h memory is not accessible. The transceivers will operate normally, but digital diagnostic monitoring will not be available.

There is no workaround to access page A2h memory. (1165733)

SKUs Affected: 10GBASE-AOC and NIM-4S

- Tunable SFP transceivers are not supported. Therefore, the interface configuration commands "transceiver frequency ..." and "transceiver channel ..." are not supported. (1165735)
SKUs Affected: 10GBASE-DWDM, 10GBASE-DWDM-T, 10GBASE-DWDM-ZR and NIM-4S
- Arista Networks branded 1000BASE-T SFP transceivers starting with serial number AHT19 or XHT19 are not compatible and will not link up.
Use Arista Networks branded 1000BASE-T SFP transceivers starting with XTH, ATH, XMD, and AMD serial numbers, or modules with XHT22 serial numbers and above. (1165753)
SKUs Affected: 1000BASE-T and NIM-4S
- Arista Networks branded AOC-S-S-10G-yM manufactured before 2016 may not link up on NIM-4S interfaces.
There is no workaround. (1181419)
SKUs Affected: 10GBASE-AOC and NIM-4S
- Support of SFP-1G-T transceivers is limited.

The following third-party SFP-1G-T SKUs are supported:

- Finisar FCLF-8521-3
- Kinnex A XSFP-T-RJ12-0101-DLL
- Avago ABCU-5710RZ
- Finisar FCLF8521P2BTL (described to be functionally equivalent to Finisar FCLF-8521-3)

The following Arista Networks SFP-1G-Ts SKUs are supported:

- Serial number with: "HN" or "PF"
- Serial number with: "TH"
- Serial number with: "MD"
- Serial number with: "HT" (1181428)

SKUs Affected: 1000BASE-T and NIM-4S

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

CCS-750 and DCS-7060X4 Series

- When re-enabling a disabled power supply using `power supply enable module PowerSupply<n>`, the power supply may not re-enable.

The power supply can be recovered by removing and re-applying the input power to the power supply. (857791)

SKUs Affected: PWR-1511-AC-BLUE, PWR-1511-AC-RED, PWR-1511-DC-RED, PWR-1512-AC-BLUE, PWR-1512-AC-RED, PWR-400-DC-RED, PWR-500AC-F, PWR-500AC-R, PWR-501AC-F, PWR-511-AC, PWR-511-AC-BLUE, PWR-511-AC-RED, PWR-511-DC, PWR-511-DC-RED, PWR-745AC-F, PWR-747AC-F and PWR-747AC-R

DCS-7170 Series

- On devices with Algomatch enabled, removing an ACL on an interface might fail because the remaining ACLs on the system may not fit in the hardware tables. (192811)
- Ingress ACLs on IPv6 packets do not work when IPv6 uRPF is configured and AlgoMatch is enabled. (512394)

DCS-7130, DCS-7132 and DCS-7135 Series

- This platform only supports 64-bit versions of EOS. (615075)
- Several link flaps may be seen as the links are coming up. Debounce time CLI do not apply to Ethernet Interfaces on dropbear products. Link monitoring, Debounce and EOAM rely on PCS link status and not enabled on these since these interfaces use network PMA to set link status. The failure this reported here is a test issue. (633947)
- In the event of power loss, the switch will not report the correct reload cause. If power is interrupted briefly, the switch will report the reload as being caused by a watchdog. If the power loss lasts more than a few seconds, the switch will report the reload as cause unknown. (700291)
Series Affected: DCS-7130 Series
- When link loss is experienced on a MetaMux input interface, MetaMux may erroneously increment the error count for that MetaMux input interface. (723415)
Series Affected: DCS-7130 Series
- SwitchApp's published routing scale limits are approximations and simplifications so it is possible for a system to run out of internal resources for certain prefixes even though the overall scale is within the published limits. (764479)
- Maximum packet size supported by MetaWatch could be limited to 9214 bytes depending on platforms (775872)
SKUs Affected: DCS-7130LBR-48S6QD, DCS-7130LBR-48S6QD-F and DCS-7130LBR-48S6QD-R
- The commands for initializing an instance of an App have been changed. The CLI commands `Arista(config)# <app>` and `Arista(config)# <app> default` have been removed. Please use `Arista(config)#<app> instance N` instead. (778693)
SKUs Affected: DCS-7130LBR-48S6QD
- CL73 Autoneg is not supported. This means auto-negotiation on copper cables at speed 40G or higher is not supported. (832922)
Series Affected: DCS-7130 Series

- Asymmetric routing and NAT sync are not supported for TCP sessions. (848815)
- Use of unidirectional L1 configuration for a front panel Ethernet interface populated with a BASE-T type transceiver is not supported.
Setting unidirectional configuration to such an interface may produce unexpected results, including traffic loss.
Workaround is to use unidirectional configuration with non BASE-T type transceivers. (860391)

Series Affected: DCS-7130 Series

- Link flaps may occur when configuring a Layer 1 loop with two Ethernet interfaces (ie. snaked cable connecting two interfaces on the front panel, with each interface sourcing each other via the `l1 source` command).

Workaround: please avoid configuring loops. (905572)

SKUs Affected: DCS-7135LB-48Y4C, DCS-7135LB-48Y4C-F and DCS-7135LB-48Y4C-R

- LLDP transmit does not work if the front panel Ethernet interface is layer1 sourcing an interface (except a Switch interface on the Switch FPGA application). Transmit Discard counters may increment on the interfaces as these packets are dropped. (941796)

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

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

- Device is unable to support 1G optics on majority of the FPGA applications. This includes SwitchApp and 1G hardware PTP on MetaWatch. WhiteRabbit on MetaWatch is not affected by this limitation. (971303)

Series Affected: DCS-7132 and DCS-7135 Series

SKUs Affected: DCS-7132LB-48Y4C, DCS-7132LB-48Y4C-F, DCS-7132LB-48Y4C-R, DCS-7135LB-48Y4C, DCS-7135LB-48Y4C-F and DCS-7135LB-48Y4C-R

- MuxIn Interfaces from non-EB profile, do not get enabled when sourcing directly from different MetaMux instance. At least one another MuxIn Interface in that non-EB profile, must be sourced from either front panel interface or, switch interface from J2C, to make it enabled (982810) (1)

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

- If counters are cleared while traffic is running, front panel counters may become out of sync with their backing interface.

Ensure the counters are only cleared when there is no traffic flowing through the link. (1024514)

SKUs Affected: DCS-7130LBR-48S6QD

- An Arista FDK application with Tscore cannot be configured alongside MetaWatch, even if the configurations are disabled.

To use an Arista FDK application with Tscore, all MetaWatch instance configurations must be removed (not just disabled).

To use MetaWatch, any Arista FDK applications with Tscore must be disabled. (1089489) (1)

- Dynamic source NAT translations in different VRFs sharing the same ACL is not supported, when the ingressing traffic flows in both VRFs have the same source and destination address and port. (1105562) (1)

Series Affected: DCS-7132 Series

- Switch to Switch 11 sourcing is not supported with 1G speed configuration (1174034)

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

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

- When a system shuts down due to a power overload, EOS does not log a 'power overload' reload cause (371160)
- Mixed supervisor types are not supported on dual-supervisor modular systems. (450053)
- With scaled VRF, SVI and Arp entries, typically beyond claimed support, the Arp agent may experience a SIGQUIT during SSO switchover, on the new active supervisor. There is no workaround for this issue and switchover is ultimately successful. (495528)
- IP reachability between the ManagementActive interface and physical management interfaces may be impacted if both sets of interfaces are configured in separate VRFs. (885866)
- If Management0 interface redundancy is enabled with ipv6 neighbor monitoring, the backup interface can become active on supervisor switchover instead of the primary interface (893164)

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

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

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

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

- The L3 MTU on interfaces is not enforced on SVIs. (11659)
Series Affected: DCS-7500R Series
- Disabling IPv6 routing on an L3 interface is not supported when IPv4 routing is also enabled on that L3 interface. (43093)
- On DCS-7500 series, stateful switchover is not hitless. (54305)

- show interfaces counters rates' for port channels may show a slightly inaccurate data rate (including above 100%). (67367)
- A VXLAN encapsulated packet after de-encapsulation is not forwarded back to the port on which the original packet was received. Similarly, a native ethernet packet after VXLAN encapsulation is not forwarded back to the original receive port towards the remote VTEP. (75075)
- Due to hardware limitation, if an egress IP ACL is applied on an interface which is part of a VLAN with an egress router ACL, the filtering behavior for egress traffic of that interface may be undefined. (82094)
- For packets that hit rules across multiple ACLs, the hit counts are accounted for only in a single ACL. PACL counters take precedence over RACLs and MAC ACLs, and RACL counters take precedence over MAC ACLs. (85440)
- When an egress IPv6 RACL is applied to a VLAN, IPv6 packets may not be subjected to MTU checks. (88778)
- An egress ACL on a destination port of a monitor session only takes effect for Rx mirrored packets which are IP forwarded. The egress ACL will not work for bridged packets or Tx mirrored routed packets. (89915)
- When a shaper is configured the output of the shaper has a tolerance of +30% with 100-byte frames. This reduces to +5% as the frame size is increased to 600 bytes. The shaper variance reduces when the port happens to be the endpoint of a tunnel. (93546)
- The egress per-queue byte counter is not exact for certain packet types. For routed packets the egress per-queue byte counter may be offset by +/-3 bytes. For any packets originating from the CPU, the per-queue byte counter may be offset by -12 bytes. (97660)
- Packets that attempt to do a GRE key lookup where the GRE key lookup misses will have egress IPv6 ACLs applied with the ingress VLAN (102455)
- GRE key forwarding packets that are recirculated will have no QoS policy applied to them. (102458)
- ACL counters may increment spuriously during modification (117509)
- If a double tagged frame is received on an untrusted or DSCP-trusted interface, the CoS is retained if the outgoing frame is tagged, unless the frame is routed (131614)
- On the DCS-7280E and DCS-7500E series, configuring QinQ or SelectiveQinQ along with configurable TPID on the same interface is not supported (131757)
- VXLAN encapsulated packets received on MLAG peer-link will not be bridged after decapsulation to any MLAG interfaces. The workaround is to redesign the

network such
that VXLAN encapsulated packets are not received on the MLAG peer-link.
(133841)

- IPv6 Egress ACL deny logging is not supported on subinterfaces. (146487)
- Disabling IPv4 routing on an L3 interface is not supported when IPv6 routing is enabled on that L3 interface. When "no ip routing ipv6 interfaces" is configured to enable the forwarding of IPv4 packets over IPv6 nexthops over interfaces that do not have IPv4 address, but only a IPv6 address, it will also allow routing IPv4 traffic over these interfaces. (151356)
- Congestion on an egress interface used by a GRE tunnel that is a mirroring destination can cause drops on forwarding ASICs where the mirrored traffic ingresses the switch. (158001)
- On the DCS-7500E and DCS-7280 series, traffic outage is expected during transition from a single member LAG to two member LAG and vice versa. Use "platform sand lag hardware-only" to avoid this outage. (160439)
- When MLAG peer mac routing feature is enabled on both MLAG peers, the IPv6 neighbor may not be resolved when source ipv6 address in solicited neighbor is link local address. (165067)
- When the 'vxlan-routing' hardware TCAM profile is configured, VXLAN traffic flows in overlay may not work alongside IPv6 traffic flows in underlay. The workaround is to use the CLI command 'no platform sand ipv6 host-route exact-match'. (190017)
- L2-ECMP over VXLAN is not supported for multi-homed hosts under EVPN. (193475)
- Ingress ACLs are not supported on decap groups. (195568)
- In Tap Aggregation mode, Tap Aggregation features may not work for 802.1BR/VN tagged packets when 802.1BR/VN tag stripping is not configured. (198798)
- 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: 7800R3-48CQM-LC

- Native VLAN setting on Tap interfaces does not take effect. (227770)
- While in Tap Aggregation mode, ingress VLAN membership filtering does not work when packets are destined to a tool interface configured with egress packet truncation. (227841)
- While in Tap Aggregation mode, interfaces configured for egress truncation will continue to consume hardware resources even when they are down. (228384)
- LLDP functionality is not supported on interfaces with ingress truncation enabled. (229983)
- Sflow agent CPU utilization is around 10% even when hardware acceleration is enabled and the agent is not processing any flow samples. (251674)
- With MTU increased beyond 9100B, there can be some multicast packets of size greater than 9100B that are dropped as a port gets close to linerate. (261446)
- Changing PMF profile on one linecard may cause a brief disruption of PMF-based features on other linecards which should not be affected by the same change. (270849)
- A QoS policy, with a class map which has either set-tc or set-dscp on the ingress port and which has DSCP rewrite map on the egress port, results in indeterministic behavior with respect to DSCP remarking of the outgoing packet. (273320)
- If a QOS and PDP Policy both having police action are applied on an interface and the traffic hits both the policies then QOS Policer Counter is messed up. The "show policy-map type qos pmap-name" will give the PDP Policer Counter, not the QOS Policer Counter. (275994)
- Rate values for subinterfaces shown by "show interface counter rates" may deviate up to 15% from corresponding rate on parent interface. (278700)
- With TCP MSS Clamping feature configured, TCP SYN packets might get dropped under high CPU usage conditions if the incoming TCP SYN packet rate is high. (279370)
- Packets larger than 10196 bytes will be discarded by the forwarding ASICs. (280459)
- PFC on DSCP trusted ports is not supported. (281214)

- Routing is not supported on the parent interface of a L2 sub-interface. (287016)
- Loop protect feature is not supported when AlgoMatch HW is enabled. (290706)
- When a port becomes a LAG member, there may be a very brief moment where a few packets may be duplicated on the LAG. (295260)
- Fragmented IP packets will not match on flow-spec rules that include a match on TCP/UDP port. (297309)
- When all the L2 subinterfaces belonging to a parent interface are shutdown, packets are bridged and mac addresses are learnt on parent interface (300202)
Series Affected: DCS-7500R Series
- VXLAN routing is not supported on a modular system on which one of the linecards is configured for TAP aggregation. (306180)
- Show interfaces counters ip only show IPv4/Ipv6/MPLS counters per physical interface and not per port-channel (306209)
- When a BGP peer session is cleared, traffic which hits existing flow-spec rules may increment the wrong counter for a brief period of time. (339523)
- If non shaped sub interfaces are configured under same interface having shaped sub interfaces, traffic will egress out of the non shaped sub interfaces in a round robin fashion. It is recommended to not have shaped and unshaped sub interfaces under the same parent. (341851)
Series Affected: DCS-7500R Series
- Hardware accelerated sflow with profile sflowaccel-v2 doesn't work with 7500E linecards. (342130)
- Stale telemetry flow-spec statistics may be seen until the next "show flow-spec" command is issued. (345195)
- The sFlow sample output interface of IPv6 packet subject to an egress ACL, is set to unknown output interface. (354746)
- IPv4 packets with options in the IP header will not be subjected to Unicast Reverse Path Forwarding (uRPF) checks, and hence will not be dropped even if the source IP is unknown. (355542)
- unicast RPF is not enforced on IP packets when the destination IP needs ARP resolution (356347)
- Software forwarding is not supported for GRE or IP-in-IP tunnel terminated packets using decap-groups. (358369)

- The egress DSCP value in the Outer IP header for packets that undergo headend replication after routing is derived from the Traffic Class (and not copied from the ingress DSCP value). (366189)
- While in Tap Aggregation mode, if the identity tag configuration of a Tap port is changed, traffic received on this port may be sent to an invalid multicast destination momentarily. This will cause the DchUnreachables counter to increment. (372757)
- TX mirroring shows VXLAN encapsulated packets that are actually dropped by split horizon. (375795)
- If more than maximum supported 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)
- Syslog message informing user about enabling a counter feature that was previously disabled due to hardware resource conflict may not be emitted. (389299)
- Configuring more than 250 BFD sessions on a physical interface may result in many of those sessions flapping because of hardware drops on the interface. (389387)
- Drop Precedence based tail drop thresholds can be ignored when multiple queues are congested and buffering resources are low. (395693)
- Queue build-up doesn't show up for 64 byte packets in case of congestion via following commands:-
 - 1) platform fap diag diag cosq non_empty_queues
 - 2) show interfaces queue length

Following command should be used instead:-
show queue-monitor length (398937)
- PIM sparse mode feature which allows installation of outgoing interfaces on the non-DR for a faster failover does not work on L3 sub-interfaces. This feature relies on an egress ACL to stop multicast traffic from going out of the interface. Multicast egress ACLs do not work on L3 sub-interfaces due to a platform limitation. If the feature is configured, there will be duplication of traffic. The workaround is not use the feature on such interfaces. (405696)
- When a policy-map with match on inner VLAN is programmed on an interface, if single tagged packet traffic is sent to that interface, the behaviour of

that packet is indeterministic as inner VLAN tag is calculated using an offset and vlan-tag-format qualifier doesn't differentiate between single and double tagged packets. (415587)

- When MPLS label is pushed on the traffic ingressing CoS trust port then EXP bits are not derived from TC based on packet's COS. (417108)
- Byte counts are not supported for IPSec tunnels. (419031)
- SflowAccel supports up to 200 VRFs. (423892)
- PAUSE frames with correct DMAC=01:80:C2:00:00:01, Ethertype=0x8808, opcode=0x0001 are never forwarded and are always consumed by the hardware. (426047)
- Header truncation for mirrored packets can not be used when the destination of a monitor session is a GRE tunnel. (428547)
- A port transmitting packets smaller than 100B that were tunnel terminated on the switch might be limited to ~90% linerate (429797)
- When the VXLAN routes are programmed using 2 hierarchical FEC levels in hardware, and the outgoing interface of a VXLAN underlay route changes from one forwarding chip to another, then it can negatively impact VXLAN overlay route convergence. (440233)
- Egress MAC ACLs are matched on the incoming packet's ethernet header and not on the rewritten ethernet headers. (458724)
- Dynamic port-channels coconfigured using LACP are not supported as destinations for monitor sessions (463186)
- SflowAccel does not send flow samples to collectors routed via BGP routes with additional backup paths. (476553)
- When the feature "flow" in a TCAM profile is configured with only the "drop" action, the profile is rejected. The workaround is to add any other action to feature "flow". (486849)
- When a LAG is used as a destination for a monitor session, a given set of packets mirrored from a Rx source will hash differently compared to the same set of packets captured from a Tx source, due to hardware limitations. Consequently, symmetric hashing will not work as expected with a monitor session which captures both Rx and Tx traffic from a single port, and using a LAG as a destination.

A workaround is to only capture traffic on the ingress (Rx direction). Another workaround is to use two different monitor sessions, and

destinations, for the two directions of the target traffic, and implement symmetric hashing on a second stage. (486994)

Series Affected: DCS-7500R Series

- When GUE is enabled with outer IP hashing, inner IP fields are not included in the load balance key of MplsOverGre transit packets in tc-counters TCAM profile. (491706)

Series Affected: DCS-7500R Series

- When "ip hardware fib optimize" is enabled for a non-nibble aligned prefix-length, local-remote MPLS Pseudowires must have Control Word enabled for MPLS decap forwarding to work correctly. If Control Word is not enabled, then ETHoMPLSoETH traffic that ingress with the first nibble of the inner DMAC starting with "0x4" will be speculative parsed as IPv4oMPLSoETH and incorrectly dropped. (496624)

Series Affected: DCS-7500R Series

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

Series Affected: DCS-7500R Series

- All asymmetric IRB traffic flows forwarded over EVPN VXLAN toward to a single multi-homed destination are not load balanced across the VTEPs that form the L2 ECMP group. However, note that the traffic destined to different hosts are statically load balanced across the VTEPs. (501343)

- VRF label termination into non-default VRF is not supported for multi-label MPLS packets. (504763)

- When BGP neighbors advertise Flowspec rules with police actions where the number of rules exceed what can be programmed in TCAM, policers will be allocated in hardware for all requested rules even though best-effort programming will only program the subset of rules that fit in TCAM. As a result, hardware policer resources will be unnecessarily consumed for rules which are not installed in hardware. (506688)

- SVI egress counters are not incremented for IPV4 or IPV6 multicast packets. (508471)

- Egress counter features configured via "hardware counter feature ..." do not include packet or byte counts for packets sent from the local CPU. (525380)

- Counter information flows from the forwarding chips to the tables that support TerminAttr in a very bursty manner. When averaged over a several minute window the rates will be accurate, but for shorter intervals calculated rates may oscillate significantly above and below the nominal rate.

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

- Hitless restart is not supported with Macsec proxy (540419)
- `show forwarding destination` may give incorrect results when the traffic would hit an egress IPv4 ACL using bridged IPv4 traffic or hit an egress IPv4 ACL using one of the following qualifiers: TCP established, fragment, and L4 port range. Note that TCP and UDP egress IPv4 ACL rules use the fragment qualifier. (554480)

Series Affected: DCS-7500R Series

- Route churn under hardware FIB exhaustion may cause BFD sessions to flap. (556089)
- VRRP IPv6 using VRRP IPv4 MAC prefix may have unexpected packet forwarding behavior when the VRRP IPv6 and IPv4 have different states and one of the states is Master. (567504)
- Different streams being rate-limited by the same group policer might exhibit bias towards some streams and fair policing across all the streams is not guaranteed. (578041)

Series Affected: DCS-7500R Series

- When all members of a port channel have "forwarding transmit disabled", traffic will not be forwarded to other VLAN members. (587648)
- While in Tap Aggregation mode, traffic steering service policy rules that match on fields from the MPLS header may result in false matches on packets without an MPLS header.

Remove the MPLS key-fields from feature 'tapagg mac' in the active TCAM profile. (592184)

- While in Tap Aggregation mode, traffic steering service policy rules that match on TCP flags may result in false matches when using the tap-aggregation-user-tcp-flags TCAM profile, or another profile based on this one.

Remove L4 key-fields from the `tapagg mac` feature in the active TCAM profile. (592191)

- While in Tap Aggregation mode, traffic steering service policies with VLAN matching may match untagged traffic erroneously.

Add the vlan-tag-format key-field to features that filter on VID in the active TCAM profile. (597603)

- VLAN matching in ACLs attached to monitor sessions may result in false matches on untagged traffic. (598186)

- Packets greater in size than the egress interface MTU (fragmented packets) are not counted in "show interfaces counters ip".

Packet size should be limited to MTU to avoid fragmentation and loss of packet and byte counts. (601808)

- For IPsec tunnel mode on Jericho2, the output of show ip security connection detail incorrectly shows the replay window as 16384, the correct replay window is 128 and fixed to that value. (609180)
- Security counters are not updated correctly when using Vxlansec with SecTag2 encapsulation format. (630807)
- This platform only supports 64-bit versions of EOS. (646592)
SKUs Affected: 7816R3-FM, DCS-7816-CH, DCS-7816-SUP and DCS-7816-SUP1S
- When 'show forwarding destination' is used in an L2 forwarding scenario, the source MAC address of the packet can be learnt in the MAC address table. This can result in MAC moves if the MAC address was already learnt.

A workaround is to clear the MAC address after running the 'show forwarding destination' command. (650968)

- For configured load-intervals of less than 1 minute, counter rates may fluctuate significantly above and below the actual rate.

The rate will become increasingly accurate as the configured load-interval approaches 5 minutes. (684152)

- Packets originating from or forwarding via the CPU through an SVI are not counted by the egress SVI counter feature ("hardware counter feature vlan-interface out"). (688993)
- Counter sampling for ingress sFlow on sub interfaces is not supported. (718833)
- IPv4 and IPv6 strict mode Unicast Reverse Path Forwarding (uRPF) does not work if the route matching source IP address resolves via an ECMP nexthop. Such packets will be dropped. (757793)
- With scale QoS scale policy-map config with ACL rules, configure-replace might CLI time-out (766811)
- For tunnel destinations, the decision to fragment is based on unencapsulated packet size, while MTU enforcement is based on the encapsulated packet size. This may cause unencapsulated packets near the MTU packet size to be dropped instead of fragmented. (768410)

- At high combined scale of VLANs and pseudowires in VPLS or L2 EVPN, exhaustion of "managedTt inLif" resources may occur. This may prevent hardware programming of any configured patch panel patches, resulting in "Unprogrammed" status for one or more patches. It may also result in VPLS packets drops in the decapsulation direction. This situation can be detected by "show hardware capability".

The workaround is to reduce the scale of these or other features which use the managedTt inLif resource and then perform a shut/no shut on the affected patch panel or subinterface. (781993)

- VXLAN routing is not supported when the maximum hierarchical FEC level in hardware is set to 1 using the "ip hardware fib hierarchical next-hop max-level 1" CLI command. (791422)
- "ip hardware fib next-hop sharing" is not supported with VXLAN. (795334)
- Traffic-policy 'redirect next-hop' and 'redirect next-hop group' actions on L2 interfaces are not supported. (821920)
- If multiple traffic policies are applied such that the TCAM usage for traffic-policy is close to scale (~100%), then the traffic-policy actions may not work if the Sand agent is restarted.

The workaround is to remove these policies from applied interfaces, and re-apply each of them. (821984)

- IP nexthop group entries are categorized as MPLS tunnels in the 'show hardware capacity' and 'show platform sand l3 summary' CLI commands. (823863)
- The VLAN interface counter feature is not compatible with dynamic VLANs used in EVPN VXLAN prefix routing setups. When the VLAN interface counter feature is enabled, a "Counters unavailable" message will be shown for each of EVPN VXLAN dynamic VLAN in the "show interface counters" CLI command output. (840994)
- The LoopProtection feature is not supported for L2 sub-interfaces at high scales. (853892)
- There may be packet drops when a MLAG utilizing Fast MAC Redirection carries more than 100Gbps of traffic (861038)
- The virtual-cable feature is unsupported, but the configuration CLI is not blocked. Configuring this unsupported feature will result in the Sand agent repeatedly restarting.

The workaround is to remove the virtual-cable feature from the running config. This will allow the Sand agent to recover. (893360)

- If 'switchport vlan tag validation' is configured, packets with more than 2 VLAN tags will be dropped. (898006)
- If a Tap Aggregation policy map is applied to an interface while running in the tap-aggregation-default profile, or any user profile that does not have the count action in the Tap Aggregation features, then the message "ACE-6-ACTION_NOT_AVAILABLE" will be logged for StatId and StatProfile. These messages mean that the policy ACL counters will not work for these policies. Other configured actions are not affected. (939008)
- Arista Networks Release Notes for EOS Release 4.31.1, 4.31.2 and 4.32.0 incorrectly documented the 7816B-CH in the New Platforms and Hardware and Supported Hardware sections. 7816B-CH is not introduced in those releases. 7816L-CH is introduced/supported in 4.31.1, 4.31.2, 4.32.0 and newer releases. Newer versions of EOS release notes after April 20, 2024 correctly document the support for 7816L-CH. (949378)
- Replacing one PBR rule with another rule (by changing the match criteria and/or actions) at the same logical priority in the PBR policy is considered both an addition of a rule, and a removal of the rule previously configured. Any policy update containing this replacement operation cannot be programmed using in-place update of TCAM. (955171)
- In a multi-chip system, rapid host moves may lead to temporary divergence of MAC address tables between chips. After host moves quiesce and at least a third of the configured MAC aging time passes, the MAC address tables will converge.

A workaround is to clear the MAC address table. (983999)

- Source VTEP pruning with VTEP-to-VTEP bridging is not supported if the underlay L3 interface is an SVI and the nexthop MAC is unknown. Workaround is to configure a static MAC address for the nexthop. (993682)
- The flowspec redirect action is not supported on bridged traffic. (999653)
- If a VLAN-to-VNI mapping configuration under VXLAN is updated while traffic is flowing in the VLAN, then some packets may be unexpectedly flooded in the old VNI during the transition.

Workaround is to perform the configuration update while traffic is stopped. (1016288)

- The EAPI output of `show interfaces counters` for subinterfaces categorizes packets as "inUcastPkts", but it actually counts all packets (unicast, broadcast, multicast). (1046158)
- User-Defined Fields, or UDFs, are defined as part of an access-list filter and are comprised of an offset, length and pattern match. The behaviour of a UDF

where the offset value exceeds the packet's size is undefined and may result in false matches. (1058051)

- While in Tap Aggregation mode, when using Traffic Policies for forwarding, if the forwarding destination has no active egress interfaces, then the matching traffic may be forwarded to the default forwarding destination on the ingress TAP interface.

If it is desirable to drop this matching traffic instead, then be sure to configure the `drop` action in a custom TCAM profile for features `dbTrafficPolicyIpv4` and `dbTrafficPolicyIpv6`. (1114661)

- LANS does not support latency recording with bursty traffic when notifying mode is not configured.
To enable notifying mode, the `queue-monitor length notifying` global config command may be used. (1121840)
- Traffic-policies will not work if "packet-type que outer-ip" is configured in load-balance policies. (1156275)
- Only remote VTEPs that are reachable get added to the VLAN flood set when "vxlan flood vtep learned data-plane" is configured under VXLAN. (1189284)
- For L2 and L3 subinterfaces, VLAN ranges are not supported, and therefore not a valid encapsulation configuration. Thus, the subinterface status will remain down and inactive. (1213201)
- Adding or removing the PFC config on an interface can generate a re-assigned profile syslog message for other traffic-classes if they have invalid profiles assigned.
The workaround is to delete invalid profiles and their corresponding mappings. (1254788)

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

- "Drop" MAC entries only drop packets whose destination MAC address matches the "drop" MAC entry. Packets whose source MAC address matches the "drop" MAC entry are not dropped by the "drop" MAC entry. (341123)
- Strict Priority scheduling for 400G interfaces don't work properly for 64 byte packets when traffic rate is above PPS limit of the chip. (388517)
- IP forwarded traffic that is mirrored to a GRE tunnel destination will have the PCP bits of their VLAN tag set to 0. (400909)
- Application of low rate shaping on multiple interfaces with line-rate traffic into the box will create a backlog in DRAM. Due to this majority of traffic will be served to egress from slower DRAM instead of SRAM. In the congestion state, even after the removal of shaping with line-rate traffic ingress from multiple ports, switch won't be able to deliver full

throughput. The recovery from DRAM to SRAM will happen over a period of time once the congestion is removed. (414481)

- Set dcsp will not apply to mpls routed packets that egress with mpls labels (418791)
- Enabling Tx mirroring may cause the counter VOQ_SRAM_ENQ_RJCT_PKTS to increment, with reason QnumNotValid, as shown in the output from "show platform fap counters pipeline". There is no impact to mirrored or forwarded traffic. (459862)
- Unidirectional links are not supported at 40G speed on QSFP100 interfaces using the CRT50216 PHY. (467672)
SKUs Affected: 7358-16C, 7368-16C, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3-48CQM-LC, DCS-7280CR3-32D4, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-96, DCS-7280CR3-96-F, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-96 and DCS-7280CR3K-96-F
- SFlow samples of packets which are dropped by an ACL will have incorrect output interface indication. (470650)
- Certain speed groups are limited to one active serdes rate configuration. These speed groups are those that show only one serdes rate in the Active column of the output of "show hardware speed-group status". (486659)
SKUs Affected: 7500R3-24D-LC, 7500R3-24P-LC, 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-36D-LC, 7800R3-36P-LC, 7800R3-48CQ-LC, 7800R3-48CQM-LC, 7800R3K-36DM-LC, 7800R3K-48CQ-LC, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-96-F, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32D4A-F, DCS-7280CR3K-32D4A-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-32P4A-F, DCS-7280CR3K-32P4A-R, DCS-7280CR3K-96-F, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4A-S-F, DCS-7280CR3MK-32D4A-S-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F and DCS-7280PR3K-24-F
- The sFlow flow samples generated for VXLAN decapsulated packets, have the first 6 header bytes of the sampled packet set to zeros. (502957)
- The "phy link serdes mapping direct" CLI command requires the 25g serdes rate to be configured, unlike what is displayed in "show hardware speed-group configuration". (523850)

SKUs Affected: 7500R3-36CQ-LC, 7500R3K-36CQ-LC, 7800R3-48CQ-LC, 7800R3K-48CQ-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32P4, DCS-7280CR3K-96, DCS-7280CR3MK-32D4 and DCS-7280CR3MK-32P4

- Encapsulating PTP packets over VXLAN tunnel is not supported. (526461)
- When Hardware Accelerated Sflow is configured the uptime in sFlow datagram is updated every 1/4 of milli sec, rather than every milli sec as prescribed in sFlow standard. (531769)
- Sequence number in sFlow datagram doesn't reset to one when hardware accelerated sFlow is disabled.
The sequence number in sFlow datagram gets reset to 1 for those sFlow sub-agent when the line-card is reset/ restarted on which the sFlow sub-agent is residing. (531983)
- With loose mode uRPF configuration, the switch doesn't drop packets with an uRPF lookup (using source IP) matching a drop route. (531998)
- When SandCounter agent is down, datagram counters are not accounted. Hence in show sflow output, the "Number of Datagrams" entry displays lesser value than expected. (536213)
- If all front-panel ports are populated, and traffic load is high, an improperly programmed current limit could be exceeded generating a spurious power-off. (536331)

SKUs Affected: DCS-7280CR3-96 and DCS-7280CR3K-96

- Symmetric hashing is always on for IPv6 addresses, when both source and destination addresses are included in the applied load-balance profile. (557714)
- In transparent two-step mode, PTP packets egressing the same port with the same sequenceId may cause transient jumps in correctionField.

The workaround is to use one-step transparent clock mode (564682)

- Switching over in SSO mode fails when initiated on a system with no linecards inserted. (604537)
- In Tap Aggregation mode, when both header removal and truncation are configured on a tap port, a packet will not egress the switch if it has a size less than 64 or has an invalid ethertype, after header stripping and truncation actions are performed. (733144)
- MAC access-lists are not applied on L2 protocol packets even when L2 Protocol Forwarding profile is applied. (735328)

- The "Tap Aggregation GRE tunnel termination" feature only works when the hardware forwarding profile is set to "system-profile-tap-aggregation". Otherwise, packets will egress without GRE header being stripped.

Note that even in this case, if the Tap Aggregation counter is enabled and the configured TCAM profile supports GRE tunnel termination counter, "show tap aggregation counters tunnel termination gre" will still show counters incrementing for the matching GRE packets. (800865)

- Ethernet frames ingressing on a routed interface with multicast MAC as destination MAC address may be unnecessarily copied to CPU. (910627)

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

- When operating under fabric egress replication mode (``platform sand multicast replication default fabric-egress``), multicast traffic egressing (or dropped in the egress pipeline of) Ethernet ports is incorrectly counted as unicast traffic in the "Egress Queue Counters" section of the ``show interfaces counters queue`` command. In other words, all traffic will be counted as unicast. (488383)
- ZTP is not supported on interfaces controlled by the 'hardware port-group' command. (904876) (1)
 SKUs Affected: DCS-7280CR3-36S and DCS-7280CR3A-32S
- The `noclear` option in the command `"show platform fap acl tcam hw bank <num> noclear"` will not work as expected. The hit indication bit will be cleared after dumping the TCAM rules. (1206198) (1)

DCS-7280R3A and DCS-7800R3A Series

- When software sFlow is enabled, ingress packet discards (seen with `"show int counters discards"`) may occur at packet sizes less than 700 bytes. (601920) (1)
- Packets mirrored from a TX source to a GRE tunnel destination will have incorrect metadata in the "key" field of the GRE header. (763403)
- PTP Two-step Transparent Clock mode is not supported on this platform. Use 1-step TC mode instead. (891460)
- Real-time Weighted Fair Queueing (Rtbfq) is not supported on Jericho2C+ systems. (1045586)
- IPsec and VXLANsec interop with other implementations do not work since while the control plane negotiates 32b sequence number mode with the peer, the data plane operates in 64b sequence number mode.

Workaround: Configure "no anti-replay detection" as a way to force both control and data planes to operate in 32b sequence number mode. (1175579) (1)

- In the output of 'show ip security connection detail' command for an IPsec or secure VXLAN tunnel, 'Packets outside replay window' field displays the cumulative count of both late and replayed packets, while 'Replay' field consistently displays 0. (1205415) (1)
- The noclear option in the command "show platform fap acl tcam hw bank <num> noclear" will not work as expected. The hit indication bit will be cleared after dumping the TCAM rules. (1206198) (1)

DCS-7280R3A Series

- When software sFlow is enabled, ingress packet discards (seen with "show int counters discards") may occur at packet sizes less than 700 bytes. (601920) (1)
- ZTP is not supported on interfaces controlled by the 'hardware port-group' command. (904876) (1)

SKUs Affected: DCS-7280CR3-36S and DCS-7280CR3A-32S

- IPsec and VXLANsec interop with other implementations do not work since while the control plane negotiates 32b sequence number mode with the peer, the data plane operates in 64b sequence number mode.

Workaround: Configure "no anti-replay detection" as a way to force both control and data planes to operate in 32b sequence number mode. (1175579) (1)

- In the output of 'show ip security connection detail' command for an IPsec or secure VXLAN tunnel, 'Packets outside replay window' field displays the cumulative count of both late and replayed packets, while 'Replay' field consistently displays 0. (1205415) (1)
- The noclear option in the command "show platform fap acl tcam hw bank <num> noclear" will not work as expected. The hit indication bit will be cleared after dumping the TCAM rules. (1206198) (1)

DCS-7280R3 Series

- Qumran2A based SKUs cannot support VXLAN decap packet matching for TCAM features, including PBR and traffic-policy. (1235166)

DCS-7280R3 Series

- ZTP is not supported on interfaces controlled by the 'hardware port-group' command. (904876) (1)
SKUs Affected: DCS-7280CR3-36S and DCS-7280CR3A-32S

DCS-7280R3A Series

- In the output of 'show ip security connection detail' command for an IPsec or secure VXLAN tunnel, 'Packets outside replay window' field displays the cumulative count of both late and replayed packets, while 'Replay' field consistently displays 0. (1205415) (1)

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

- Priority tagged packets (802.1Q tagging with the VLAN set to 0) being reflected by an Ethernet reflector interface configured in the egress direction with MAC address swap action will have the priority tag stripped from the packet before they are transmitted. (354190)
- Only port-channel interfaces are supported as peer links. Single member port-channel interfaces are supported. (409958)
- If a packet that is decapsulated with a decap group is mirrored to a GRE tunnel, the packet's ingress VLAN will be set to 0 in the upper 2 bytes of the GRE key. (430986)
- Egress ACLs does not match packets with UDP ports 4754 and 4755. (455922)
- Traffic matching PBR rules will be dropped if it violates the MTU of the PBR destination port (471548)
- A learnt MAC address won't age out in case there is only continuously running unicast RPF drop traffic from the source MAC. The workaround is to clear the dynamically learnt MAC from the MAC address table. Once the MAC address is cleared, the MAC address won't be learnt if all traffic with the source MAC address is dropped by unicast RPF. (536215)
- If there are multiple mirror to GRE monitor sessions sharing the same GRE tunnel parameters (i.e. source and destination IP address, DSCP values, TTL, protocol, all matching), and if any of these sessions has rate limiting enabled, then all these monitor sessions may either have rate limiting functionality disabled, or share the rate limit configured in one of these sessions, following a switch reload.

A workaround is to remove these monitor sessions and reconfigure them again. (548973)

- During EVPN MPLS encapsulation, the VLAN priority bits of payload Dot1Q header will be marked using the traffic-class to CoS map configuration in the global QoS map. (550987)
- QoS policy-map is not supported on Pseudowire decapsulated traffic flows. (574600)
- On switches with B52 external PHY devices, the interface command "traffic-loopback source system device phy" will result in packets being both transmitted and looped back. (582409)
- When there is congestion in Tc6/Tc7 across multiple queues. Tc6/Tc7 Queues can hog up entire DRAM. This can lead to DRAM rejects of lower TC traffic of unrelated stream thus causing some discards. In the current case we had 2 to 3% lower instead of 100% throughput of data traffic (591523)
- Renamed TCAM Qualifiers l4ops_ip_pacl, l4ops_ip_qos, l4ops_ip6_pacl, and l4ops_ip_racl into l4ops_24b, l4ops_7b, l4ops_3b and l4ops_18b. This can cause customer scripts using `show platform fap pmf` command to fail if they are using any of the above renamed TCAM qualifiers. (595868)
- On Gen4 Faps , SSO for TCAM based ACL rules is hit full. There is a brief window where ACLs are not applied to the flows. (606756)
- On Gen4 Faps , SSO is hit full for TCAM based ACL rules. (606764)
- When outer IP hashing is enabled, GUE decap followed by MPLS NULL label termination is not supported. (607977)
- Configuration of "encapsulation dot1q vlan" is only supported on sub-interfaces. (608383)
- Software forwarding of MPLS-over-UDP terminated packets is not supported. (613102)
- If a packet matches both a Tap Aggregation steering rule, as well as either a security ACL or traffic-policy rule to drop the packet, the steering policy counter will still increment. (620429)
- SandTm restart is hitful for named TC to DSCP QoS maps. (632481)
- SandOam may restart with 200+ DM sessions running at 10 milliseconds. The workaround is to use higher tx-interval like 100 milliseconds, 1 second, 10 seconds, 1 minute, 10 minutes. (642106)
- SandL3Unicast on the peer DUT causes a few RMEPs to be marked as unreachable on the local DUT. A SandOam restart on the local DUT recovers from the problematic state. (646125)
- IP UDFs in ACLs may not work correctly on IPv6 packets with a routing extension header when matching on bits beyond the IPv6 header. (650286)

- VXLAN is not supported with greater than 256-way ECMP in the underlay. (652266)
- GUE tunnel decapsulation does not support UDP destination port 2152. (675837)
- In RFC2544 throughput test, oversubscribing the ingress interface causes hardware test engine lockup and test traffic won't be generated. Restart the SandFapNi agent. (684101)
- A pseudowire with Cos-To-Tc or Dscp-To-Tc QoS map applied will classify all of the decap traffic to default traffic-class of 1 in the presence of entropy or the null label. (694370)
- While RFC2544 test is in progress, received and sent packet counts may not match. (694683)
- If egress MAC/IPv4 ACL is applied to the front panel port and egress IPv4/IPv6 ACL is applied to SVI with ACL sharing enabled for the SVIs, only the ACL applied to the SVI will match the routed traffic. (695784)
- When using IPv4 route optimization using "ip hardware fib optimize .." command, the number of VRFs supported on the system is reduced to 30 non-default VRFs. (720042)
 SKUs Affected: 7800R3-36D-LC, 7800R3-36P-LC, 7800R3-48CQ-LC, 7800R3-48CQ2-LC, 7800R3-48CQM-LC, 7800R3-48CQM2-LC, 7800R3-48CQMS-LC, 7800R3A-36D-LC, 7800R3A-36D2-LC, 7800R3A-36DM-LC, 7800R3A-36DM2-LC, 7800R3AK-36DM-LC, 7800R3AK-36DM2-LC, 7800R3K-36DM-LC, 7800R3K-48CQ-LC, 7800R3K-72Y-LC, DCS-7280CR3-32D4, DCS-7280CR3-32P4, DCS-7280CR3-36S, DCS-7280CR3-96, DCS-7280CR3K-32D4, DCS-7280CR3K-32D4A, DCS-7280CR3K-32P4, DCS-7280CR3K-32P4A, DCS-7280CR3K-96, DCS-7280CR3MK-32D4, DCS-7280CR3MK-32D4S, DCS-7280CR3MK-32P4, DCS-7280CR3MK-32P4S, DCS-7280DR3-24, DCS-7280DR3K-24, DCS-7280PR3-24, DCS-7280PR3K-24, DCS-7280SR3-48YC8, DCS-7280SR3K-48YC8, DCS-7280SR3K-48YC8A and DCS-7280SR3M-48YC8
- The MPLS no-php mode is not supported with L3VPN and VPWS services. (720213)
- The packet counting logic for CFM Loss Measurement adheres to Version 1 of the protocol although the LMM is sent with version 0. This can result in an incorrect loss being reported. (720859)
- OAM doesn't generate SLR's in response to SLM's if Test ID in the SLM is more than 1048575 (0xFFFFF), instead it generates LMR's. (726371)
- CFM Loss Measurement will report incorrect loss if the loss between consecutive LMM/LMRs is greater than 32K packets. (729327)
- When using IPv4 route optimization using "ip hardware fib optimize .." command, the number of VRFs supported on the system is reduced to 14 non-

default VRFs. (731721)

SKUs Affected: DCS-7280SR3-40YC6 and DCS-7280SR3E-40YC6

- Traffic throughput on an internal recirculation interface is capped at 100Gbps. (738474)
- When policy configuration changes are made with segment-security, some counts may be missed or misclassified under "show segment-security counters" for up to 30 seconds. (769949)
- Post SandOam restart RFC2544 generator test is initiated for subinterfaces that are operationally down. (772320)
- CFM PDUs received on the Up-MEP interfaces that are to trapped/dropped get counted in the following platform counters: "Tx Total Err On Sop Pkt Counter", "Tx Total Pkt With Err Counter", "Total Byte With Err Counter", "Tx Total Err On Sop Byte Counter", "TxTotalPktWithErrOnEopCounter" and "TxBytesWithErrOnEopCounter". (796338)
- When ETH/MPLS/IP/GTP packets are MPLS terminated, hashing on GTP TEID will not occur.

A workaround to hash on GTP TEID is to use the custom UDP payload hashing feature. (813466)

- Hardware resources are exhausted when hardware accelerated sFlow is enabled with scaled subinterfaces, predominantly with port-channel subinterfaces. Possible workarounds are disabling sFlow hardware acceleration or disabling subinterface ifindex stamping by running "no sflow sample input subinterface". (813602)

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

- ECN re-marking is not supported on shaped sub-interfaces (826500)
- When PTP is enabled, the command 'vxlan qos dscp ecn rewrite bridged enabled' does not take effect. (830876)
- Powering off a Linecard which is transmitting high bandwidth on traffic class 6 and 7 may lead to flap of protocol session on other Linecards. Workaround is to shutting down all interfaces of the Linecard before its removal. (832211)
- When the device is configured with 'vxlan qos dscp ecn propagation decapsulation disabled', the inner DSCP for VXLAN decapsulated packets is retained even if the ingress port is in cos or untrusted mode and global DSCP rewrite is enabled. (836505)
- Mirror copies, or sFlow samples, of decapsulated VXLAN traffic will have bytes missing from the underlay headers. (843636)

- A MAC address that has not been refreshed within the configured aging time will age out even if it has been refreshed shortly after the configured aging time has elapsed.

A workaround is to increase the configured aging period via "mac address-table aging <AgingPeriod>" (858321)

- The switch does not learn the MAC address of routed packets ingressing on VLAN interfaces when using IPv4 route optimization involving two prefix lengths.

The workaround is to set the ARP aging timer to be smaller than the MAC address aging timer, so that ARP packets refresh MAC addresses. (868827)

- After changing configuration related to hardware accelerated sFlow, it takes 90 seconds for packet sampling to resume. (876765)
- The following two configurations do not work together:
 1. IPv4 optimize configuration of internet profile or any other compression configuration - "ip hardware fib optimize prefixes profile internet" or "ip hardware fib optimize prefix-length 21 compress <>".
 2. PBR policy matching on nexthop-group.


```
policy-map type pbr PMAP1
  10 match ip any any nexthop-group <>
```

The work around is to do either of the below:

1. Remove the optimize configuration
 2. Use an optimize configuration without compression
 3. Use urpf-internet profile. (893189)
- Even with custom CoS-To-Tc map applied on the ingress interface, DSCP remarking follows the global CoS-To-Tc map (on ingress) and global Tc-To-DSCP map (on egress). (902952)
 - Custom CosToTc map doesn't consider defaultCos on interface, instead maps to traffic class 1 (903830)
 - CFM protocols do not work if egress sFlow is active on the interface on which CFM MEP with direction Up is defined.
Disable sFlow on CFM UP MEP interface. (905833)
 - RFC2544 Initiator benchmarking tests with 64 byte packet size is only able to test max throughput of 256Gbps because of hardware limitation.
400Gbps RFC2544 Initiator throughput testing is supported for packet sizes greater than 100bytes. (915701)

- A traffic-policy match which includes 'encapsulation gtpv1 ipv4', but no inner match criteria will match on both IPv4 and IPv6 traffic. This same behavior is also seen for 'encapsulation gtpv1 ipv6'.
Workaround: add qualifiers inner-src-ip-label, and inner-dst-ip-label, to feature "traffic-policy port ipv4/6" in the TCAM profile (917902)

- Traffic-class to CoS rewrite map does not work for multicast and broadcast traffic on L3 subinterfaces. (924908)
Series Affected: DCS-7280CR3, DCS-7280DR3, DCS-7280PR3, DCS-7280SR3, DCS-7500R3 and DCS-7800R3 Series

- PseudoWire Attachment Circuit subinterfaces do not inherit the custom CosToTC map configured on the parent physical interface.

Workaround: Configure all subinterface AC explicitly with the custom CosToTc map. (928486)

- When mirror-on-drop packet-buffer drops are mirrored, if the dropped packet was originally destined to a non-shaped subinterface, the egress interface reported will be the parent interface. (931518)
- In case of EEDB hardware resource overflow, the switch might incorrectly acknowledge the corresponding MPLS route or Nexthop Group as programmed successfully in hardware. (944356)
- When hardware accelerated sFlow is enabled, the UDP checksum field of sFlow datagrams is set to 0. (949490)
- For ports 1-12, the secondary port will be errdisabled instead of inactive when a 200G transceiver is inserted into the corresponding primary port. (951503)
SKUs Affected: DCS-7280CR3A-32S, DCS-7280CR3A-32S-F, DCS-7280CR3A-32S-R, DCS-7280CR3AK-32S, DCS-7280CR3AK-32S-F, DCS-7280CR3AK-32S-R, DCS-7280CR3AM-32S, DCS-7280CR3AM-32S-F and DCS-7280CR3AM-32S-R
- Egress L2 subinterface counters include packets that are filtered by the split-horizon group feature. (953456)
- When a split horizon group is applied to an L2 subinterface used in an active/active multihoming MPLS-EVPN service, only BUM traffic will be filtered by the split horizon rules. Unicast traffic between members of the same split horizon group will not be filtered. (953805)
- When using the mirror-on-drop packet buffer drop reporting feature, if for a given forwarding chip and core there are several high frequency packet buffer drops with different drop reasons occurring, or if there are high frequency packet buffer drops occurring across multiple interfaces connected to that forwarding chip and core, then lower frequency packet buffer drops may not be mirrored. (957516)

- If the header removal feature is configured in a monitor session with a GRE tunnel destination, the metadata field in the GRE header will be set to 0. (963970)

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

- Untagged traffic entering VPWS port-based connectors are always classified as traffic-class 1, and ignores any QoS mapping configurations. (967855)
- Packets entering a VLAN interface via an L2 subinterface will not be matched by a traffic-policy applied on the VLAN interface. (974325)
- Packets entering a VLAN via an L2 subinterface will not be matched by segment security. (974329)
- In Tap Aggregation mode, the Mirroring with Header Removal feature is not supported in any monitor sessions where the destination is also a tool or tap-tool port. (974452)

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

- For hardware accelerated sFlow, flow samples ending in the pattern 0xFFFFFFFF may be truncated by 4B. (989032)
- When VRF Selection Policy is applied to an interface along with another feature like PBR, it is possible that a flap of the interface can lead to longer re-install times of the VRF Selection Policy resulting in a small window of FIB routing. (1001055)
- The scale of VTEP decap counters is limited to 1365 VTEPs. (1015265)
- On multi-chip systems, some MAC entries may prematurely age out resulting in unexpected flooding until these entries are relearned. (1030082)
- Redirect action is not support for BGP Flowspec polices applied to a L2 interface. (1037957)
- Due to hardware limitations, it is not possible to mirror the entire VXLAN packet when VXLAN is enabled on the switch. (1044064)
- For hardware-accelerated sFlow, if multiple agent IP addresses are configured for different VRFs, only one of the configured addresses will be used on the sFlow datagrams regardless of the VRF from which the datagrams is sent out. The selected agent IP address will be shown in the output of "show sflow hardware status". (1047309)
- RFC2544 test fails due to VOQ latency based packet drops for packet generated from SAT.
Workaround : disable VOQ latency based drop using CLI knob "platform sand VOQ tail-drop latency disabled" (1048704)
- When the system enters the ECMP or FEC overflow state, and certain ECMP next hops are programmed in hardware as transient non-ECMP, any change in the

next hop resolution of an ECMP adjacency (initially programmed as ECMP in hardware) may cause it to be reprogrammed as a transient non-ECMP. Consequently, the ECMP FEC previously associated with that adjacency may be reassigned to another next hop that was originally programmed as a transient non-ECMP adjacency in hardware. (1079393)

- At 200G or 400G speeds, the FEC degraded SER indication may be transmitted to the peer. The link otherwise operates normally. (1101261)
SKUs Affected: 7500R3-24D-LC, 7500R3-24P-LC, 7500R3-48Y4D-LC, 7500R3K-48Y4D-LC, 7800R3-36D-LC, 7800R3-36P-LC, DCS-7280CR3-32D4-F, DCS-7280CR3-32D4-M-F, DCS-7280CR3-32D4-M-R, DCS-7280CR3-32D4-R, DCS-7280CR3-32P4-F, DCS-7280CR3-32P4-M-F, DCS-7280CR3-32P4-M-R, DCS-7280CR3-32P4-R, DCS-7280CR3-36S-F, DCS-7280CR3-36S-R, DCS-7280CR3K-32D4-F, DCS-7280CR3K-32D4-R, DCS-7280CR3K-32D4A-F, DCS-7280CR3K-32D4A-R, DCS-7280CR3K-32P4-F, DCS-7280CR3K-32P4-R, DCS-7280CR3K-32P4A-F, DCS-7280CR3K-32P4A-R, DCS-7280CR3MK-32D4-F, DCS-7280CR3MK-32D4-R, DCS-7280CR3MK-32D4A-S-F, DCS-7280CR3MK-32D4A-S-R, DCS-7280CR3MK-32D4S-F, DCS-7280CR3MK-32D4S-R, DCS-7280CR3MK-32P4-F, DCS-7280CR3MK-32P4-R, DCS-7280CR3MK-32P4S-F, DCS-7280CR3MK-32P4S-R, DCS-7280DR3-24-F, DCS-7280DR3-24-M-F, DCS-7280DR3K-24-F, DCS-7280PR3-24-F, DCS-7280PR3-24-M-F and DCS-7280PR3K-24-F
- When using Hardware Accelerated Sflow, the uptime in sFlow datagrams start with a random value. (1150365)
- If the 'drop-pseudowire' action is configured in the TCAM profile for the traffic-policy feature and applied to non-pseudowire interfaces, it will drop L2 protocol packets like micro BFD that is not normally subject to filtering by traffic policies. (1153002)
- Link Qualification testing is not supported on a generator interface with the SFP-10G-MRA-T rate-adapting PHY running at a speed lower than 10Gbps. (1165217)
- VOQ counters for a shaped port-channel subinterface gets cleared when its anchor interface changes. This may occur when a rebalancing event occurs on that port-channel. (1173157)
- The "redistribute igmp" CLI command is required in MAC VRF VLAN aware bundle configuration is required on EVPM OISM multihoming setups. (1174359)
- Traffic-policy applied on psuedowires configured with the 'drop-pseudowire' action in the TCAM profile can take precedence over L2 protocol forwarding. (1187657)
- MAC entries do not age out in the MAC address table if continuous traffic is seen from these learned MAC addresses, even in the presence of a deny ACL denying these specific learned MAC flows. (1187950)

- Backup entries for MPLS-over-GUEv6 nexthop groups tunnels is not supported. (1199917)
- Egress sFlow does not report the ingress interface for flooded traffic. (1262066)
- Due to hardware limitations, egress traffic policy and uRPF cannot be configured together. (1287860)
- When IPv4 and/or IPv6 FlexRoute optimization is enabled on multiple VRFs, a high number of common prefixes across these VRFs can lead to a hardware overflow in the LEM table. This overflow may prevent some prefixes from being installed in the hardware and can cause persistent traffic loss for the affected routes.

LEM reservation configuration may be used as a workaround to avoid persistent traffic loss due to physical LEM table overflow. The workaround involves first determining the percentage of LEM table size at which overflow is seen and then

use the LEM reservation command to reserve LEM entries for one or more dummy VRFs such that the available unreserved LEM space reduces to below the overflow

detected size. With this, LEM table full is detected before overflow and routes

are programmed in LPM without leading to hardware overflow of the LEM table. (1290255)

7300X3, 7358, 7368, 7388, CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-720XP, CCS-722XPM, CCS-750, DCS-7010TX, DCS-7050X3, DCS-7050X4, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7060X5, DCS-7060X6, DCS-7260X3 and DCS-7300X Series

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

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

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

- QoS shaping and guaranteed bandwidth will only work in store-and-forward mode. They are not supported in cut-through mode. This includes 'shape rate X' and the 'bandwidth guaranteed X' commands. (56790)
- Accelerated Software Upgrade (ASU) from 4.13.x (x < 6) to 4.13.6 and above requires a special upgrade procedure due to a software image format change. (90097)
- 10 Gbps packet replication on all SFP+ ports at the same time may lead to packet discard. (91149)
- Egress L2 MTU violations in cut-through mode may result in unexpected discards of other frames. (95577)
- VXLAN encapsulated packets cannot be TX mirrored at the egress properly. The mirrored packets will have an incorrect MAC header. (97272)
- When IGMP snooping is enabled a single unknown multicast floodset is shared across all VLANs with IGMP snooping enabled. This can cause unexpected flooding to a trunk port that has no multicast routers on a given VLAN, but has multicast router attached to the same trunk port on another VLAN. (105188)
- When "hardware access-list resource sharing vlan in" is configured, Link Local Multicast traffic will not be intercepted by the ACL attached to the VLAN interface. (105504)
- There may be a momentary traffic disruption every time the nexthop specified in a PBR policy resolves to a new destination: affected packets may be forwarded by normal L3 lookup, or be dropped.

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

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

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

- If IP-in-IP decap groups are configured, "qos trust dscp" configuration is not supported on traffic matching the decap groups. (107579)
- Vxlan and MPLS features may not be configured at the same time. (109330)
- When "hardware access-list resource sharing vlan in" is configured and TCAM is close to being full, restart of the forwarding agent or a reboot of the box, can cause some ACLs to not get programmed properly. (109876)
- The Accelerated Server Upgrade feature is not supported with the "hardware access-list resource sharing vlan in" configuration. (110114)
- IGMP snooping is not supported on Vxlan enabled VLANs. The multicast packets will always be flooded locally and to remote VTEPs. (110479)
- MPLS is not supported on subinterfaces. (113110)
- Packets being VXLAN encapsulated are constrained to a single underlay physical nexthop per physical egress port.

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

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

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

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

- L2 flooding of BUM packets on Vxlan vlans uses L3 replication tables, which are also used by IP multicast routing. The replication tables have an entry for each vlan, physical port combination (lag or non lag). We will run out of this table resources if we need more than 64k entries. (114517)
- If the switch is in cut-through mode, L2 MTU violations will not be counted if the packet headers are changed while forwarding. (116760)
- QoS policies are not supported for L3 Unicast flooded packets (routed packets with nexthop MAC not learned) (123883)

- When a PBR is attached to any interface, Layer-3 packets which are getting flooded in the egress VLAN will not be treated with Egress Router ACL if one is configured on that egress VLAN. (133144)
- VXLAN routing is not supported on systems with BCM56750 fabric chips. Such switches can be identified using the CLI command "show platform trident l3 summary", which will show bcm56750 for "Fabric chip model". (134625)
- Mirroring of incoming multicast traffic to a GRE tunnel may result in data traffic loss on the egress interface in an oversubscription scenario. (139591)
- Egress IP/IPv6 router ACL is not supported for VXLAN encapsulated traffic. (148642)
Series Affected: DCS-7050X3 and DCS-7300X Series
- IP ACLs configured on SVIs to match routed IP multicast traffic may also match bridged IP multicast traffic in that SVI. (152523)
- Up to 4 mirroring sessions are supported.

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

- 1) Multicast traffic matching reserved IPv4 multicast address range 224.0.0.X is not counted by L3 interface counters.
2) L3 interface counters not supported for IPv6 multicast traffic.
3) Unicast traffic to CPU is counted by L3 interface counters only if routing is enabled (i.e. ip routing for IPv4 traffic and ipv6 unicast-routing for IPv6 traffic) (160930)
- Mirroring of Vxlan encapsulated packets in the egress direction is not supported. (167189)
- If a trunk port is configured as a member of both the underlay and overlay VLANs, the following limitations apply:
 - If packets received on the trunk port are flooded in the overlay VLAN, the VxLAN encapsulated copy is not forwarded on the ingress trunk port.
 - Vxlan encapsulated packets received on the trunk port, if flooded in the overlay VLAN after decapsulation, are not forwarded on the ingress trunk port. (185321)
Series Affected: DCS-7050CX3, DCS-7050SX3 and DCS-7300X Series
- If ingress ACL is applied on an interface, traffic sent to CPU from that interface is falsely counted as InAclDrop in "show int counters ingress acl drop". (203465)

- Clause 37 auto-negotiation on SFP28 and QSFP28 ports is not supported. A workaround is to use 1G forced speed configuration on the affected ports. (206778)
Series Affected: 7300X3, CCS-720XP, CCS-750, DCS-7050CX3, DCS-7050SX3, DCS-7050TX3, DCS-7050X3, DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X, DCS-7060X2 and DCS-7260CX3 Series
- When "reload fast-boot" is performed from a release prior to 4.17.2, switches with BGP peering connections with the switch undergoing "reload fast-boot" could experience BGP KeepAlive timeouts before they see port flaps. This could result in longer than expected data plane downtime during the upgrade. (214700)
Series Affected: DCS-7060X and DCS-7060X2 Series
- A multi-bit ECC error in forwarding ASIC packet memory cannot be automatically corrected. A hitful restart of the slice agent managing the forwarding ASIC or a reboot of the system is required to clear the condition. (214950)
- If ip-ttl is the only enabled field for port channel hashing, after reload hitless, packets may get hashed to a different link than before the reload. (223568)
Series Affected: DCS-7060X Series
- A configured mirroring ACL may not be applied in hardware. This can be confirmed by checking the output of 'show platform trident tcam mirror'. (236908)
- When a VXLAN encapsulated packet is received on a VXLAN VLAN where PBR policy is applied, the switch will not PBR forward the decapsulated packet. (244093)
- Sflow output interface determination will not work when the output interface is a subinterface. (250424)
- Cut-through mode is not supported on SFP+ interfaces. (252731)
Series Affected: DCS-7050X3, DCS-7050X4, DCS-7060X4, DCS-7060X5 and DCS-7260X3 Series
- A switch may not be able to bridge or route VXLAN traffic if it hits one or all of following conditions:
 - Nexthop resources are exhausted as indicated by syslog `VXLAN_HWHWER_NEXTHOP_RESOURCE_FULL`.
 - MAC table overflow as indicated by `show platform trident mac-address-table` missing.

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

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

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

SKUs Affected: DCS-7050SX3-48YC8-F, DCS-7050SX3-48YC8-R, DCS-7050TX3-48C8-F and DCS-7050TX3-48C8-R

- When MLAG peer MAC routing is enabled, OSPF neighborship over VXLAN overlay may never get established. (349242)

- Ip-length based rules are not supported on Egress ACLs. (353247)
- Packets dropped in the egress pipeline may still be egress mirrored successfully. (357609)
- Routed unicast traffic egressing a SVI for which destination MAC is not learnt are sent to CPU and then flooded on the egress VLAN. (368228)
Series Affected: 7358, 7368, 7388, DCS-7050X4, DCS-7060DX4, DCS-7060PX4, DCS-7060X4 and DCS-7060X5 Series
- Ports with ingress MAC ACL still learn MAC addresses although traffic may be dropped (400211)
- Packets originated from the CPU are not included in egress SVI/ subinterface counters. The CPU originated packet/octet count can be obtained from the switch bash shell using the linux ifconfig command. (400885)
- Reload hitless from EOS-4.22 or earlier results in /2 and /4 ports on Ethernet 1-12, 21-44 and 53-64 to remain inactive when /1 and /3 are configured in 10G or 25G. In order to activate /2 and /4 ports, configure 40G, 50G or 100G /1, then configure 10G or 25G back. These ports are also activated following a forwarding agent restart. (409057)
Series Affected: DCS-7260CX3 Series
- "reload fast-boot" is not supported with virtual IP address config. (411323)
- A 10G or 25G SFP transceiver plugged into a QSFP port using a QSFP-SFP Adapter (QSA) reserves 4 logical ports. (415439)
SKUs Affected: DCS-7050SX3-48C8, DCS-7050SX3-48YC8, DCS-7050TX3-48C8, DCS-7260CX3-64 and DCS-7260CX3-64E
- Speed change on a port with macsec enabled is not supported and could cause the port to get into a persistently un-authenticated state. Workaround is to disable macsec on the port, change the speed and then re-enable macsec on the port. (415737)
SKUs Affected: DCS-7050CX3M-32S-F and DCS-7050CX3M-32S-R
- When a VXLAN VTEP is reached through a resilient ECMP route, (VXLAN bridging or routing) traffic to that VTEP may get dropped in HW. Workaround is to remove the Eesilient ECMP config. (421216)
- Performing SSU from any release prior to 4.22.2 can lead to extended outage and mis-forwarding. (424812)
SKUs Affected: DCS-7050SX3-48YC8-F and DCS-7050SX3-48YC8-R
- IP ACL match over MPLS Encapsulated Packets is not supported. (426413)
- During Mpls PHP operation inner payload type gets ignored. (428546)

- When both "hardware counter feature vni decap" and "hardware counter feature vni encap" are enabled, the egress flexcounter pools will only be released after both features are disabled. (441887)
Series Affected: 7300X3, 7368, CCS-720XP, DCS-7050X3, DCS-7060X, DCS-7060X2, DCS-7060X4, DCS-7260X3 and DCS-7300X Series
- When 'hardware nat static access-list resource sharing' is in the switch's config, and a static NAT filter ACL and a security ACL have overlapping rules, the security ACL's hit counter might not increment in all cases. Regardless of the hit counter, both the security ACL rule and the NAT filter ACL rule will be applied correctly.

Workaround is to disable static NAT ACL resource sharing with:
'no hardware nat static access-list resource sharing'
Alternative workaround is to disable port ACL in VFP with:
'platform trident tcam port-acl vfp disabled' (450080)
- When an interface undergoes a hitless speed change, sFlow sample numbers and sequence counts for that interface may be reset. (452949)
Series Affected: 7300X3, 7368, DCS-7050X3, DCS-7060X4 and DCS-7260CX3 Series
- BUM traffic is not sampled by egress sFlow. (459902)
- If the incoming packet is VLAN tagged, then the truncated mirrored packet will be 4 bytes shorter than expected. (463986)
Series Affected: 7300X3, 7368, CCS-720XP, DCS-7050X3 and DCS-7060X4 Series
- If virtual mac address of the switch is learnt against a front panel port, it may not age out. The workaround is to clear the mac address manually by issuing 'clear mac address-table dynamic vlan <VLAN> address <ADDRESS>' command. (470965)
- When IPV4 and IPV6 PBR are both programmed in TCAM, forwarding plane agent restart may result in IPV6 PBR rules not being programmed in hardware. (473582)
- By default, configuring a static NAT rule with an empty access-list is treated the same as no access-list: for source NAT, the switch will ignore the destination IP, and vice versa for destination NAT.
With "hardware nat static access-list resource sharing" enabled, configuring a static NAT rule with an empty access-list is treated the same as an undefined access-list: some hardware resources may be consumed, but the traffic will be forwarded untranslated. (473783)
Series Affected: CCS-720XP, DCS-7050CX3 and DCS-7050X3 Series

- The total length field in the IP header is not populated correctly when operating in cut-through mode. The switch should be operated in store-and-forward mode when mirroring to a GRE tunnel. (475273)
- Mirrored packets that violate the MTU requirements will not be fragmented. (481513)
- When persistent port security is enabled on an interface, MAC addresses will persist when an interface is down. While the interface remains down, traffic destined to these MAC addresses will be dropped.

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

- If dynamic load balancing is enabled for an ECMP group, it's member link flap might cause poor traffic distribution. Traffic distribution may improve over time. (487595)

Series Affected: DCS-7060X4 and DCS-7260X3 Series

- If dynamic load balancing is enabled for ECMP groups while the traffic is flowing, then the load on the ECMP members may not be distributed uniformly. Traffic distribution may improve over time. (490418)

Series Affected: DCS-7060X4 and DCS-7260X3 Series

- When dynamic load balancing is enabled, If multiple traffic flows destined to the same ECMP group arrive at the same time, then the load on the ECMP members may not be distributed uniformly. Traffic distribution may improve over time. (490716)

Series Affected: DCS-7060X4 and DCS-7260X3 Series

- Packets sent from the CPU cannot be egress mirrored unless they are being sent out of an SVI. (495146)
- If the speed of an interface changes with traffic running, it is possible for "show interface counters" and "show interface counters fast-poll" to diverge on that interface. (502252)
- Mirror on drop does not mirror dynamic NAT traffic drops. (506047)
- Egress mirroring of multicast packets that are going to be encapsulated on the switch, independent of mirroring, is not supported. (512880)
- Mirror on drop does not mirror VXLAN IPv6 underlay traffic drops. (516212)
- Issuing "reboot -f" from the shell can cause continuous machine check errors, requiring a power cycle.

Instead, use the "reload" CLI command, or "reboot" in bash. (525314)

SKUs Affected: DCS-7010TX-48, DCS-7010TX-48-DC, DCS-7010TX-48-DC-F, DCS-7010TX-48-DC-R, DCS-7010TX-48-F and DCS-7010TX-48-R

- The mirroring logic can only capture and mirror packets seen by the ingress pipeline. Packets generated by the memory-management unit do not traverse the ingress pipeline and hence cannot be mirrored. This applies to all types of mirroring. (527668)
- When policer is attached to the class-default of an egress policy-map consisting of both IPV4 and IPV6 ACLs, IPV4 and IPV6 traffic hitting class-default would be policed individually with rate configured on default class-map (541843)
- L3 MTU configuration is not honored on a NAT enabled interface for multicast traffic. (546193)
 SKUs Affected: CCS-720XP-24Y6, CCS-720XP-24ZY4, CCS-720XP-48Y6, CCS-720XP-48ZC2, DCS-7050CX3-32S, DCS-7050CX3M-32S, DCS-7050SX3-48C8, DCS-7050SX3-48YC12, DCS-7050SX3-48YC8 and DCS-7050TX3-48C8
- The Strata agent restarts after running "platform trident diag psr <port>" command. (546588)
- Routing between VXLAN VLANs with VRRP is not supported. (548160)
 Series Affected: DCS-7060CX, DCS-7060CX2, DCS-7060SX2, DCS-7060X and DCS-7060X2 Series
- The traffic load across ports comprising a port channel(s) may be uneven for certain traffic patterns and/or port channel configuration and state.

Workaround:

- change the seed in the hashing algorithm until satisfactory result is achieved by executing configuration CLI command:

"arista(config)# port-channel load-balance trident <seed>" (555833)

- Enabling L2 protocol forwarding of LLDP on a given port will implicitly enable forwarding of EAPOL and MKA packets on the same. (569478)
- Configure replace operation is not supported with TCAM profile configuration. Changing TCAM profile with configure-replace may cause an unexpected restart of forwarding agent(s) if TCAM based features like ACL, policy QoS are also part of the configuration change. (571912)
- Deny action is not supported on ACLs configured for filtering based packet sampling used for mirroring packets to the collector. (582014)
- After performing SSU from a PTP unsupported release to a PTP supported release, PTP will not be functional.
 The workaround is to reset the forwarding switch by issuing the CLI command "platform trident reset" after the SSU and prior to enabling PTP. This will cause all interfaces to flap and a brief traffic outage to occur. (582970)
 Series Affected: DCS-7060X4 Series

- "no hardware access-list update default-result permit" is not supported. (585298)
SKUs Affected: CCS-755-X3-SC
- BGP packets are ignored by IP counters if hardware IP counters are not enabled. (599275)
- When Port Security configuration is removed from an interface while StrataL2 agent is initializing , we can end up in a state where some MAC addresses on the interface do not get aged out and can also not be cleared via configuration.
A workaround is to flap the interface to remove the MAC addresses. (635439)
- Packets egress mirrored to a Greenspan destination with subinterface nexthops may have the internal VLAN in the inner VLAN portion of the packet. (642351)
- The 32-bit image of EOS does not support the 7050X4, 7358X4, 7060X5, and 7388X5 series switches. Use the 64-bit version instead (646587)
- When MACsec encryption traffic spans across multiple queues encrypted packets may be sent out of order. (669251)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- The outPktCtrl MACsec counters will not increment. (671159)
Series Affected: CCS-750 Series
- When the packets matching FBPS rule are also marked for drop due to other ACL rules, the packets are not mirrored and dropped in hardware. Workaround is to remove the configuration that results in dropping these packets. (673800)
SKUs Affected: CCS-720XP-48ZC2-F
- Mirroring packets to CPU can cause congestion on CPU queues and disrupt control plane traffic. It is not recommended to leave this feature on for extended periods of time. (675773)
- When attempting to VLAN mirror packets ingressing a VXLAN edge VLAN, the VXLAN encapsulated packets that ingress the core interface and are destined to be decapsulated to the edge VLAN will not be mirrored. (679353)
Series Affected: DCS-7260CX3 Series
- The maximum amount of VLAN sources allowed in a VLAN mirror session depends on the configuration's TCAM usage and the number of VLAN mirror sessions configured. A maximum of up to 3 mirror sessions are supported. Using a configuration that utilizes more VLAN sources than supported will result in VLAN mirroring sessions not being programmed. (680940)
Series Affected: CCS-710P, CCS-720XP, DCS-7010TX, DCS-7050CX3, DCS-7050SX3 and DCS-7050TX3 Series

- DirectFlow redirecting to interface with 'switchport source-interface tx' configured is unsupported. (681960)
- The outPktsUntagged counter will increment for unencrypted traffic egressing an interface even when the interface is not MACsec enabled. (682111)
Series Affected: CCS-750 Series
- The inPktsBadTag MACsec counter will not increment. Instead the drops will be accounted under the inPktsNotValid MACsec counter. (695748)
Series Affected: CCS-750 Series
- A chassis may only contain all 7300X3 or all 7300X linecards. Combining 7300X3 or 7300X linecards may result in repeated restarts of the forwarding agents, and links may stay down. Power cycle the system after removing the incompatible cards to recover. (699880)
Series Affected: DCS-7300X and DCS-7300X3 Series
- If "dst-ip" key-field is removed from "acl vlan ip ingress" feature in TCAM profile then, implicit link-local-multicast TCAM entry will not be programmed for ACL applied to SVI in unshared mode.

Workaround:

To overcome this issue add "dst-ip" key-field to "acl vlan ip ingress" feature in TCAM profile. (704466)

Series Affected: 7300X3, 7368, CCS-710P, CCS-720XP, CCS-750, DCS-7050CX3, DCS-7050SX3, DCS-7060DX4, DCS-7060PX4 and DCS-7260CX3 Series

- 7300-SUP/7300-SUP-D and 7300-SUP2-D supervisors are incompatible with each other and cannot be mixed in the same chassis. (728201)
SKUs Affected: 7300-SUP, 7300-SUP-D and 7300-SUP2-D
- The CLI "show hardware resource IpostLagLagMember" displays current usage statistics of hardware resources.
That hardware table may appear partially used even after some/all LAG(s) that were active before have been deleted. This is expected behaviour with no other adverse effects. (729756)
- In PTP boundary or generalized PTP mode, configuring the 'ptp forward-unicast' feature with interfaces using IPv6 transport is not supported. A workaround would be to use the IPv4 or L2 transport. (730817)
- If any currently unsupported mirroring configuration options such as ACLs, truncation, multiple destinations, or rate limits are entered in such a way that bypasses the CLI guards (such as config replace or editing the startup configuration), these options may be displayed as configured and in-use but will not be in effect.

Workaround: Remove any unsupported mirroring options from the startup configuration and/or un-configure them via the CLI. (731708)

- In cases of failed autonegotiated link up, link training may report as "off" on the primary lane. This is cosmetic. (731948)
- Configuring PTP MACsec bypass using the per-profile "ptp bypass" command is not supported.

Enable the "ptp bypass" command in MAC security configuration mode to have PTP packets bypass MACsec processing. (755578)

SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8, CCS-722XPM-48ZY8-F, CCS-755-CH and CCS-758-CH

- Only 100G copper downlink and AOC uplinks are supported on this Reverse Airflow platform. (758569)
SKUs Affected: DCS-7060DX5-64S-R
- When the switch is configured in ALPM mode and when the default route is forwarded over a VXLAN tunnel, L3 agent forwarding restart may result in some routes not getting programmed in the hardware. There is no workaround for this issue (764586)
- Per interface L3 MTU CLI configuration is not supported.

Workaround is to use global L3 MTU configuration in "interface defaults" mode. (777699)

Series Affected: CCS-710P, CCS-720DF, CCS-720DP, CCS-720DT, CCS-722XPM, CCS-750 and DCS-7010TX Series

- DLB can only be enabled on first 127 ECMP groups, that is ECMP groups that have ID between 1-127. (779882)
SKUs Affected: DCS-7050DX4-32S-F, DCS-7050PX4-32S-F, DCS-7060DX5-64-F, DCS-7060DX5-64S-F, DCS-7060DX5-64S-R and DCS-7060PX5-64E-F
- If the same ACL is applied on control plane and data plane, the byte counter for the ACL will be incremented for only data plane packet hits while the packet counter will be incremented for both data plane and control plane packet hits.

A workaround is to create different ACLs to apply on control plane and data plane. (788296)

- Mirroring to CPU dataplane traffic using a non-default traffic-class will use the CPU queue of that traffic class, instead of the "Other" CPU queue. Workaround is to mirror to port instead of CPU destination. (791599)
Series Affected: 7368, DCS-7060X and DCS-7060X2 Series
- Counters / statistics for egress ACL(s) applied to L3 port(s) do not account for fragmented packets. (798886)

- A feature using TCAM resources will not be able to use the last entry in the last slice allocated to it. For example, if a TCAM slice has 512 entries free, and if there are 2 unused slices, a feature using these slices, will be able to program only $(512 \times 2) - 1 = 1023$ entries as the last entry in last slice can not be used. (815178)

Series Affected: 7388 Series

SKUs Affected: 7358X4-SC, 7388X5-SC, DCS-7050CX4-24D8, DCS-7050CX4-24D8-F, DCS-7050CX4-24D8-R, DCS-7050DX4-32S, DCS-7050DX4-32S-F, DCS-7050DX4M-32S, DCS-7050DX4M-32S-F, DCS-7050PX4-32S, DCS-7050PX4-32S-F, DCS-7050SDX4-48D8, DCS-7050SDX4-48D8-F, DCS-7050SDX4-48D8-R, DCS-7050SPX4-48D8, DCS-7050SPX4-48D8-F, DCS-7050SPX4-48D8-R, DCS-7060DX5-64, DCS-7060DX5-64-F, DCS-7060DX5-64S, DCS-7060DX5-64S-F, DCS-7060DX5-64S-R, DCS-7060PX5-64E and DCS-7060PX5-64E-F

- This supervisor module is only supported with 7300X3 linecards. (826371)
SKUs Affected: 7300-SUP2-D
- ACL with deny and log rule which is applied to an Ethernet or LAG interface will not syslog the packet(s) if the packet also matches the ACL applied to its sub-interface. (830841)

- The 64-bit image of EOS does not support the 7010T series of switches. Use the 32-bit version instead (846011)

- The outPktCtrl MACsec counters will not increment. (924462)

SKUs Affected: DCS-7050CX4M-48D8, DCS-7050CX4M-48D8-F and DCS-7050CX4M-48D8-R

- MTU configured on the egress interface is not enforced for the routed IP multicast traffic. This affects both regular IP multicast traffic as well as multicast traffic routed through tunnels. (938808)
- Packets destined to an L3 interface in a VRF that does not have routing enabled, are copied to the CPU on the `other` queue instead of the `self-ip` queue. This can potentially lead to packet loss for sessions on that L3 interface.

Workaround is to enable IP routing on that VRF. (938852)

- If a custom TCAM profile is applied which does not have "l4-src-port" key-field for feature "acl port ipv6 ingress" then, L4SrcPort missing syslog is seen when "hardware access-list ipv6 implicit-permit icmpv6 neighbor-discovery" is also configured.

Workaround:

"l4-src-port" should be part of qset for feature "acl port ipv6 ingress" in TCAM profile when security ACLs need to allow IPv6 neighbor-discovery.

(969034)

- When switch is configured with multi VTEP MLAG, all traffic that are head end replicated will use the MLAG shared IP address. (989109)
- With DLB enabled and Strata-FixedSystem agent shutdown, link flaps from peer will cause complete traffic drops due to DLB being non functional (1000399)
- Ingress Filtered Mirroring with Inner VLAN ID is not supported (1003362)
- In case of CPU queue exhaustion, some features in startup config may not work due to CPU queue unavailability in a deterministic way. (1043368)
- MSS-G and common VARP MAC configuration in the network may result in traffic drops on L2 VLANs.

Workaround: Configure different VARP MAC on affected devices with L2 VLANs. (1045376)

- Port-channel negotiating LACP goes down when egress MAC ACL is applied on the port-channel interface.
To workaround this issue either the ACL can be removed from the interface or, a permit LACP rule can be added to the end of the ACL. (1051026)
- QoS policy-map having two rate three color (trTCM) policers does not support red marked packet count and displays value as 0 (1108223)
- Ports with the speed set to 10Mbps do not recirculate packets when they are used in Recirc Channels. (1192725)
- Egress security ACLs does not support UDF filtering. Configuring IP access-list with UDF filtering on egress direction will skip UDF data/mask programming. Rest of the filters will be programmed as expected. (1207111)
- Storm control drop counters do not distinguish between data plane and control plane drop counts. They share the same counters. (1301580)

CCS-750 Series

- Mirroring PTP packets when PTP is enabled is not supported. (738080)
- PTP two-step E2E and P2P transparent clock mode is not supported. (738097)
- Dropping PTP management message using the per-interface `ptp management drop` configuration is only supported in boundary clock mode. (752342)
- When PTP is configured, PTPv1 packet forwarding is not supported. (752367)
- Switches acting as either a Boundary Clock or Transparent Clock, downstream devices may see higher jitter in PTP time of up to +/- 1000ns (1025195)
SKUs Affected: CCS-755-CH and CCS-758-CH

7300X3 and DCS-7060X Series

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

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

- Cut-through mode is not supported on the SFP+ interfaces. (182728)
- Cut-through mode is not supported at 1G speed. (219427)
- Mirror To GRE destinations with multiple ECMP nexthops or LAG is not supported. A single link of the LAG or ECMP nexthops will receive all mirror traffic. (639980)
Series Affected: 7368 and 7388 Series
SKUs Affected: DCS-7060DX4-32 and DCS-7060PX4-32
- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- VXLAN encapsulated traffic UDP source port is generated in the range 1-65535 instead of 49152-65535. (1002498) (1)

DCS-7060X2 Series

- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- VXLAN encapsulated traffic UDP source port is generated in the range 1-65535 instead of 49152-65535. (1002498) (1)

DCS-7260X3 Series

- Configuring 100G, 50G, 40G, 25G and 10G simultaneously might result in unexpected packet forwarding behavior. (235075)
- Load Balance Number based hashing is not supported with L3 EVPN. (837409)
- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- Packets dropped by ACL rules are included in the count performed by the ECMP monitor snapshot feature. (991089) (1)
- Enabling the MAC hash collision detection feature caused an increase in CPU utilization for agents Strata and StrataL2 under linerate traffic. (996460)
SKUs Affected: DCS-7260CX3-64-F, DCS-7260CX3-64-R, DCS-7260CX3-64E-F, DCS-7260CX3-64E-R, DCS-7260CX3-64LQ-F and DCS-7260CX3-64LQ-R

- VXLAN encapsulated traffic UDP source port is generated in the range 1-65535 instead of 49152-65535. (1002498) (1)
- AGM snapshots only reflect ECMP member hashing before CLB redirection. (1193385) (1)

7368 and DCS-7060X4 Series

- Multicast replication resources in TH3 are only 10% of their unicast counterparts. Due to this multicast replication in TH3 is limited to 1.2 Tbps. (297317)
- Layer2 source MAC learning may occur on packets received with CRC errors. (325507)
- When 'hardware counter feature gre tunnel interface in' feature is enabled, the gre tunnel interface decap counters are not incremented when subinterface is used as the ingress interface. (464972)
- Per vlan routing and sub interface are not supported together. (486439)
- 3m QSFP-DD passive copper cables are not supported in ports 1-4 and 29-32 (639613)
SKUs Affected: DCS-7060DX4-32
- Double Dot1Q tagged packets with IP payload may not be classified as IP packets in hardware. (813027)
- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- IP Source Guard is not supported on interfaces that are configured as Dot1q tunnels. If the QinQ packets need to be dropped, user may configure the command "switchport vlan tag validation" on the interface. (1032292)
- When a device behind phone/hub/switch connected to an Arista switch is authenticated via Dot1x MAC Based Authentication, it cannot be moved to different port if the other port already has a different supplicant authenticated on it.

A workaround is to configure 'dot1x timeout idle-host' (1066149)

DCS-7060X5 Series

- Ports 1-8 and 25-32 do not support 400G-4, 200G-2, and 100G-1 speeds with copper (CR) transceivers. (774664)
SKUs Affected: DCS-7060DX5-64, DCS-7060DX5-64-F, DCS-7060PX5-64E and DCS-7060PX5-64E-F

- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- Port-channel sub-interface will consume more hardware resources than sub-interface on routed physical interfaces (1001437) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- Config replace does not bring subinterface to operation state which was previously failed due to hardware resource full. (1009696) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- VXLAN is not supported when system has sub-interfaces configured. (1015740) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F, DCS-7060DX5-64S-R and DCS-7060PX5-64E-F
- MPLS over L3 subinterfaces is not supported. (1029083) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R

7388 and DCS-7060X5 Series

- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- Port-channel sub-interface will consume more hardware resources than sub-interface on routed physical interfaces (1001437) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- Config replace does not bring subinterface to operation state which was previously failed due to hardware resource full. (1009696) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- VXLAN is not supported when system has sub-interfaces configured. (1015740) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R,

DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F, DCS-7060DX5-64S-R and DCS-7060PX5-64E-F

- MPLS over L3 subinterfaces is not supported. (1029083) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R

DCS-7060CX5 and DCS-7060DX5 Series

- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- Port-channel sub-interface will consume more hardware resources than sub-interface on routed physical interfaces (1001437) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- Config replace does not bring subinterface to operation state which was previously failed due to hardware resource full. (1009696) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- VXLAN is not supported when system has sub-interfaces configured. (1015740) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F, DCS-7060DX5-64S-R and DCS-7060PX5-64E-F
- MPLS over L3 subinterfaces is not supported. (1029083) (1)
Series Affected: 7388 Series
SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R

DCS-7060X6 Series

- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- Enable standalone link training or auto-negotiation if supported for better performance with 800G copper cables. (1039049) (1)
SKUs Affected: DCS-7060X6-64PE and DCS-7060X6-64PE-F
- Persistent port security MACs are cleared when the StrataL2 agent is restarted. (1089853) (1)
Series Affected: 7358 and DCS-7060X6 Series

SKUs Affected: 7358X4-SC, DCS-7050CX4-24D8, DCS-7050CX4M-48D8,
DCS-7050DX4-32S-F, DCS-7050DX4M-32S-F, DCS-7050PX4-32S-F and
DCS-7050SDX4-48D8

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

- Triple-wide TCAMs are not supported. (603246)
Series Affected: CCS-710P Series
- Generation of PFC frames is not supported. (699868) (1)
- The IPv6 VARP VTEP feature may not work after switch reload or Strata slice agent restart.
Configure "platform trident tcam port-acl vfp disabled" before configuring user-defined ACLs. This will prevent rearrangement of ACLs and exhaustion of TCAM slices used by VXLAN. (1005600) (1)

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

- XPN cipher suites are not supported. (590170)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- Non-unicast traffic egressing on Port-Channel interfaces that are MACsec enabled do not use hashing to load balance. A single Port-Channel member will be designated to transmit non-unicast traffic for a given replication group (VLAN or multicast group). (604539)
Series Affected: CCS-722XPM Series
- "show platform trident counters" output is misleading for certain counters with MACsec encryption enabled. The workaround is to use the "show interfaces counters discards" and "show mac security counters interface" commands where possible. (623564)
SKUs Affected: CCS-722XPM-48Y4, CCS-722XPM-48Y4-F, CCS-722XPM-48ZY8 and CCS-722XPM-48ZY8-F
- Generation of PFC frames is not supported. (699868) (1)
- PTP boundary clock mode is not supported when MACsec is enabled and the configuration "no ptp bypass" is turned on.
The workaround is to enable "ptp bypass" or to use interfaces without MACsec (755609)
Series Affected: CCS-722XPM Series
SKUs Affected: CCS-720XPM-48TH-6SY
- Default BGP Graceful Restart timeout of 300 seconds is insufficient.
Workaround is to configure this to 600 seconds via "graceful-restart restart-time 600" (796535)

- VLAN tagged MACsec is not supported (852797)
Series Affected: CCS-722XPM Series
- DirectFlow is not supported on MACsec enabled interfaces. (907033)
Series Affected: CCS-722XPM Series
- The IPv6 VARP VTEP feature may not work after switch reload or Strata slice agent restart.
Configure "platform trident tcam port-acl vfp disabled" before configuring user-defined ACLs. This will prevent rearrangement of ACLs and exhaustion of TCAM slices used by VXLAN. (1005600) (1)
- If "traffic unprotected allow" is configured, and MACsec falls back to unprotected traffic, any encrypted packets will be dropped. (1007567)
Series Affected: CCS-722XPM Series
- "forwarding l3 source-interface drop" command is not supported when MACsec is enabled on the port. Configuring it may result in undesirable double delivery of the packet. (1086542)
- When "no hardware access-list update default-result permit" is configured for an egress ACL on an interface and MACsec is configured on the same interface, traffic is not blocked while ACL is updated. (1212387)
Series Affected: CCS-722XPM Series
SKUs Affected: CCS-720XPM-48TH-6SY

CCS-720DF, CCS-720DP and CCS-720XP Series

- When flow group ACL rules are changed, flow group counters displayed in "show platform trident flow counters" will continue to have counters incremented for active flows corresponding to older flow groups till the flows are aged out. (552450)
- On interfaces with speeds less than 10G, PFC cannot be transmitted. (689087)
- The IPv6 VARP VTEP feature may not work after switch reload or Strata slice agent restart.
Configure "platform trident tcam port-acl vfp disabled" before configuring user-defined ACLs. This will prevent rearrangement of ACLs and exhaustion of TCAM slices used by VXLAN. (1005600) (1)

CCS-750 Series

- The switch is not capable of generating pause flowcontrol frames. (480688)
Series Affected: CCS-750 Series
- Per interface L3 MTU configuration is not supported.
The workaround is to use the global L3 MTU configuration command instead. (520825)

- The supervisor uplink SCD does not support hitless FPGA upgrade. (528417)
SKUs Affected: CCS-750-SUP100 and CCS-750-SUP25
- The switch is not capable of generating Priority Flow Control (PFC) frames. (555294)
Series Affected: CCS-750 Series
- Shaper below 1 Mbps is not supported. (605407)
- MACsec encrypted flow control packets and priority-flow-control packets that are received will not be visible in "show flowcontrol" or "show priority-flow-control counters". (670539)
- When a mirror session has a LAG member as the source, packets sent to the LAG will not be mirrored. (712186)
- TCAM profile does not effect any saving of TCAM slices / banks for Egress MAC ACL and IP ACL applied to Ethernet, LAG ports (723449)
SKUs Affected: CCS-755-X3-SC and CCS-758-X3-SC
- Ports on this linecard share a total sustainable bandwidth.

Though this linecard can absorb concurrent bursts received on all ports, the total sustained bandwidth is limited to a total of 148 Gbps for an average frame size of at least 276 Bytes.

To optimize sustained bandwidth availability, the recommendation is to distribute the bandwidth equally across the groups of eight ports. The groups are: Ethernet1-8, Ethernet9-16, Ethernet17-24, Ethernet25-32, Ethernet33-40, Ethernet41-48

For example, if there are 24x5G and 24x2.5G devices to connect, it is recommended that 4x5G and 4x2.5G be connected to each of the 6 groups of front-panel ports. (788129)

SKUs Affected: CCS-750X-48ZXP-LC

- 4 extra bytes per packet is seen in egress ACL byte counter for packets sent on down link ports. (792001)
- Stateful Switchover triggered by removing the active supervisor will cause a long data-plane outage for uplink traffic destined to the removed uplinks. Workaround is to unplug the uplink interfaces, wait for the link to go down and then remove the active supervisor. (839923)
Series Affected: CCS-750 Series
- When Greenspan session is configured with trunk port or sub-interface as next hop destination, the mirrored packets will have missing dot1q and VLAN tag. (979999)

DCS-7050X3 Series

- The IPv6 VARP VTEP feature may not work after switch reload or Strata slice agent restart.
Configure "platform trident tcam port-acl vfp disabled" before configuring user-defined ACLs. This will prevent rearrangement of ACLs and exhaustion of TCAM slices used by VXLAN. (1005600) (1)

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

- The egress queue and per hop latency of the INT terminator node are always shown as zero. (562000)
Series Affected: DCS-7050CX3 and DCS-7050SX3 Series
- The IPv6 VARP VTEP feature may not work after switch reload or Strata slice agent restart.
Configure "platform trident tcam port-acl vfp disabled" before configuring user-defined ACLs. This will prevent rearrangement of ACLs and exhaustion of TCAM slices used by VXLAN. (1005600) (1)

7358 and DCS-7050X4 Series

- Sampled counters are not required. (437955)
- Color based DSCP or CoS assignment is not supported. (679840)
- Packets destined for the control plane may appear to be dropped within the ingress pipeline. This will appear in the output of `show platform trident counters` as a RX drop, and is expected due to more drop events available in the pipeline. (685445)
- On interfaces Ethernet33 and Ethernet34, PFC cannot be transmitted. (737611)
SKUs Affected: DCS-7050DX4-32S, DCS-7050DX4-32S-F and DCS-7050PX4-32S-F
- Egress ACL does not match on policy rewritten DSCP nor remarked ECN values. (741424)
Series Affected: DCS-7050X4 Series
- Maximum replications for a non unicast packet is limited to 1023. (750699)
Series Affected: DCS-7050X4 Series
- Setting different actions for packets marked as green and yellow by two rate three color (trTCM) QoS policy-map does not work (768572) (1)
- A maximum of two QoS features can be configured on Trident4X platforms. (772245)
- Mirror to GRE with forwarding drop is not supported on T4.
Workaround is to use Mirror on Drop feature for capturing drop packets. (930242) (1)

- Mirror to GRE with forwarding drop is not supported on T4X11.
Workaround is to use Mirror on Drop feature for capturing drop packets.
(930244)
Series Affected: DCS-7050X4 Series
SKUs Affected: 7358X4-SC
- Persistent port security MACs are cleared when the StrataL2 agent is restarted. (1089853) (1)
Series Affected: 7358 and DCS-7060X6 Series
SKUs Affected: 7358X4-SC, DCS-7050CX4-24D8, DCS-7050CX4M-48D8, DCS-7050DX4-32S-F, DCS-7050DX4M-32S-F, DCS-7050PX4-32S-F and DCS-7050SDX4-48D8

DCS-7050X4 Series

- An interface that acts as a Dot1x supplicant cannot be paired with a third party device that acts as a Dot1x authenticator. (643209) (1)
Series Affected: CCS-750 Series
SKUs Affected: DCS-7050CX4M-48D8
- Setting different actions for packets marked as green and yellow by two rate three color (trTCM) QoS policy-map does not work (768572) (1)
- The inPktsBadTag MACsec counter will not increment. Instead the drops will be accounted under the inPktsNotValid MACsec counter. (926900)
SKUs Affected: DCS-7050CX4M-48D8
- MACsec encrypted flow control packets and priority-flow-control packets that are received will not be visible in "show flowcontrol" or "show priority-flow-control counters". (927568)
SKUs Affected: DCS-7050CX4M-48D8
- Mirror to GRE with forwarding drop is not supported on T4.
Workaround is to use Mirror on Drop feature for capturing drop packets.
(930242) (1)
- Persistent port security MACs are cleared when the StrataL2 agent is restarted. (1089853) (1)
Series Affected: 7358 and DCS-7060X6 Series
SKUs Affected: 7358X4-SC, DCS-7050CX4-24D8, DCS-7050CX4M-48D8, DCS-7050DX4-32S-F, DCS-7050DX4M-32S-F, DCS-7050PX4-32S-F and DCS-7050SDX4-48D8

7388 and DCS-7060X5 Series

- On MLAG setup, pruned BUM packets Counters are not visible from CLI (841924)
- Matching on VRF in direct-flow policy is not supported. (845677)

- (*,G) multicast routes are unsupported (869527)
- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- DLB on L3 LAG interface is not supported in hardware (920830) (1)
- Mirror to GRE with forwarding drop is not supported on Tomahawk4. Workaround is to use Mirror on Drop feature for capturing drop packets. (930240)
- Metadata fields of LNS, Congestion, QueueID, Queue Depth, Queue Transmit Byte count and MMU Stats are not supported (934700)
- GRE Tunnels are not programmed on exceeding hardware capacity of 256 Tunnels. This is indicated by "Hardware Programming Failure" under "show interface tunnel id". (968201)
- Metadata is stamped only on Layer 3 routed packets in IFA 2.0 transit mode (990808)
- Packets dropped by ACL rules are included in the count performed by the ECMP monitor snapshot feature. (991089) (1)
- Port-channel sub-interface will consume more hardware resources than sub-interface on routed physical interfaces (1001437) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- Config replace does not bring subinterface to operation state which was previously failed due to hardware resource full. (1009696) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F and DCS-7060DX5-64S-R
- VXLAN is not supported when system has sub-interfaces configured. (1015740) (1)
 Series Affected: 7388 Series
 SKUs Affected: 7388X5-SC, DCS-7060DX5-32-F, DCS-7060DX5-32-R, DCS-7060DX5-64-F, DCS-7060DX5-64E-F, DCS-7060DX5-64S-F, DCS-7060DX5-64S-R and DCS-7060PX5-64E-F
- The traffic loss will be observed on parent interface when last subinterface is flapped under the parent port. (1038667)
 Series Affected: DCS-7060X Series
- AGM snapshots only reflect ECMP member hashing before CLB redirection. (1193385) (1)

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

DCS-7060X6 Series

- A flow transmitted over LAG and/or ECMP may move to another LAG/ECMP member after UDF filtered mirroring or UDF ACL is applied. (912829) (1)
- DLB on L3 LAG interface is not supported in hardware (920830) (1)
- In the event of an unexpected Strata agent restart, the StrataCounters agent may also restart. (963956)
- Packet corruption can occur if ingress MTU size is multiple of 4.

As a workaround, use egress MTU size or do not use a multiple of 4 as the MTU size (990433)

- D-ECN profile exhaustion may occur when modifying the D-ECN configuration on a range of interfaces and tx-queues when there are one or fewer D-ECN profile resources available.

Remove the D-ECN configuration on the range of interfaces and tx-queues and then configure the desired D-ECN configuration. (1006441)

- DLB on an ECMP group with more than 64 members may initially have a bias towards a subset of ECMP members. (1030731)
Series Affected: DCS-7060X6 Series
SKUs Affected: DCS-7060X6-64PE and DCS-7060X6-64PE-F
- Enable standalone link training or auto-negotiation if supported for better performance with 800G copper cables. (1039049) (1)
SKUs Affected: DCS-7060X6-64PE and DCS-7060X6-64PE-F
- Lossless traffic is not subject to dynamic ECN thresholds. (1252726)
- If a packet latency exceeds 4.2 seconds, the internal latency timer will roll over to zero. This means latencies greater than 4.2 seconds may not have latency-based ECN marking applied as configured, and could egress marked or unmarked in this circumstance. (1271932)

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

Transceiver Series

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

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

- * "speed auto 100full" : enable autoneg, advertise only 100M
 - * "speed 100full" : disable autoneg, force speed to 100M (159401)
 - SKUs Affected: 1000BASE-T
- 25G transceivers could suffer signal degradation when using the MAM1Q00A-QSA QSFP-to-SFP adapter.

Use a MAM1Q00A-QSA28 QSFP28-to-SFP28 adapter instead. (253324)
- During reboot, a peer may see a link-flap on ports using 1000BASE-T transceivers.
The workaround is to configure the link-debounce period to 30 seconds for the link-up transition on the peer interfaces. (268080)
SKUs Affected: 1000BASE-T
- For SKUs CAB-O-2Q-400G, CAB-O-4Q-400G, CAB-D-2Q-400G and CAB-D-4Q-400G, when operating at 25G or 50G NRZ speeds, forward error-correction should be manually configured (enabled or disabled) on both sides of the link as the default error-correction on both ends may not match. (392124)
- SFP-10G-T transceivers can only advertise one speed at a time. Desired speed and negotiation must be explicitly configured. Recommended commands are:
 - * "speed auto" / "speed auto 10gfull"/"speed auto 10000full" : enable autoneg, advertise only 10G
 - * "speed auto 1gfull" / "speed auto 1000full" : enable autoneg, advertise only 1G
 - * "speed auto 100full" : enable autoneg, advertise only 100M
 - * "speed 100full" : disable autoneg, force speed to 100M (413941)
 SKUs Affected: SFP-10G-T
- Revision 20 OSFP-400G-2FR4 transceivers may not correctly report the operational value of the configured Tx Input Equalization. This can be observed in 'show interfaces transceiver tuning detail', in which the Operational Tx Input Equalization may not match the Configured value. If the 'Tx Input Equalization' Configured column shows auto, automatic Tx Input Equalization is enabled, despite the Operational value shown. (423045)
SKUs Affected: OSFP-400G-2FR4
- Using a default speed configuration on any interface associated with a 400G optical transceiver may prevent all interfaces associated with the transceiver from linking up.
The workaround is to explicitly configure speed on all interfaces associated with the 400G optical transceiver. (426232)
- FEC should be configured manually when using OSFP/QSFP-DD to QSFP200/SFP50 breakout cables at 25G or 50G-2 speeds (432776)
SKUs Affected: CAB-D-2Q-400G, CAB-D-4Q-400G, CAB-D-8S-400G, CAB-O-2Q-400G, CAB-O-4Q-400G and CAB-O-8S-400G

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

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

- Interfaces with BCM82391 PHY do not support 100GBASE-T transceivers. (535121)
- QSFP200 cables and optics encoded using the CMIS specification are not recognized in QSFP100 ports. (536318)
- Some Arista and 3rd party 400GBASE-ZR/ZR+ transceivers may display incorrect data displayed in the 'show transceiver performance monitoring' command.

Impacted data is read from page 35h (coherent page). Example: Rx total/channel power, Tx power, ESNR, OSNR etc.
pre-FEC BER and uncorrected frames are not impacted

Impacted transceivers: Arista 400GBASE-ZR/ZR+, manufactured by Acacia, 3rd party Acacia/Cisco 400GBASE ZR/ZR+

Workaround: PM data is updated every PM period (defaults to 15 min)
(645571)

SKUs Affected: OSFP-400G-ZR, OSFP-400G-ZRP, QDD-400G-ZR and QDD-400G-ZRP

- OSFP & QSFP-DD port LEDs are not supported in hardware LED CLI config mode (650615)
- 3rd party 400GBASE-ZR transceivers, manufactured by Inphi/Marvell, did not implement DOM Rx channel power

Workaround: use Rx total power, if the link is direct or through mux
(650772)

SKUs Affected: QDD-400G-ZR

- Interfaces with the SFP-10G-MRA-T 10GBASE-T module may link up during system reload if the peer is enabled. During the reload the port will not pass traffic. This will occur on enabled or disabled ports. (651386)
- SFP 10GBASE-T transceivers which do not support Soft RX_LOS or Soft TX_DISABLE are not supported on DSFP ports. (702025)

SKUs Affected: SFP-10G-T

- The rate adapting transceivers SFP-10G-MRA-T, SFP-10G-RA-1G-LX, and SFP-10G-RA-1G-SX will not allow PAUSE or PFC frames to pass to and from the peer device. (881411)

SKUs Affected: SFP-10G-MRA-T

- 400GBASE-ZR transceivers do not support Zero Touch Provisioning. (882209)
SKUs Affected: OSFP-400G-ZR, OSFP-400G-ZRP, QDD-400G-ZR and QDD-400G-ZRP
- The SFP-1G-T pluggable 1000BASE-T SFP is not capable of advertising PAUSE or ASM PAUSE during auto negotiation. This does not impact the ability to send or receive PAUSE per the specific switch's abilities and flow control configuration. (904116)
SKUs Affected: 1000BASE-T
- After using "reload fast-boot" to upgrade to a new EOS release, certain transceivers may have new default tuning values that are not applied until the next system reboot or after the transceiver is reinserted. (1032377)
- The SFP-10G-T may become error disabled with cause xcvr-unsupported if inserted into a slot with insufficient power (less than 2W).
A number of workarounds exist:
 - Use a lower power 10GBASE-T SFP such as the SFP-10G-T-RP (1.5W). For 1G and 100M speeds, the SFP-1G-T may be used.
 - Verify with Arista that it is safe to override the power limit with 'transceiver power ignore'. You may verify this by checking the Arista Transceiver Guide for limitations. In cases where the platform of interest is not listed, contact Arista TAC.
 - Disable the transceiver in EOS with 'transceiver diag simulate removed' command and save to the startup config. This will allow ASU reloads to proceed at the expense of the SFP-10G-T ports until it has been confirmed that it is safe to override the power limit. (1133243)
SKUs Affected: SFP-10G-T

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

- The DirectFlow "set vlan action" command is not supported. (237046)
- When configuring a DirectFlow action, specifying "interface hardware-loopback" as the destination of "action output" is not supported. (243031)
- IPv6 packets with WESP payload (next header 141) bypass the IPv6 security ACLs and service policies. (243387)
- A UDF rule matching on the pattern 0 might cause false positive matches if the offset of the UDF lies in a region of the packet where UDF extraction is not possible. (344985)
- Directflow configuration with "set vlan", but no "action set output interface" configured is not supported. (345852)
- MAC Egress ACL configurations matching IP packets may not work when the IPv4/v6 per-interface counter feature is also configured. (411243)

- CLI output of "show platform trident 13 shadow mpls-vc-and-swap-label" will show push as MplsLabelAction even for swap operation. (416134)
- Packets dropped by the MACsec engine may not be reported in the CLI (420085)
SKUs Affected: DCS-7050CX3M-32S
- MPLS packet undergoing swap operation does not honor the MTU configured on ingress/egress interface. (422668)
- Only 2K MPLS routes are supported (425691)
Series Affected: CCS-720XP Series
- If the packet ingresses on one line card and egresses on a different line card, then a truncated egress mirrored packets will be 4 bytes shorter than expected. (445200)
SKUs Affected: DCS-7304 and DCS-7308
- When IPV6 VXLAN underlay core interface is an access port on a SVI, VXLAN encapsulated packets may be sent out with VLAN tag in the outer ethernet header. The work around is to configure the port as a trunk port. (533929)
- Packets mirrored to GRE with a size within 16 bytes of the MTU limit may be dropped on egress due to exceeding the MTU size after the mirrored packet is encapsulated. (592957)
- "Dont fragment" bit is not set in the outer IP header of the VXLAN encapsulated packets when the payload is non-IP. (605003)
- When Postcard Telemetry is enabled, EVPN Multicast will not be supported. (614072)
Series Affected: 7300X3, CCS-720XP, DCS-7050SX3, DCS-7050TX3 and DCS-7050X3 Series
- ACL resource sharing feature is not supported for NAT translations of tunneled packets like VXLAN-tunneled packets.

Workaround is to turn off ACL resource sharing: "no hardware nat static access-list resource sharing" (620222)
- If the destination MAC address of a NAT flow is not learnt on an interface, then the packet will be dropped. Flooding the translated packet on all interfaces belonging to the VLAN is not supported.

Workaround is to adjust the ARP and MAC address timeouts, so that MAC addresses is learnt again. (642094)
- DirectFlow actions on interfaces configured as dot1q-tunnel are not supported. (686339)

- If an egress ACL is filtering on the NAT flow's source IP, then the ACL rules have to permit both original source IP and translated IP to avoid NAT flows getting dropped. (688883)
- With EVPN Multicast with underlay as HER configured, known multicast traffic might not be load balanced across all ECMP links. (709053)
- Ingress mirroring double tagged packets to CPU will have the outer VLAN tag stripped.
Workaround is to use a port mirroring destination instead. (722891)
- Packets with Broadcast source MAC may be copied to CPU instead of being dropped. (745001)
- For features which use egress TCAM the match is on outer header for a tunnelled encapsulated packet. There is no workaround for this bug because this is a hardware limitation. (801113)
- When ingress sFlow is enabled along with VXLAN data plane learning, VXLAN remote VTEP Learning could fail resulting in remote MACs not being learnt for the sampled packets. (825322)
- Layer 2 Multicast forwarding will experience higher traffic outage during SSU. Outage will continue till the router re-learns the multicast receivers post SSU operation. There is no workaround for this problem (833449)
- With EVPN multicast configured with IPv4 underlay tunnels, overlay IPv6 unknown multicast traffic may not get flooded in broadcast domain.
Work-around is to use VLANs without IGMP-proxy for unknown IPv6 multicast traffic. (869000)
- Packets that match on a dynamic NAT ACL will be sent to CPU if the packets are egressing out of any other interface that also has NAT rules configured. The workaround is to make the NAT ACL less greedy by either reducing the scope of matching IPs or by adding deny rules for specific destination IPs that should not be NATed. (900348)
- A warning will not be emitted when 'dot1x mac based access-list' is missing, but 'dot1x aaa unresponsive action traffic allow access-list' or 'dot1x authentication failure action traffic allow access-list' is configured. (918137)
- When one or more L2 MAC addresses cannot be programmed in the hardware due to events such as hash collisions, user will not be able to see them using 'show platform trident mac-address-table missing' command. There is no workaround for this CLI output issue. (968260)
- Same dynamic twice NAT rule under multiple interfaces (that may be part of an ECMP group) is not supported. (1086985)

- On a root bridge, if Port Isolation is configured on all ports connected to a forwarding STP state port, it might lead to complete traffic loss. Similarly on a non root bridge, if Port Isolation is configured on a forwarding STP state port, it might lead to traffic loss in both outgoing and incoming directions on that port. (1108888)
- If the "switchport vlan translation out required" and "macsec profile" commands are configured on the same interface, all packets will egressing the interface will be dropped (1207817)
 Series Affected: CCS-722XPM Series
 SKUs Affected: CCS-720XPM-48TH-6SY

7358 and DCS-7050X4 Series

- Setting different actions for packets marked as green and yellow by two rate three color (trTCM) QoS policy-map does not work (768572) (1)
- If 300K routes with 128-way ECMP are programmed, system can encounter out of memory conditions. Agents may restart as the scale exceeds platform capabilities. There is no workaround for this issue. (785832)
- Greenspan packets where the size of the mirror payload plus the Greenspan header exceeds the configured MTU size of the ingress interface will be fragmented on a second pass of the pipeline.

Note that high rates of fragmentation can result in high CPU utilization. (791223)

- The interface with selective dot1q tunnel feature configured will fail to remove the service tag when transmitting a packet if that packet was received on another interface where selective dot1q tunnel feature was also configured on the same VLAN(s). (808169)
- Egress filtered mirroring is not supported due to hardware limitation. (809908)
- By default IPv6 packets are set to not fragment. If an added IPv6 Greenspan header causes a packet to exceed the MTU size of the destination next hop, then it will be dropped. Workaround is to increase MTU size of the nexthop port. (818665)
- The number of monitoring sessions is limited to 4 even though the hardware may be capable of more. (824563)
- Unfiltered ingress mirroring of individual port-channel members are not supported.

Workaround: Ingress mirror the entire port channel, egress mirror the

destination port, or an an access-group (ACL) to the port channel member source. (837229)

- Filtered mirrored Greenspan packets cannot be egress sampled by sFlow. (863261)
- Configuring VXLAN bridging in unsupported profile with "forwarding vxlan disabled" may result in undesired behavior. Workaround is to remove VXLAN bridging configuration or to configure VXLAN supported profile by removing "forwarding vxlan disabled". (863281)
- For ports configured at 25G speeds using fire-code encoding FEC, PTP ingress timestamps are not correctly recorded.

Workaround: Configure the port to use either reed-solomon encoding FEC, or disable error-correction using 'no error-correction encoding' (891861)

- Trident4 does not have configurable debug counters, so IP Counters are only possible with Flex Counters. (892598)
- Per VRF PIM Bidir group configuration is not supported. Configure the bidir groups in the default VRF. (898306)
- Early shipments of system bundles with this switch card were loaded with the incorrect EOS-4.29.0.2F binary (32-bit) instead of the EOS64-4.29.0.2F binary (64-bit). Unexpected traffic drops may be seen if the 32-bit binary is used. Upgrade to a 64-bit version of EOS to avoid such traffic drops. (909191)

SKUs Affected: 7358X4-SC

- Mirror to GRE with forwarding drop is not supported on T4. Workaround is to use Mirror on Drop feature for capturing drop packets. (930242) (1)
- LLC Validation and L2 Protocol Transparency on an interface are not supported together. (961473)
- Hardware HER will generate a hit to SPLIT_HORIZON_DROP counter for each replicated packet (1073018)

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

Conditions and Impacts

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

Condition: Rib agent restart

Impact: When the Rib agent goes down, all routing sessions are terminated; all BGP sessions drop, all OSPF adjacencies are lost, and neighbors stop forwarding traffic to us. When the Rib agent restarts, adjacencies are re-formed, and, after protocol convergence, traffic flows through us again. In most cases where there is adequate network-level redundancy, application-visible impact should be minimal

Condition: Rib agent hang

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

Condition: kernel crash

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

Condition: forwarding agent restart

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