

K-Series[®]

Customer Release Note

Firmware Version 8.63.09.0001

Table of Contents

Customer Release Notes	2
Installing an I/O or I/O Fabric Module	4
Chassis Minimum FW Version Required:.....	4
System Behavior	5
Multi-slot Chassis User Capacities.....	6
Maximum User Capacity:	6
K-E OS-PPC - Port Capacities License	6
Port Mirroring	6
Class of Service:.....	7
Class of Service Support.....	7
Link Aggregation (LAG)	7
Multi-User 802.1X.....	7
Power over Ethernet Controller Code Upgrade.....	7
Power Supply Behavior	8
Features, Scale and Capacity.....	8
Policy Capacities	8
Multicast Capacities	9
DHCP Capacities	9
Licensed Features.....	9
100Mb Optics:.....	10
Capacity Reductions in 8.63.05.0004	15
Problems Corrected in 8.63.09.0001.....	15
Problems Corrected in 8.63.08.0002.....	15
Problems Corrected in 8.63.07.0003.....	16
Problems Corrected in 8.63.06.0005.....	19
Problems Corrected in 8.63.05.0004	19
Problems Corrected in 8.63.04.0003	21
Problems Corrected in 8.63.03.0002.....	21
Problems Corrected in 8.63.02.0004.....	23
Problems Corrected in 8.63.01.0019	24
Problems Corrected in 8.62.04.0001.....	26
Problems Corrected in 8.62.03.0006.....	26
Problems Corrected in 8.62.02.0022	28
Features Enhancements 8.62.01.0034.....	30
Problems Corrected in 8.62.01.0034	30
Problems Corrected in 8.61.02.0001.....	33
Features Enhancements 8.61.01.0018	33
Problems Corrected in 8.61.01.0018	33

Problems Corrected in 8.42.06.0001.....	34
Problems Corrected in 8.42.05.0003.....	34
Problems Corrected in 8.42.04.0016.....	35
Problems Corrected in 8.42.03.0006.....	36
Features Enhancements 8.42.02.0012.....	38
Problems Corrected in 8.42.02.0012.....	38
Problems Corrected in 8.42.01.0007.....	38
Features Enhancements 8.41.01.0004.....	39
Problems Corrected in 8.41.01.0004.....	41
Features Enhancements 8.32.02.0008.....	45
Problems Corrected in 8.32.02.0008.....	46
Features Enhancements 8.32.01.0021.....	46
Problems Corrected in 8.32.01.0021.....	47
Problems Corrected in 8.31.03.0001.....	52
Problems Corrected in 8.31.02.0014.....	52
Features Enhancements 8.31.01.0006.....	57
Problems Corrected in 8.31.01.0006.....	58
Problems Corrected in 8.21.03.0001.....	71
Feature Enhancements in 8.21.02.0001.....	71
Problems Corrected in 8.21.02.0001.....	72
Feature Enhancements in 8.11.05.0006.....	82
Problems Corrected in 8.11.03.0005.....	87
Feature Enhancements in 8.02.01.0012.....	88
Feature Enhancements in 8.02.01.0012.....	88
Problems Corrected in 8.02.01.0012.....	91



6480 Via Del Oro
San Jose, CA 95119
+1 408-579-2800

Customer Release Notes

K-Series

Firmware Version 8.63.09.0001

September 2021

INTRODUCTION:

This document provides specific information for version 08.63.09.0001 of firmware for the K-Series. The K-Series modules may be installed in the K6 and K10 chassis. This version of firmware supports the following K-Series model numbers:

K-Series Fabric Modules	
KK2008-0204-F2	K10 Management/Fabric Module (4) 10GB via SFP+
KK2008-0204-F1	K6 Management/Fabric Module (4) 10GB via SFP+

K-Series I/O Modules	
KT2006-0224	K-Series (24) Port 10/100/1000 802.3at RJ45 PoE IOM
KG2001-0224	K-Series (24) Port 1Gb SFP IOM
KK2008-0204	K-Series (4) Port 10Gb SFP+ IOM
KT2010-0224	K-Series (24) Port 10/100/1000 802.3at Mini-RJ21 PoE IOM

Extreme Networks recommends that you thoroughly review this document prior to installing or upgrading this product.

For the latest firmware versions, visit <http://support.extremenetworks.com>

PRODUCT FIRMWARE SUPPORT:

Status	Firmware Version	Product Type	Release Date
Current Version	8.63.09.0001	Customer Release	September 2021
Previous Version	8.63.08.0002	Customer Release	March 2020
Previous Version	8.63.07.0003	Customer Release	December 2019
Previous Version	8.63.06.0005	Customer Release	July 2019
Previous Version	8.63.05.0004	Customer Release	April 2019
Previous Version	8.63.04.0003	Customer Release	December 2018
Previous Version	8.63.03.0002	Customer Release	September 2018
Previous Version	8.63.02.0004	Customer Release	June 2018
Previous Version	8.63.01.0019	Customer Release	April 2018
Previous Version	8.62.04.0001	Customer Release	June 2017
Previous Version	8.62.03.0006	Customer Release	May 2017
Previous Version	8.62.02.0022	Customer Release	March 2017
Previous Version	8.62.01.0034	Customer Release	October 2016
Previous Version	8.61.02.0001	Customer Release	August 2016
Previous Version	8.61.01.0018	Customer Release	May 2016

Previous Version	8.42.03.0006	Customer Release	April 2016
Previous Version	8.42.02.0012	Customer Release	January 2016
Previous Version	8.42.01.0007	Customer Release	October 2015
Previous Version	8.41.01.0004	Customer Release	September 2015
Previous Version	8.32.02.0008	Customer Release	May 2015
Previous Version	8.32.01.0021	Customer Release	March 2015
Previous Version	8.31.03.0001	Customer Release	January 2015
Previous Version	8.31.02.0014	Customer Release	November 2014
Previous Version	8.31.01.0006	Customer Release	September 2014
Previous Version	8.22.03.0006	Customer Release	July 2014
Previous Version	8.22.02.0012	Customer Release	June 2014
Previous Version	8.22.01.0022	Customer Release	April 2014
Previous Version	8.21.03.0001	Customer Release	February 2014
Previous Version	8.21.02.0001	Customer Release	December 2013
Previous Version	8.11.05.0006	Customer Release	December 2013
Previous Version	8.11.04.0005	Customer Release	October 2013
Previous Version	8.11.03.0005	Customer Release	September 2013
Previous Version	8.02.01.0012	Customer Release	April 2013

HIGH AVAILABILITY UPGRADE (HAU) FW COMPATIBILITY:

This version is HAU-compatible with any future release whose HAU-compatibility key is:

```
274dcb324cb75138fe7bbc5b7656d0fce60c1848
```

(The HAU key is reported using the CLI command 'dir images').

In an effort to reduce out of service time as much as possible for customers, HAU key changes are kept at a minimum. When HAU keys must change within a period of 18 months, maintenance releases will be available for the previous release. A maintenance release for the 8.4X train will be posted shortly after 8.61.01 is posted.

HARDWARE COMPATIBILITY:

This version of firmware is supported on all hardware revisions.

BOOT PROM COMPATIBILITY:

This version of firmware is compatible with all boot prom versions.

INSTALLATION INFORMATION:**Installing an I/O or I/O Fabric Module****Chassis Minimum FW Version Required:**

Multislot Chassis	
K6-Chassis	07.31.02.0008
K10-Chassis	
K-AC-PS-1400W	07.31.02.0008
K-AC-PS	07.31.02.0008, (07.42.02.0002 displays model number properly)
K-Series Fabrics	
KK2008-0204-F2	07.31.02.0008
KK2008-0204-F1	
K-Series I/O Modules	
KT2006-0224	07.31.02.0008
KG2001-0224	
KK2008-0204	
KT2010-0224	07.42.02.0002

System Behavior

The K-Series fabric module is responsible for management, configuration, file storage, fabric service, and flow setup for all ports in the chassis. When a chassis boots, the fabric is first initialized, and then the line cards are powered up and initialized. The system may become manageable prior to the line cards completing initialization. Use the command “show version” to determine which line cards have been initialized.

The operating configuration is stored in binary format on the fabric module. When replacing a fabric module you must transfer an ASCII version of the configuration to the new fabric, and then execute a “configure” command to restore the configuration. Each chassis has a USB connector which supports a mass storage flash stick that simplifies this step. For a K10 chassis:

1. Using original system with USB device installed, issue “show config outfile slot11/usb1/myconfig.txt”.
2. Install new fabric module—this clears the configuration.
3. Install flash stick.
4. Issue “config append slot11/usb1/myconfig.txt”.

K10 chassis requires a K10 fabric module; K6 chassis accepts either K6 or K10 fabric modules. K10 chassis’ line card slots are numbered 1–10 and the fabric numbered 11 (1–6, and 7 respectively for K6). Restart a line card by issuing the “reset <slot>” command. “reset system” and “reset 11” (7 for K6) are equivalent.

A K chassis supports a maximum of 12 × 10Gb ports. In addition, slot 10 in a K10 chassis only supports 10Gb line cards. Unsupported configurations cause an error message to appear to the console and the problem line card does not powered up.

Line cards may be inserted at any time and are normally initialized by the fabric module. A fabric module clears all configuration settings stored on it each time the module is moved to a different chassis. A status LED on the front of each line card indicates its status. If the LED is green, the line card is operational. The system should be notified prior to removing a line card. To do this, press the line card’s “off line” button with a small diameter object such as a paper clip. The Status LED changes from green to amber, and then turns off to indicate that the line card can now be removed.

Removing a line card while its status LED is lit may result in numerous hardware errors appearing and may temporarily interfere with the normal operation of control plane protocols, and in some cases other line cards or the fabric module. Wait for the line card status LED to turn off before removing the card.

Line card slots may also remain permanently in the shutdown state by issuing the command “set system module <slot> disable”. The disable is persistent through chassis reboots and with line card removal or reinstallation. System administrators may use this feature to take a suspect line card offline or to prevent empty slots in a chassis from becoming operational if accidentally populated.

Line cards in the disabled or powered down state display “Hardware is not physically present” for serial and version fields of “show version”. “show system module” displays disable/enable status for each slot. Disabling a slot does not remove its configuration from the system.

When a line card is removed or halted, the ports for that line card slot and any settings are retained. Interface Operstatus for its ports becomes “not-present”. The ports remain visible in MIBs and CLI, and configuration setting changes can be applied to these ports.

Inserting a new type line card into a previously configured slot clears configuration for ports on that slot.

The “show linecard” command displays information about the line card type that is configured for the slot. “clear linecard <slot>” both clears configuration and removes all ports for the specified line card slot. If the slot is populated, the line card is shut down and rebooted into the system. After the reboot completes, the ports are once again present in the system and have default configuration. The fabric module slot cannot be specified with this command.

Summary of conditions that cause current configuration to be lost:

- All configuration is lost when:
 - Fabric module changes chassis
 - “clear config” or “configure” commands or equivalent MIBs are issued
 - Switch 7 is toggled
- Line card specific configuration is lost when:
 - Type of installed line card changes
 - Line card type is changed with “linecard” CLI commands

Multi-slot Chassis User Capacities

Each K-Series multi-slot chassis (K6/K10) has a maximum authenticated user capacity.

Maximum User Capacity:

Chassis Type	Maximum User Capacity
K6-Chassis	1152
K10-Chassis	1920

All ports in the K-Series have an 8 user/port capacity limit by default. A ‘K-EOS-PPC’ license can be applied to the chassis to remove the per port limits, allowing for up to 256 users/port.

When present, the PPC license defaults the user-capacity at 256 users per port. You can modify this value using the CLI command ‘set multiauth port numusers’.

K-EOS-PPC - Port Capacities License

A license is required for each K chassis requiring additional port user capacity.

The license removes the per port restriction of 8 users per port for all ports in the chassis allowing up to 256 users/port. The total authenticated users in the chassis may not exceed the chassis user-capacities maximum described above.

Port Mirroring

The K-Series device provides support for 4 mirror instances.

A mirror could be a:

- “One-to-one” port mirror
- “One-to-many” port mirror
- “Many-to-one” port mirror
- IDS mirror*
- Policy mirror†
- Mirror N Packet mirror
- Remote Port Mirror

For the “one-to-many,” there is no limit to the number of destination ports.

For the “many-to-one,” there is no limit to the number of source ports.

For the port mirror case, the source ports(s) can be a physical port or VLAN.

Frames with layer 1 errors, such as bad CRC, are never mirrored.

* Support for no more than 1 IDS mirror. An IDS mirror specifies a LAG port as its destination. This LAG port may contain up to 10 physical ports.

† Destination ports of a policy mirror can be single or multiple (no limit) ports.

Remote Port Mirrors are supported and provide the ability to send port mirror traffic to a remote destination across the IP network. Traffic is encapsulated in an L2 GRE tunnel, and can be routed across the network.

Note that the preceding examples illustrate the number and types of mirrors supported, and how they can be used concurrently. The mirror configurations are not limited to these examples.

Class of Service:

Class of Service (CoS) is supported with, and without, policy enabled. Policy provides access to classes 8–255. Without policy, classes 0–7 are available.

Class of Service Support

- Supports up to 256 Classes of Service
- ToS rewrite
- 802.1D/P Priority
- K-Series supports 12 Transmit Queues per port (1 reserved for control-plane traffic)
 - Queues support Strict, WFQ and Hybrid Arbitration
 - All queues support rate-shaping
- 32 Inbound-Rate-Limiters per port
- 16 Outbound-Rate-Limiters per port
- Support for Flood-Limiting controls for Broadcast, Multicast, and Unknown Unicast per port.
- Management
 - Support for Enterasys CoS MIB

Link Aggregation (LAG)

The K-Series multi-slot chassis, K6 and K10, supports a total of 62 LAGs per chassis with up to 8 ports per LAG.

Multi-User 802.1X

Authentication of multiple 802.1X clients on a single port is supported. This feature only operates correctly when the intermediate switch forwards EAP frames, regardless of destination MAC address (addressed to either unicast or reserve multicast MAC).

To be standards compliant, a switch is required to filter frames with the reserved multicast DA. To be fully multi-user 802.1X compatible, the intermediary switch must either violate the standard by default, or offer a configuration option to enable the non-standard behavior. Some switches may require the Spanning Tree Protocol to be disabled to activate pass-through.

Use of a non-compatible intermediary switch results in the 802.1X authenticator missing multicast destined users' logoff and logon messages. Systems used by multiple consecutive users remain authenticated as the original user until the re-authentication period expires.

The multi-user 802.1X authenticator must respond to EAP frames with directed (unicast) responses. It must also challenge new user MAC addresses discovered by the multi-user authentication/policy implementation.

Compatible supplicants include Microsoft Window XP/2000/Vista, Symantec Sygate Security Agent, and Check Point Integrity Client. Other supplicants may be compatible.

The enterasys-8021x-extensions-mib and associated CLI is required to display and manage multiple users (stations) on a single port.

Power over Ethernet Controller Code Upgrade

Each release of K-Series firmware contains a copy of PoE microcontroller code. This code is installed in the microcontroller's flash memory system any time the K-Series boots and discovers the installed code is not the appropriate version. When up- or down-grading K-Series firmware, you may experience an additional delay in PoE delivery of a few minutes while this upgrade step completes.

Power Supply Behavior

A single power supply in the K-Series provides both system (12V) and PoE (54V) power. Up to four K-AC-PS-1400W power supplies can be installed in either the K6 or K10.

Using a 120V AC source the supply supports 1000W of power output, (600W system and 400W PoE).

Using a 220V AC source the supply supports 1400W of power output, (600W system and 800W PoE).

Available Power Table

Installed Supplies	System Power (120–240V)	PoE Power (120V)	PoE Power (220V)
1	600W	400W	800W
2	1200W	800W	1600W
3	1800W	1200W	2400W
4	2400W	1600W	3200W

Features, Scale and Capacity

Each release of K-Series firmware contains specific features and associated capacities or limits. The CLI command "show limits" provides a detailed description of the features and capacity limits available on your specific hardware with its current licensing. Use this command to get a complete list of capacities for this release.

Policy Capacities

Roles/Profile	127
Rules	8,191

IGMP Capacities

IGMP Groups	5K
DVMRP Routes	1K

Router Capacities

The following table defines the router capacities:

ARP/ND(Neighbor Discovery) Entries	16,000 system, 4,000 (per interface)
Static ARP/ND(Neighbor Discovery) Entries	512
IPv4 : Route Table Entries	30,000
IPv6 : Route Table Entries	25,000
IPv4: Router interfaces	256
IPv6: Router interfaces	256
OSPF Areas	8
OSPF LSA(s)	30,000
OSPFv3 LSA(s)	16,000
OSPF Neighbors	60
Static Routes	256
RIP Routes	2,500
Configured RIP Nets	300
VRRP Interfaces	32
Routed Interfaces	128
ACLs	1,000

Access Rules	5,000
Access Rules – Per ACL	1,000
Policy Based Routing Entries	100
ECMP Paths	8
Static VRF	64
Dynamic VRFs	8
Router Links in Area	100
Secondaries per Interface	128
Secondary Interfaces per Router	2,048
IP Helper addresses (per router/per interface)	5,120/20
Tracked object sessions ICMP/TCP/UDP	2000
BFD sessions	64

Multicast Capacities

IGMP/MLD Static Entries	64
IGMP/MLD *,G and S,G Groups ¹	64K
IGMP/MLD Snooping Flow Capacity	5K
Multicast Routing (PIM/DVMRP flows)	4K
IGMP/MLD Clients ²	64K

1. Group entries may be consumed for each egress VLAN of a routed flow.

2. A client is defined as a reporter subscribing to a *,G or S,G group, or sourcing a multicast flow.

DHCP Capacities

DHCP Server Leases	5,000
DHCP Pools	100

Shortest Path Bridging

SPBv (constrained by 4094 VLANs)	Up to 100 VLANs mapped as base VLANs	Up to 50 SPBv nodes in SPB region
----------------------------------	--------------------------------------	-----------------------------------

Tunneling Capacities

Total Number of Tunnels	16 *Licensed
-------------------------	-----------------

Some of the limits listed in the tables above may **not** be enforced by the firmware and may cause unknown results if exceeded.

Licensed Features

The Advanced Routing license, model number K-EOS-L3 is applied per chassis and provides support for the following features:

- OSPF/OSPFv3, PIM-SM/PIM-SSM, PIM-DM and VRF
- IP-Tunnels

The Chassis Bonding license, K-EOS-VSB, allows two like slot count chassis to be bonded together. Two K6 chassis or two K10 chassis may be chassis bonded using a unique license per chassis. (Two licenses are required for the chassis bonded pair.)

NETWORK MANAGEMENT SOFTWARE:

NMS	Version No.
NetSight Suite	6.1 or greater

NOTE:

If you install this image, you may not have control of all the latest features of this product until the next version(s) of network management software. Review the software release notes for your specific network.

PLUGGABLE PORTS SUPPORTED:**100Mb Optics:**

SFP Optics	Description
MGBIC-N-LC04	100 Mb, 100Base-FX, IEEE 802.3 MM, 1310 nm Long Wave Length, 2 Km, LC SFP
MGBIC-LC04	100 Mb, 100Base-FX, IEEE 802.3 MM, 1310 nm Long Wave Length, 2 Km, LC SFP
MGBIC-LC05	100 Mb, 100Base-LX10, IEEE 802.3 SM, 1310 nm Long Wave Length, 10 Km, LC SFP
MGBIC-100BT	100 Mb, 100BASE-T Copper twisted pair, 100 m, RJ45 SFP

1Gb Optics:

MGBICs	Description
MGBIC-LC01	1 Gb, 1000Base-SX, IEEE 802.3 MM, 850 nm Short Wave Length, 220/550 M, LC SFP
MGBIC-LC03	1 Gb, 1000Base-SX-LX/LH, MM, 1310 nm Long Wave Length, 2 Km, LC SFP
MGBIC-LC07	1 Gb, 1000Base-EZX, IEEE 802.3 SM, 1550 nm Long Wave Length, 110 Km, LC SFP (Extended Long Reach)
MGBIC-LC09	1 Gb, 1000Base-LX, IEEE 802.3 SM, 1310 nm Long Wave Length, 10 Km, LC SFP
MGBIC-MT01	1 Gb, 1000Base-SX, IEEE 802.3 MM, 850 nm Short Wave Length, 220/550 M, MTRJ SFP
MGBIC-02 ³	1 Gb, 1000Base-T, IEEE 802.3 Cat5, Copper Twisted Pair, 100 m, RJ 45 SFP
MGBIC-08	1 Gb, 1000Base-LX/LH, IEEE 802.3 SM, 1550 nm Long Wave Length, 80 km, LC SFP
MGBIC-BX10-U	1 Gb, 1000Base-BX10-U Single Fiber SM, Bidirectional 1310nm Tx / 1490nm Rx, 10 km, Simplex LC SFP (must be paired with MGBIC-BX10-D)
MGBIC-BX10-D	1 Gb, 1000Base-BX10-D Single Fiber SM, Bidirectional, 1490nm Tx / 1310nm Rx, 10 km, Simplex LC SFP (must be paired with MGBIC-BX10-U)
MGBIC-BX40-U	1 Gb, 1000Base-BX40-U Single Fiber SM, Bidirectional, 1310nm Tx / 1490nm Rx, 40 km, Simplex LC SFP (must be paired with MGBIC-BX40-D)
MGBIC-BX40-D	1 Gb, 1000Base-BX40-D Single Fiber SM, Bidirectional, 1490nm Tx / 1310nm Rx, 40 km, Simplex LC SFP (must be paired with MGBIC-BX40-U)
MGBIC-BX120-U	1 Gb, 1000Base-BX120-U Single Fiber SM, Bidirectional, 1490nm Tx / 1590nm Rx, 120 km, Simplex LC SFP (must be paired with MGBIC-BX120-D)
MGBIC-BX120-D	1 Gb, 1000Base-BX120-D Single Fiber SM, Bidirectional, 1590nm Tx / 1490nm Rx, 120 km, Simplex LC SFP (must be paired with MGBIC-BX120-U)

10Gb Optics:

SFP+ Optics	Description
10GB-SR-SFPP	10 Gb, 10GBASE-SR, IEEE 802.3 MM, 850 nm Short Wave Length, 33/82 m , LC SFP+

³ 100Mb speed is also supported for MGBIC-02 on S-Series & K-Series.

10GB-LR-SFPP	10 Gb, 10GBASE-LR, IEEE 802.3 SM, 1310 nm Long Wave Length, 10 km , LC SFP+
10GB-ER-SFPP	10 Gb, 10GBASE-ER, IEEE 802.3 SM, 1550 nm Long Wave Length, 40 km , LC SFP+
10GB-LRM-SFPP	10 Gb, 10GBASE-LRM, IEEE 802.3 MM, 1310 nm Short Wave Length, 220 m , LC SFP+
10GB-ZR-SFPP	10 Gb, 10GBASE-ZR, SM, 1550 nm, 80 km , LC SFP+
10GB-USR-SFPP	10Gb, 10GBASE-USR MM 850nm, LC SFP+
10GB-SRSX-SFPP	10Gb / 1Gb DUAL RATE, MM 850nm 10GBASE-SR / 1000BASE-SX, LC SFP+
10GB-LRLX-SFPP	10Gb / 1Gb DUAL RATE, SM 1310nm 10GBASE-LR / 1000BASE-LX, 10 km LC SFP+
10GB-BX10-D	10Gb, Single Fiber SM, Bidirectional, 1330nm Tx / 1270nm Rx, 10 km SFP+
10GB-BX10-U	10Gb, Single Fiber SM, Bidirectional, 1270nm Tx / 1330nm Rx, 10 km SFP+
10GB-BX40-D	10Gb, Single Fiber SM, Bidirectional, 1330nm Tx / 1270nm Rx, 40 km SFP+
10GB-BX40-U	10Gb, Single Fiber SM, Bidirectional, 1270nm Tx / 1330nm Rx, 40 km SFP+
SFP+ Copper	Description
10GB-C01-SFPP	10Gb pluggable copper cable assembly with integrated SFP+ transceivers, 1 m
10GB-C03-SFPP	10Gb pluggable copper cable assembly with integrated SFP+ transceivers, 3 m
10GB-C10-SFPP	10Gb pluggable copper cable assembly with integrated SFP+ transceivers, 10 m
SFP+ Laserwire	Description
10GB-LW-SFPP	SFP+ Laserwire Transceiver Adapter
10GB-LW-03	Laserwire Cable 3 m
10GB-LW-05	Laserwire Cable 5 m
10GB-LW-10	Laserwire Cable 10 m
10GB-LW-20	Laserwire Cable 20 m
10GB-F10-SFPP	10Gb, Active optical direct attach cable with 2 integrated SFP+ transceivers, 10m
10GB-F20-SFPP	10Gb, Active optical direct attach cable with 2 integrated SFP+ transceivers, 20m
SFP+ DWDM Optics	Description
10GB-ER21-SFPP	10GB-ER, DWDM CH21 SFP+
10GB-ER23-SFPP	10GB-ER, DWDM CH23 SFP+
10GB-ER24-SFPP	10GB-ER, DWDM CH24 SFP+
10GB-ER29-SFPP	10GB-ER, DWDM CH29 SFP+
10GB-ER31-SFPP	10GB-ER, DWDM CH31 SFP+
10GB-ER33-SFPP	10GB-ER, DWDM CH33 SFP+
10325	10 Gb, 10GBASE-ZR 102 channel DWDM tunable SFP+ Transceiver
SFP+ CWDM Optics	Description
10GB-LR271-SFPP	10Gb, CWDM SM, 1271 nm, 10 km, LC SFP+
10GB-LR291-SFPP	10Gb, CWDM SM, 1291 nm, 10 km, LC SFP+
10GB-LR311-SFPP	10Gb, CWDM SM, 1311 nm, 10 km, LC SFP+
10GB-LR331-SFPP	10Gb, CWDM SM, 1331 nm, 10 km, LC SFP+

Dual speed operation: The SFP+ ports support the use of SFP+ transceivers and SFP transceivers. (10Gb/1Gb)
The SFP ports support the use of SFP transceivers and 100Mb transceivers. (1Gb/100Mb)

NOTE:

Installing third-party or unknown pluggable ports may cause the device to malfunction and display MGBIC description, type, speed, and duplex setting errors.

SUPPORTED FUNCTIONALITY:

Features		
Multiple Authentication Types Per Port - 802.1X, PWA+, MAC	Layer 2 through 4 VLAN Classification	Entity MIB
Multiple Authenticated Users Per Port - 802.1X, PWA+, MAC	Layer 2 through 4 Priority Classification	IP Routing
DVMRP	Dynamic VLAN/Port Egress Configuration	Static Routes
SNTP	Ingress VLAN Tag Re-write	RIP v2
Web-based configuration (WebView)	VLAN-to-Policy Mapping	OSPF/OSPFv3
Multiple local user account management	RMON – Statistic, History, Alarms, Host, HostTopN,	OSPF ECMP
Denial of Service (DoS) Detection	RMON Matrix groups, Host, HostTopN, Events, Capture and Filter	OSPF Alternate ABR
Passive OSPF support	SMON – VLAN and Priority Statistics	Graceful OSPF Restart (RFC 3623)
802.1X – Authentication	SNMP v1/v2c/v3	RIP ECMP, CIDR configuration
802.1D – 1998	Port Mirroring	Virtual Router Redundancy Protocol (VRRP)
802.1Q – Virtual Bridged Local Area Networking	Flow Setup Throttling	ICMP
GARP VLAN Registration Protocol (GVRP)	MAC locking (Static/Dynamic)	Protocol Independent Multicast - Sparse Mode (PIM-SM)
802.1p – Traffic Class Expediting	Node/Alias table	Proxy ARP
802.1w – Rapid Reconfiguration of Spanning Tree	Policy-Based Routing	Basic Access Control Lists
802.1s – Multiple Spanning Trees	SSH v2	Extended ACLs
802.1t – Path Cost Amendment to 802.1D	OSPF NSSA, equal cost multi-path	Auto MDI-X Media Dependent Interface Crossover Detect (Enhanced for non auto negotiating ports)
802.3 – 2002	Audit trail logging	DHCP Server
802.3ad – Link Aggregation	RADIUS Client	DHCP Relay w/ option 82
802.3x – Flow Control	FTP/TFTP Client	Jumbo Frame support
Static Multicast Configuration	Telnet – Inbound/Outbound	Directed Broadcast

Broadcast Suppression	Configuration File Upload/Download	Cisco CDP v1/2
Inbound and Outbound Rate Limiting	Text-based Configuration Files	CLI Management
Strict and Weighted Round Robin Queuing	Syslog	CPU and task Debugging
IGMP v1/v2 and Querier support	Span Guard	RADIUS (Accounting, Snooping)
SMON Port and VLAN Redirect	RAD (Remote Address Discovery)	Split RADIUS management and authentication
Spanning Tree Loop Protection	Cabletron Discovery Protocol (CDP)	Link Flap detection

Features		
TACACS+	NetFlow v5/v9	Daylight Savings Time
Type of Service (ToS) Re-write	LLDP and LLDP-MED	RFC 3580 with Policy support
Multi-VRF (IPv4)	VRF-Aware Policy Based Routing	Flex-Edge
VRF IPv4/IPv6 Static Route Leaking	IPv6 Static Routing	VRF-Aware DHCP Relay
IPv6 Policy Based Routing	IPv6 DHCP Relay	IPv6 ACLs
RIPng	PIM-SM v6	OSPFv3
PIM-SSM v6	MLDv1	PIM-SSM
IGMPv3	MLDv2	IPsec support for OSPFv3
Multi-VRF (IPv6)	802.3-2008 Clause 57 (Ethernet OAM – Link Layer OAM)	802.1Qaz ETS, (Data Center Bridging – Enhanced Transmission Selection)
Tracked Objects	Remote Port Mirror	User Tracking and Control
Zero Config - Proxy Web	IEEE 802.1ak MVRP (Multiple VLAN Registration Protocol)	VLAN Provider Bridging (Q-in-Q)
Unidirectional Link Detection	Dynamic Arp Inspection (DAI)	IEEE 802.1Q-2011 (Connectivity Fault Management)
DHCP Snooping	RADIUS Server Load Balancing	IP Source Guard
IP Service Level Agreements	Routing as a Service (RaaS)	802.1aq-2012 Shortest Path Bridging (SPBv)
VXLAN (RFC 7348)		

FIRMWARE CHANGES AND ENHANCEMENTS:

Features Requests in 8. 62.03.0006

MGBICs	Description	Introduced in Version:
MGBIC-LC01	Optic vendor add/update for Model: MGBIC-LC01 Avago AFBR-5715APZ-EN1	Unknown
	Optic vendor add/update for Model: MGBIC-LC01 Finisar FTLF8518P4BTL-EN	N 8.41.01
	Optic vendor add/update for Model: MGBIC-LC01 WTD RTX191-552-C72	N 8.41.01
	Optic vendor add/update for Model: MGBIC-LC01, 10051H Formerica TSD-S2CH1-C11M	N 8.41.01
MGBIC-LC09	Optic vendor add/update for Model: MGBIC-LC09 WTD RTX191-404-C72	N 8.41.01
	Optic vendor add/update for Model: MGBIC-LC09, 10052H Formerica TSD-S2CA1-	N 8.41.01

	F11M	
--	------	--

10051H	Description	Introduced in Version:
10051H	Optic vendor add/update for Model: 10051H WTD RTX191-552-C71	N 8.41.01
	Optic vendor add/update for Model: 10051H Finisar FTLF8518P4BTL-EX	N 8.41.01
	Optic vendor add/update for Model: 10051H Avago AFBR-5715APZ-EX1	N 8.41.01

10301 Finisar	Description	Introduced in Version:
10301	Optic vendor add/update for Model: 10301 Finisar FTLX8574D3BCL-EX	N 8.41.01

SFP+ Optics	Description	Introduced in Version:
10GB-SR-SFPP	Optic vendor add/update for Model: 10GB-SR-SFPP Finisar FTLX8574D3BCL-EN	N 8.41.01
10GB-LR-SFPP	Optic vendor add/update for Model: 10GB-LR-SFPP WTD RTX228-401-C62	N 8.41.01
10GB-F10-SFPP	New vendor Finisar parts FCBG110SD1C10-EN, and FCBG110SD1C20-EN worked for the Model: 10GB-F10-SFPP, 10GB-F20-SFPP.	N 8.62.01
10GB-F20-SFPP	New vendor Finisar parts FCBG110SD1C10-EN, and FCBG110SD1C20-EN worked for the Model: 10GB-F10-SFPP, 10GB-F20-SFPP.	N 8.62.01

10302 WTD	Description	Introduced in Version:
10302 WTD	Optic vendor add/update for Model: 10302 WTD RTX228-401-C61	N 8.41.01

10052H	Description	Introduced in Version:
10052H Finisar	Optic vendor add/update for Model: 10052H Finisar FTLF1318P3BTL-EX	N 8.41.01
10052H WTD	Optic vendor add/update for Model: 10052H WTD RTX191-404-C71	N 8.41.01

Capacity Reductions in 8.63.05.0004

Shortest Path Bridging Capacity Reductions:
Maximum number of SPBv nodes reduced from 100 to 50.
Maximum number of ECT algorithms being used reduced from 16 to 4.

Connectivity Fault Management Reductions:
Number of MEPS supported reduced from 50 to 25.
Number of MFHS supported reduced from 250 to 125.

Problems Corrected in 8.63.09.0001

Tunnel Manager Problem Corrected in 8.63.09.0001	Introduced in Version:
The reframer database manager might add an already free entry to the age list, resulting in an error or a crash when this entry is re-used.	08.32.01

Multicast Protocols Problems Corrected in 8.63.09.0001	Introduced in Version:
MSDP SA-filter in/out does not correctly filter 'excludes' in the ACL list.	07.91.01
While running MSDP, the host main slot becomes unresponsive requiring a manual reset.	07.79.01

NetFlow Problem Corrected in 8.63.09.0001	Introduced in Version:
For S chassis, and VSB K chassis, the start and end times for netflow exported records may not be accurate. They may show values up to a few minutes before or after the actual time. The inaccuracies in reported times may get larger with longer chassis up times.	07.00.01

Problems Corrected in 8.63.08.0002

Multicast Protocols Problems Corrected in 8.63.08.0002	Introduced in Version:
IP Multicast flows may not be set up in hardware, resulting in soft-forwarding of the multicast flow, which may lead to data loss.	8.31.01
With an active K-EOS-L3 license, MSDP is not supported.	07.00.01

Link Aggregation Problems Corrected in 8.63.08.0002	Introduced in Version:
If a malformed LACP PDU packet is received, the switch may reset, leaving the following message:	07.00.01
Message 9/146 Exception PPC750 Info 08.61.01.0018 11/20/2019 03:54:03 Exc Vector: DSI exception (0x00000300)	

<p>Thread Name: tLagRx Exc Addr: 0x00bdfa58 Thread Stack: 0x2e416000..0x2e40e000 Stack Pointer: 0x2e415a80 Traceback Stack: [0] 0x00000000 [1] 0x00be056c [2] 0x00be06b0 [3] 0x00be08d8 [4] 0x00bdd7f4 [5] 0x00bcfb48 [6] 0x01c7ff64 [7] 0xeeeeeeee</p> <p>GENERAL EXCEPTION INFO: srr0 : 0x00bdfa58 srr1 : 0x0000b032 dar : 0x2e51701c cr : 0x24000888 xer : 0x00000000 fpcsr:0x00000000 dsisr:0x40000000</p> <p>GENERAL REGISTER INFO: msr : 0x0000b032 lr : 0x00be056c ctr : 0x00000000 pc : 0x00bdfa58 cr : 0x24000888 xer : 0x20000000 r[0]:0x0000003f r[1]:0x2e415e00 r[2]:0x04628e80 r[3]:0x2e3009a0 r[4]:0x2e516f40 r[5]:0x2e415e4c r[6]:0x2e415e84 r[7]:0x000000d3 r[8]:0x2e516fbc r[9]:0x00000000 r[10]:0x2e516f5a r[11]:0x0000001e r[12]:0x24000284 r[13]:0x0613d5a0 r[14]:0x0629a468 r[15]:0x0629a470 r[16]:0x2e415f2c r[17]:0x2e415f30 r[18]:0x2e415f34 r[19]:0x0399ea70 r[20]:0x0399eaa0 r[21]:0x0399ea3c r[22]:0x2e415e60 r[23]:0x2e415e4c r[24]:0x00000000 r[25]:0x00000001 r[26]:0x06279100 r[27]:0x062b8ecc r[28]:0x2e3009a0 r[29]:0x2e516f40 r[30]:0x00000037 r[31]:0x000000dc</p> <p>In the message log.</p>	
---	--

VRF Problems Corrected in 8.63.08.0002	Introduced in Version:
Infrequently, after a VLAN interface is created in a multiple VR configuration, the interface will not be fully functional.	07.00.01

IP Interface Manager Problems Corrected in 8.63.08.0002	Introduced in Version:
The MIB2 in discard counter for port "host.0.1" may be incremented when it should not be.	07.00.01

Problems Corrected in 8.63.07.0003

L3 Problem Corrected in 8.63.07.0003	Introduced in Version:
<p>When router is processing high rates of DVMRP frames, any frames processed by IP stack could be dropped. This includes:</p> <ul style="list-style-type: none"> - OSPF frames. Potentially leading to enough OSPF Hello packets being dropped to cause OSPF Neighbor Drops. - All management protocols supported by IP stack (SSH, telnet, SNMP, etc). 	7.00.01

Platform Problem Corrected in 8.63.07.0003	Introduced in Version:
<p>Infrequently, there could be a hardware failure that impacts packet forwarding and leads to logging of messages similar to:</p> <pre> ===== ===== Message 1/202 Syslog Message 08.63.04.0003 07/22/2019 07:22:09 <3>Fuji[7.tFujiMon]fujiMonitoringTask: Fuji 7 TXQ CRC error check failed. ===== ===== Message 4/202 Syslog Message 08.63.04.0003 07/22/2019 07:21:49 <3>Fuji[7.tFujiMon]fujiMonitoringTask: Fuji 7 Host Rx0/Rx1 CRC error check failed </pre> <p>But the S Chassis blade, or K Chassis Linecard is not reset, in attempt to resolve the problem.</p>	7.00.01

Host Services Problem Corrected in 8.63.07.0003	Introduced in Version:
<p>Very infrequently, after a reboot, the following MIBS will not function correctly:</p> <ul style="list-style-type: none"> • LSNAT • NAT • TWCB <p>MIB GET's will not return any data and MIB SET's will fail.</p>	7.00.01

VxWorks Operating System Vulnerabilities Corrected in 8.63.07.0003	Introduced in Version:
<p>Recently, Armis Labs conducted testing on the VxWorks operating system used by K-Series. They discovered 11 different vulnerabilities to potential DOS attacks, which are documented here: https://www.armis.com/urgent11</p> <p>The following vulnerabilities have not been present in any versions of K-Series, due to the versions of VxWorks being used:</p> <p>CVE-2019-12256 IPNET security vulnerability: Stack overflow in the parsing of IPv4 packets IP options</p> <p>CVE-2019-12259 IPNET security vulnerability: DoS via NULL dereference in IGMP parsing</p> <p>CVE-2019-12260 IPNET security vulnerability: TCP Urgent Pointer state confusion caused by malformed TCP AO option</p> <p>The following vulnerabilities were present in all previous versions of K-Series, and have been fixed in the 08.63.07.0003 release:</p> <p>CVE-2019-12255 IPNET security vulnerability: TCP Urgent Pointer = 0 leads to integer underflow</p> <p>CVE-2019-12257 IPNET security vulnerability: Heap overflow in DHCP Offer/ACK parsing inside ipdhcpc</p> <p>CVE-2019-12258 IPNET security vulnerability: DoS of TCP connection via malformed TCP options</p> <p>CVE-2019-12261 IPNET security vulnerability: TCP Urgent Pointer state confusion during connect() to a remote host</p> <p>CVE-2019-12262 IPNET security vulnerability: Handling of unsolicited Reverse ARP replies (Logical Flaw).</p> <p>CVE-2019-12263 IPNET security vulnerability: TCP Urgent Pointer state confusion due to race condition.</p> <p>CVE-2019-12264 IPNET security vulnerability: Logical flaw in IPv4 assignment by the ipdhcpc DHCP client</p> <p>CVE-2019-12265 IPNET security vulnerability: IGMP Information leak via IGMPv3 specific membership report</p>	7.00.01

Problems Corrected in 8.63.06.0005

Host Services Problems Corrected in 8.63.06.0005	Introduced in Version:
For a Multi Blade S Chassis, a VSB K chassis, or a stacked 7100 switch, multiple link traps and syslog, on for each blade or switch in stack, are generated for Lag and Tunnel Bridge port link changes.	7.00.01

L3 Problems Corrected in 8.63.06.0005	Introduced in Version:
Directed broadcast copy-to function does not work. Packets are not broadcast to the copy-to vlan.	8.20.02
With large amounts of passive interfaces spread across multiple VRF's, upon reboot there maybe continual resets, leaving messages similar to: <pre>Message 77/274 Syslog Message 08.63.05.0004 06/04/2019 07:13:15 <1>DistServ[4.tDSserv5]moveToSyncR.5(Host) client 5(Host) not ready in 1 8013 (0x00ffa1a8 0x008437e0 0x0084dd20 0x0084e740 0x01c881e4 0xe0000000) in the message log.</pre>	8.63.05

Platform Problems Corrected in 8.63.06.0005	Introduced in Version:
When only one of either a VLAN based or non-vxlan based tunnel type was configured, traffic passing through the device for the other tunnel type could cause high CPU and hardware flow counts.	8.41.01

Problems Corrected in 8.63.05.0004

Management Problems Corrected in 8.63.05.0004	Introduced in Version:
Previous versions of firmware only support the following SSH Key Exchange algorithms (1024-bit keys): <pre>diffie-hellman-group1-sha1 diffie-hellman-group-exchange-sha1</pre> This release adds the following algorithm, which uses a 2048-bit key. <pre>diffie-hellman-group14-sha1</pre> This change applies to both the SSH Server (i.e., SSH from somewhere else to the switch) and SSH Client (i.e., SSH from the switch to somewhere else).	7.00.01

Auto Negotiation Problems Corrected in 8.63.05.0004	Introduced in Version:
When linked to an Extreme C35 Wireless Controller via a 1G link, sometimes the link will not come up at 1G speed, instead it will automatically drop down to 100M.	7.00.01

Hardware Problems Corrected in 8.63.05.0004	Introduced in Version:
<p>On a K chassis, if COS rate limiting is enabled, and a line card is either:</p> <ul style="list-style-type: none"> - Physically removed - Reset via command line interface (CLI) - Reset by firmware due to a hardware failure <p>Very infrequently, the K chassis may reset and leave a message similar to the following in the message log:</p> <pre> Exc Vector: DSI exception (0x00000300) Thread Name: tCosEvent Exc Addr: 0x016211a8 Thread Stack: 0x151e8000..0x151e5000 Stack Pointer: 0x151e7c80 Traceback Stack: [0] 0x0162119c [1] 0x00450470 [2] 0x024ec6c4 [3] 0x024e1cfc [4] 0x024e9138 [5] 0x01815324 [6] 0xe0000000 [7] 0x00000000 GENERAL EXCEPTION INFO: srr0 : 0x016211a8 srr1 : 0x0000b032 dar : 0xe00050b1 cr : 0x48000882 xer : 0x00000000 fpcsr:0x00000000 dsisr:0x42000000 GENERAL REGISTER INFO: msr : 0x0000b032 lr : 0x0162119c ctr : 0x01800810 pc : 0x016211a8 cr : 0x48000882 xer : 0x20000000 r[0]:0x00000000 r[1]:0x151e7e00 r[2]:0x037f21d0 r[3]:0x00000001 r[4]:0x054d1054 r[5]:0x00000001 r[6]:0x00000001 r[7]:0x00000000 r[8]:0x2df4f6d4 r[9]:0xe0005075 r[10]:0x15173338 r[11]:0x151e7c80 r[12]:0x28000882 r[13]:0x04450be0 r[14]:0x00000032 r[15]:0x00000000 r[16]:0x00000023 r[17]:0x00000000 r[18]:0x151e7f14 r[19]:0x00000000 r[20]:0x00000000 r[21]:0x00000000 r[22]:0x00000000 r[23]:0x0cb167bf r[24]:0x00000000 r[25]:0x000000ea r[26]:0x151e7e4a r[27]:0x151e7e48 r[28]:0x00000000 r[29]:0x00000000 r[30]:0x00005070 r[31]:0x29d85bb0 </pre>	7.00.01

Routing Protocols Problems Corrected in 8.63.05.0004	Introduced in Version:
When the OSPF passive-interface default command is used, after a reboot it is possible that not all interfaces will come up in passive mode	8.20.02
When OSPFv2 is configured in a non-default VRF along with a distribute-list route-map filter, after a "clear ip ospf process" is executed all routes are denied from the route table.	07.00.01

Problems Corrected in 8.63.04.0003

IPV6 Forwarding Problems Corrected in 8.63.04.0003	Introduced in Version:
If an IPV6 interface is enabled, the router will transition to layer 4 flow mode, even though it does not need to. This will lead to increased flow counts and corresponding increased CPU usage on router.	7.00.01

Problems Corrected in 8.63.03.0002

IPv6 ND Problems Corrected in 8.63.03.0002	Introduced in Version:
<p>If L3 IPV6 Interfaces, with Router Advertisements, are active, a memory leak may occur. This eventually leads to a reset, leaving a message similar to:</p> <pre> Exc Vector: DSI exception (0x00000300) Thread Name: tNet0 Exc Addr: 0x01c0f5d4 Thread Stack: 0x1f952000..0x1f94f000 Stack Pointer: 0x1f94f310 Traceback Stack: [0] 0xe0000000 [1] 0x01be4df8 [2] 0x01be5980 [3] 0x01be6044 [4] 0x01087080 </pre>	7.00.01

[5]	0x01cb2080
[6]	0x01d01be4
[7]	0x01ced1b0
[8]	0x01d027a8
[9]	0x01d02a78
[10]	0x01cf56f4
[11]	0x01cb26bc
[12]	0x01d02018
[13]	0x01cb20f0
[14]	0x01d01be4
[15]	0x01ced1b0

GENERAL EXCEPTION INFO:

srr0: 0x01c0f5d4 srr1: 0x0000b032 dar: 0x1f94efc0

cr: 0x24042224 xer: 0x00000000 fpcsr:0x00000000

dsisr:0x42000000

GENERAL REGISTER INFO:

msr: 0x0000b032 lr: 0x01be4cf4 ctr: 0x01c84d30

pc: 0x01c0f5d4 cr: 0x24042224 xer: 0x00000000

r[0]:0x6a34b2b8 r[1]:0x1f94f040 r[2]:0x04630050 r[3]:0x072ce704

r[4]:0x12bf4fd0 r[5]:0x6a34b2b0 r[6]:0x1f94f098 r[7]:0x00000010

r[8]:0xffffffff r[9]:0x302291b0 r[10]:0x1fa23fb0 r[11]:0x1f94f090

r[12]:0x01c84d30 r[13]:0x061443c0 r[14]:0x2fb82704 r[15]:0x0613dc7c

<p>r[16]:0x2fb4d9c0 r[17]:0x0381e158 r[18]:0x20dce048 r[19]:0x1f951818</p> <p>r[20]:0x0000005e r[21]:0x01cf2ed0 r[22]:0x00000000 r[23]:0x00000003</p> <p>r[24]:0x00000020 r[25]:0x00000010 r[26]:0x00000000 r[27]:0x01d02328</p> <p>r[28]:0x6b40a010 r[29]:0x00000d00 r[30]:0x6b40a010 r[31]:0x072ce6c0</p> <p>-----</p> <p>In the message log.</p>	
---	--

Multicast Protocols Problems Corrected in 8.63.03.0002	Introduced in Version:
IGMP does not process a report that is received on the proper VLAN with "fabric-routing" enabled.	7.30.01

Problems Corrected in 8.63.02.0004

IPv6 ND Problems Corrected in 8.63.02.0004	Introduced in Version:
ND router advertisement may stop after system has been up for more than 497 days.	7.31.02

UDP Forwarding Problems Corrected in 8.63.02.0004	Introduced in Version:
<p>IPv4 DHCP Relay packet sent to server with UDP SRC port of 68 instead of 67.</p> <p>Added VRF and interface cli command to set the SRC port to 67: ip dhcp relay information option port-67</p>	7.03.03

Static Routes Problems Corrected in 8.63.02.0004	Introduced in Version:
It is possible for a static route to be added with an IPv4 address and an IPv6 next hop.	7.03.03

Tracked Objects Problems Corrected in 8.63.02.0004	Introduced in Version:
ICMP probe may stop sending ICMP requests and stay "up" causing it to report an invalid status.	8.20.02

NetFlow Problems Corrected in 8.63.02.0004	Introduced in Version:
For flows on a GRE tunnel, if the inner (customer) header is non-IP, netflow records with invalid L3 and L4 parameters are generated.	8.41.01

Management Problems Corrected in 8.63.02.0004	Introduced in Version:
When an authenticated SSH client sends a specific sequence of SSH messages to the switch's SSH server, the switch will core. This set SSH messages is atypical and can only be sent after SSH client has authenticated. Therefore SSH clients without authorization credentials cannot exploit this defect.	8.01.01

DHCP Problems Corrected in 8.63.02.0004	Introduced in Version:
With DHCP server configured, receipt of a DHCP request with client id greater than 22 bytes may cause a system reset.	7.40.00

Packet Dispatch Problems Corrected in 8.63.02.0004	Introduced in Version:
Router wide rate limiters for OSPF Hello frames have been increased from 300 to 500 packets per second, allowing router to now process up to 500 OSPF hellos per second.	8.20.02

Problems Corrected in 8.63.01.0019

Management Problems Corrected in 8.63.01.0019	Introduced in Version:
The switch's embedded SSH server is susceptible to a double free that will cause the switch to core. The more frequently an application or user connects and disconnects to a switch via SSH, the more susceptible they are to this defect.	7.00.01

Boot Config Problems Corrected in 8.63.01.0019	Introduced in Version:
When performing a config reset on a bonded K chassis all the line cards may not be initialized before the system is configured, resulting in port strings not being valid.	8.31.02

Netflow Problems Corrected in 8.63.01.0019	Introduced in Version:
<p>For flows with mac in mac headers present, previously no netflow records would have been generated. Now will generate netflow records for the inner (customer) L3/L4 headers. In addition:</p> <ul style="list-style-type: none"> - If exporting of MAC' is enabled, the CMAC's from the mac in mac headers will be the MAC's that are exported. - The ISID from the mac in mac header will be exported. Netflow does not currently have and fields defined for this. <p>The existing layer2SegmentId (Information Element 351) will be used for this. A value of 10 in the upper byte of that field will be used to indicate a mac in mac header. The ISID will be in the lower 3 bytes of that field.</p>	7.00.01
Netflow - fragmented IPV4 and IPV6 flows will have netflow records with invalid L4 source and destination flows exported.	7.00.01

Netflow - If export of vlans is enabled, for non-routed flows the egress vlan field in netflow records may not contain a valid value.	7.00.01
Enabling netflow leads building of flows past tunnel headers (Vxlan, GRE, MPLS, Mac in Mac, etc) to the inner customer headers. And netflow records are exported for the inner headers. In some environments that may lead negative consequences like exceeding flow limits, high CPU usage, and dropped packets. A new CLI command has been added: <code>set netflow export-inner-headers [enable/disable]</code> to give user control over weather flows are built past tunnel headers (and netflow records exported for inner headers).	8.41.01
Netflow records exported for inner headers of GRE tunnel encapsulated flows may contain invalid data.	8.41.01

Routing Protocols Problems Corrected in 8.63.01.0019	Introduced in Version:
When running OSPF with a large number of link state advertisements, an assert may occur in tRtrPtcls thread, with the following in the log message, "SMS assert in qodmuti2.c at line 237 : != database_en□try_ptr->entry_status 0 QODM_STATUS_UPDATE 0"	7.00.01

Static Routes Protocols Problems Corrected in 8.63.01.0019	Introduced in Version:
A recursive route may not display correctly if preceded by a VRF route.	8.31.01

BFD Problems Corrected in 8.63.01.0019	Introduced in Version:
The BFD protocol may not function correctly when running on oversubscribed interfaces. The protocol may go down. If attached to an OSPF interface this will cause OSPF to bounce.	8.31.01

Stacking Problems Corrected in 8.63.01.0019	Introduced in Version:
In Software VSB K series switches, and Software VSB S series switches, if ports become over-subscribed, OSPF could become unstable bounce.	7.60.01

Problems Corrected in 8.62.04.0001

Multicast Protocols Problems Corrected in 8.62.04.0001	Introduced in Version:
IP Multicast is not supported on tunnel ingress, however in some cases it is being processed and if PIM is enabled on that VLAN memory leaks could occur.	8.21.01

Tunnel Manager Problems Corrected in 8.62.04.0001	Introduced in Version:
With more than one L2 tunnel configured or a mirror tunnel configured resets with message similar to: Fuji[2.tNimIntr]Switch Chip 0 (Slot 2 Mainboard) detected fatal condition Fuji[2.tNimIntr]Fuji EDF MAIN intr: Fuji=0, Adr=0, Reg=0x00010000	8.62.03

ARP Problems Corrected in 8.62.04.0001	Introduced in Version:
The host may not send an ARP request for a secondary address.	8.62.03

Problems Corrected in 8.62.03.0006

Tunnel Manager Problems Corrected in 8.62.03.0006	Introduced in Version:
Routing between VXLAN tunnels is supported.	8.41.01
Receipt of an ARP or Neighbor Solicitation through an L2 tunnel may cause a reset.	8.61.01

Multicast Protocols Problems Corrected in 8.62.03.0006	Introduced in Version:
PIM protocol packet rateLimiter prevents maintaining neighbor adjacencies.	7.00.01

Security Problems Corrected in 8.62.03.0006	Introduced in Version:
<p>Updated cryptography library from OpenSSL 1.0.2h to 1.0.2k. Details of the security fixes can be found on OpenSSL's website:</p> <p>URL:</p> <p>https://www.openssl.org/news/cl102.txt</p> <p>Applicable Sections:</p> <p>Changes between 1.0.2j and 1.0.2k [26 Jan 2017]</p> <p>Changes between 1.0.2i and 1.0.2j [26 Sep 2016]</p> <p>Changes between 1.0.2h and 1.0.2i [22 Sep 2016]</p> <p>Vulnerabilities Resolved:</p> <p>(CVE-2017-3731)</p> <p>(CVE-2017-3732)</p> <p>(CVE-2016-7055)</p> <p>(CVE-2016-7052)</p> <p>(CVE-2016-6304)</p> <p>(CVE-2016-2183)</p> <p>(CVE-2016-6303)</p> <p>(CVE-2016-6302)</p> <p>(CVE-2016-2182)</p> <p>(CVE-2016-2180)</p> <p>(CVE-2016-2177)</p> <p>(CVE-2016-2178)</p> <p>(CVE-2016-2179)</p> <p>(CVE-2016-2181)</p> <p>(CVE-2016-6306)</p>	8.62.01

Memory failures on K-Series line cards may result in log messages similar to the following: <3>bcmStrat[11.tLcIntr]MEM_FAIL interrupt occurred on chip 4! Repeating failures may produce continuous log messages and may also result in high CPU usage.	8.32.01
---	---------

ARP Problems Corrected in 8.62.02.0022	Introduced in Version:
At times ARPs for directed Broadcast packets may not be sent out to all interfaces.	8.42.04

Port Status/Control Problems Corrected in 8.62.02.0022	Introduced in Version:
A message indicating board replacement is necessary because MAC error counters may be reported unnecessarily.	8.42.04

Mirroring Problems Corrected in 8.62.02.0022	Introduced in Version:
During a system boot, an "out-of-sync mirror detected" error may result if there are several mirror ports configured.	8.20.02

Management Problems Corrected in 8.62.02.0022	Introduced in Version:
Add a new command to display the underlying port of a lag : show lacp underlyingPort <lagPortString> { da-sa { <da-mac> <sa-mac> }dip-sip { <dip> <sip> } }}	7.00.01
Example: show lacp underlyingPort lag.0.20 dip-sip 172.22.71.2 169.254.2.2 Port [ge.2.20] will be used for that flow.	
If on a bonded system each chassis is analyzed for the underlying Port: show lacp underlyingPort lag.0.20 dip-sip 172.22.71.2 169.254.2.2 Port [ge.2.20] will be used for that flow if ingressing chassis 1. Port [ge.2.21] will be used for that flow if ingressing chassis 2.	
When TLS parameters are not configured, the outputs of "show config" and "show config tls" contain incomplete "set tls" commands. If these incomplete commands are fed back to the switch (e.g. via "configure <filename>") they will be ignored and the switch will correctly retain the default (i.e., empty) TLS configuration.	8.61.01
The "webview" feature has been removed. The "show webview" and "set webview" CLI commands have been removed, and the switch will no longer listen for or accept webview connections.	7.00.01

ACL Problems Corrected in 8.62.02.0022	Introduced in Version:
An interface ACL on multi slots configured with different 'all-traffic' option may not sync up correctly resulting in the option being present or not depending on which slot is master.	8.11.01

NodeAlias Problems Corrected in 8.62.02.0022	Introduced in Version:
Node and Alias BPDU entry types may display invalid VLAN for CLI and MIB.	7.00.01
The ctAliasAddressText field for all the Node and Alias MIB tables used is not formatted correctly for IPV6 entries. It is in IPV4 format, rather than IPV6 format.	7.00.01

SYSLOG Problems Corrected in 8.62.02.0022	Introduced in Version:
In policy syslogs, which rule was hit may not be included.	8.32.01

DHCP Problems Corrected in 8.62.02.0022	Introduced in Version:
A client DHCP packet without an IPDHCP_OPTCODE_END "255" could result in a reset.	N 8.01.01

Features Enhancements 8.62.01.0034

802.1D Filter Database Enhancement in 8.62.01.0034
A customer configurable mode for EOS which will disable learning the Source MAC address of control frames has been added. With control-frame-learning disabled the source MAC address for control frames will not be learned. Control-frame-learning mode is configurable for the entire chassis.

Remote Port Mirror Enhancements in 8.62.01.0034
Remote tunnel mirroring CPU usage reduced by half.

VXLAN Enhancements in 8.62.01.0034
VLAN names can now be used to map VNIs to VLANs.

Problems Corrected in 8.62.01.0034

ACL Problems Corrected in 8.62.01.0034	Introduced in Version:
While running a <code>show config</code> command, an ACL mismatch may cause the comment to enter a loop and display the following error message: "Unable to perform access list entry show config".	7.63.01
A policy ACL can be created with the same name as an IPv4 extended ACL.	8.32.02

CLI Problems Corrected in 8.62.01.0034	Introduced in Version:
While displaying Port VLAN information, an extra "," may be placed at the end of the string.	7.03.03
When display port egress the -verbose option may display the wrong vlan for non static vlans.	7.91.01

Cryptography Problems Corrected in 8.62.01.0034	Introduced in Version:
The switch's cryptographic library pseudo-random number generator (PRNG) is seeded with random data from various sources during bootup. The entropy (randomness) of the data provided by one of these sources was being overestimated.	7.40.01
ICMP Problems Corrected in 8.62.01.0034	Introduced in Version:
The help for ping and traceroute (interface) option is not descriptive.	7.01.02
IPv4 Forwarding Problems Corrected in 8.62.01.0034	Introduced in Version:
Recognizing an IP address as both a source and destination address may incorrectly produce an ICMP error.	8.32.01
Management Problems Corrected in 8.62.01.0034	Introduced in Version:
The switch generates a new SSH Server hostkey whenever a new system boots, whenever the hostkey type is changed to/from DSA/RSA, whenever the the security mode is switched from non-FIPS to FIPS and vice versa, and whenever a "set ssh hostkey reinitialize" command is executed. SSH Clients attempting to connect to the switch must explicitly trust each new hostkey. To prevent a Man-in-the-Middle (MITM) attack the SSH Server should provide its hostkey (or a fingerprint of the hostkey) out-of-band to SSH Client users.	7.00.01
The switch now facilitates an out-of-band fingerprint check. The "show ssh" command displays the hostkey's type (DSA or RSA) and fingerprint (MD5 hash in non-FIPS mode, SHA1 hash in FIPS mode).	
To counteract MITM attacks users should connect to a new switch via a console, enter the command "show ssh", make a note of the hostkey fingerprint, and then verify that this fingerprint matches the fingerprint of the hostkey presented during the SSH handshake.	
If the final 10gig port does not have an egress assigned lag ports may not display with the command "show port egress -v".	7.91.01
Memory is leaked when an SSH client attempts to connect to the switch and the Diffie-Hellman key exchange does not succeed. Numerous key exchange failures deplete memory and cause the switch to reset.	7.00.01
Mirroring Problems Corrected in 8.62.01.0034	Introduced in Version:
When a tx port mirror is created or deleted, and it is not the first port in the mirror (i.e. is a many to one mirror), the mirror create or delete action will not be applied to active flows.	7.30.01
Multicast Protocol Problems Corrected in 8.62.01.0034	Introduced in Version:
IGMP/MLD protocol packets are not forwarded to querier and/or router on a LAG port after physical ports of the LAG have state/forwarding change.	8.01.01
NAT Problems Corrected in 8.62.01.0034	Introduced in Version:
If an ICMP error packet matches a NAT list rule the inner packet may not be properly natted back to the original source.	7.00.01

LLDP Problems Corrected in 8.62.01.0034	Introduced in Version:
"PSE Allocated Power Value" returned in "Power Via MDI" LLDP TLV might be greater than "PD Requested Power Value" potentially causing some PDs to ignore the allocated power.	7.00.01
LLDP commands that set port tx-tlv to "all" are not executed.	8.41.01

Tunneling Problems Corrected in 8.62.01.0034	Introduced in Version:
If all the boards' reframer databases are not matching during a chassis sync, it is possible that connections would not get removed when updating entries. This would cause flows using those connection to go out malformed or not transmit.	8.61.01
Routing into or from a VLAN where the packet destination or source resides across an L2 VXLAN tunnel may result in a packet drop.	8.61.01

VRF Problems Corrected in 8.62.01.0034	Introduced in Version:
If all ports of a chassis bond are severed when the bond is reestablished the VRRP may not be properly configured.	7.40.01
<p>In rare cases, when a Virtual Router is being removed and VRRP is active on interface(s) in that router, a blade may reset, leaving a message similar to the following in the message log:</p> <pre> Exc Vector: DSI exception (0x00000300) Thread Name: tVrrpRX Exc Addr: 0x0058cb70 Thread Stack: 0x25a74000..0x25a70000 Stack Pointer: 0x25a73c40 Traceback Stack: [0] 0x0058c924 [1] 0x0057d188 [2] 0x01af23e4 [3] 0xeeeeeeee GENERAL EXCEPTION INFO: srr0 : 0x0058cb70 srr1 : 0x0000b032 dar : 0x0000010c cr : 0x44000822 xer : 0x00000000 fpcsr:0x00000000 dsisr:0x0a000000 GENERAL REGISTER INFO: msr : 0x0000b032 lr : 0x0058cb40 ctr : 0x00000000 pc : 0x0058cb70 cr : 0x44000822 xer : 0x00000000 r[0]:0x00000003 r[1]:0x25a73c40 r[2]:0x042ce870 r[3]:0x0000ca69 r[4]:0x4f129ce2 r[5]:0x0000000c r[6]:0x25a73d20 r[7]:0x00000010 r[8]:0x00000003 r[9]:0x00023594 r[10]:0x00000000 r[11]:0x00000008 r[12]:0x24000844 r[13]:0x05dd41f0 r[14]:0x0354d680 r[15]:0x25a73e6c r[16]:0x00000020 r[17]:0x0000000c r[18]:0x00000000 r[19]:0x00000000 r[20]:0x00000000 r[21]:0x00000016 r[22]:0x00000000 r[23]:0x00000001 r[24]:0x0000004b r[25]:0x00000000 r[26]:0x00000000 r[27]:0x4f129cc2 r[28]:0x0000000c r[29]:0x00000014 r[30]:0x4f129cd6 r[31]:0x258c3bb0 ----- </pre>	8.31.01

VSB Problems Corrected in 8.62.01.0034	Introduced in Version:
In rare cases, a software VSB chassis may have low throughput between VSB bonded chassis.	7.40.00

Problems Corrected in 8.61.02.0001

Management Problems Corrected in 8.61.02.0001	Introduced in Version:
Failed SSH client attempts cause a memory leak. Numerous key exchange failures deplete memory and cause the switch to reset.	7.00.01
Node Alias Problems Corrected in 8.61.02.0001	Introduced in Version:
Node Alias MDNS, LLMNR, and SSDP entries are not recognized properly in IPv6 packets.	8.11.01

Features Enhancements 8.61.01.0018

Transceiver Enhancements in 8.61.01.0018
Added support for 10G tunable DWDM single mode SFP+ transceiver (part number 10325).
VXLAN Overlay Enhancements in 8.61.01.0018
Extensions have been added to OSPF that allow automatic discovery of VXLAN VTEPs and VNIs.

Problems Corrected in 8.61.01.0018

ACL Problems Corrected in 8.61.01.0018	Introduced in Version:
While restoring an <i>ip access-group</i> ACL within an interface in a non-global VRF, the restoration may fail and display a message similar to the following: "Failed to restore the apply of ipv4 access list in non global VRF ifindex 38 vr 7"	8.21.01
Broadcast Problems Corrected in 8.61.01.0018	Introduced in Version:
If the router receives a subnet broadcast and the Layer 2 destination address is also broadcast, the router does not forward the frame.	7.63.01
Management Problems Corrected in 8.61.01.0018	Introduced in Version:
Scheduling a system reset more than 248 days in advance crashes and resets the system immediately. Resets scheduled using the CLI (<i>reset at <hh:mm> [<mm/dd>] [reason]</i>) as well as any delayed configuration management change operation configured using SNMP (ENTERASYS-CONFIGURATION-MANAGEMENT-MIB's <i>etsysConifgMgmtChangeDelayTime</i> object) are susceptible to this issue. To resolve the issue, the system does not allow delays longer than 248 days.	7.00.01
On the command line, if you enter '!' followed by some special character (not alpha or numeric; for example "#"), unexpected output may be echoed to your session.	1.07.19
The SSH protocol allows an SSH client to specify an optional command that is to be executed on the remote host. For example: <i>ssh <username>@<hostname> [<command> <arg1> <arg2> ...]</i> The SSH server on EOS switches never executes any supplied command. However, a requested command that contains 20 or more arguments causes the switch to stop responding. This issue only occurs if user authentication succeeds; therefore, unauthorized users cannot cause this issue on a switch.	7.00.01

OSPFv3 Problems Corrected in 8.61.01.0018	Introduced in Version:
OSPFv3 may not originate a new Intra-Area-Prefix-LSA if the set of addresses for an active OSPF interface is changed.	8.31.01
Tunneling Problems Corrected in 8.61.01.0018	Introduced in Version:
The IPv6 tunnel code is looking for Traffic Class, when it needs to check for Traffic Class or TOS to encapsulate the packet, causing malformed packets when tunneling ICMPv4 using IPv6 GRE with TOS rewrite.	8.21.01
Tunneled packets looping through an encapsulating device may cause a crash.	8.21.01
Transceiver Problems Corrected in 8.61.01.0018	Introduced in Version:
The MGBIC-02 copper SFP fails to link in 10G ports.	8.32.01

Problems Corrected in 8.42.06.0001

System Problems Corrected in 8.42.06.0001	Introduced in Version:
<p>If you are using a previous version, in rare instances a linecard may halt and display a message on the console similar to the following: This device has encountered a hardware failure. Petra[0] exceeded maximum soft resets</p> <p>The causes of this are mostly recoverable events. For these recoverable events the firmware will no longer halt, and continue to carry out normal operation. For non-recoverable events, the firmware will now reset (not halt), and log a message similar to the following:</p> <p>Dune[11.tConsole]assert: /firmware/common/dune/08_30_23/src/petraChip .cxx:4651 Petra[0] exceeded maximum soft resets, cause = PETRA_IPT_CRC_ERR_PKT</p> <p>If this occurs, contact Extreme Networks GTAC.</p>	7.03.03

Problems Corrected in 8.42.05.0003

CLI Problems Corrected in 8.42.05.0003	Introduced in Version:
While displaying Port VLAN information, an extra "," may be placed at the end of the string.	7.03.03
Management Problems Corrected in 8.42.05.0003	Introduced in Version:
If the final 10gig port does not have an egress assigned, lag ports may not display with the command "show port egress -v".	7.91.01
Multicast Protocol Problems Corrected in 8.42.05.0003	Introduced in Version:
IGMP/MLD protocol packets are not forwarded to querier and/or router on a LAG port after physical ports of the LAG have state/forwarding change.	8.01.01

VRF Problems Corrected in 8.42.05.0003	Introduced in Version:
If all ports of a chassis bond are severed when the bond is reestablished, the VRRP may not be properly configured.	7.40.01
<p>In rare cases, when a Virtual Router is being removed and VRRP is active on interface(s) in that router, a blade may reset, leaving a message similar to the following in the message log:</p> <pre> Exc Vector: DSI exception (0x00000300) Thread Name: tVrrpRX Exc Addr: 0x0058cb70 Thread Stack: 0x25a74000..0x25a70000 Stack Pointer: 0x25a73c40 Traceback Stack: [0] 0x0058c924 [1] 0x0057d188 [2] 0x01af23e4 [3] 0xeeeeeeee GENERAL EXCEPTION INFO: srr0 : 0x0058cb70 srr1 : 0x0000b032 dar : 0x0000010c cr : 0x44000822 xer : 0x00000000 fpcsr:0x00000000 dsisr:0x0a000000 GENERAL REGISTER INFO: msr : 0x0000b032 lr : 0x0058cb40 ctr : 0x00000000 pc : 0x0058cb70 cr : 0x44000822 xer : 0x00000000 r[0]:0x00000003 r[1]:0x25a73c40 r[2]:0x042ce870 r[3]:0x0000ca69 r[4]:0x4f129ce2 r[5]:0x0000000c r[6]:0x25a73d20 r[7]:0x00000010 r[8]:0x00000003 r[9]:0x00023594 r[10]:0x00000000 r[11]:0x00000008 r[12]:0x24000844 r[13]:0x05dd41f0 r[14]:0x0354d680 r[15]:0x25a73e6c r[16]:0x00000020 r[17]:0x0000000c r[18]:0x00000000 r[19]:0x00000000 r[20]:0x00000000 r[21]:0x00000016 r[22]:0x00000000 r[23]:0x00000001 r[24]:0x0000004b r[25]:0x00000000 r[26]:0x00000000 r[27]:0x4f129cc2 r[28]:0x0000000c r[29]:0x00000014 r[30]:0x4f129cd6 r[31]:0x258c3bb0 </pre>	8.31.01

Problems Corrected in 8.42.04.0016

ACL Problems Corrected in 8.42.04.0016	Introduced in Version:
While running a <i>show config</i> command, an ACL mismatch may cause the comment to enter a loop and instead display the error, "Unable to perform access list entry show config".	7.63.01
A policy ACL can be created with the same name as an IPv4 extended ACL.	8.32.02

Directed Broadcast Problems Corrected in 8.42.04.0016	Introduced in Version:
If the router receives a subnet broadcast packet which also has a Layer 2 broadcast destination address, the router will not forward the frame.	7.63.01

Host Services Problems Corrected in 8.42.04.0016	Introduced in Version:
Scheduling a system reset greater than 248 days in advance causes the system to reset. Resets scheduled using the command <i>reset at <hh:mm> [<mm/dd>] [reason]</i> as well as any delayed configuration management change operations configured using SNMP (ENTERASYS-	7.00.01

Host Services Problems Corrected in 8.42.04.0016	Introduced in Version:
CONFIGURATION-MANAGEMENT-MIB's etsysConifgMgmtChangeDelayTime object) cause this problem.	
IPv4 Forwarding Problems Corrected in 8.42.04.0016	Introduced in Version:
Recognizing an IP address as both a source and destination address may incorrectly produce an ICMP error.	8.32.01
LLDP Problems Corrected in 8.42.04.0016	Introduced in Version:
LLDP CLI commands that set port tx-tlv to "all" are not executed.	8.41.01
"PSE Allocated Power Value" returned in "Power Via MDI" LLDP TLV might be greater than "PD Requested Power Value" potentially casing some PDs to ignore the allocated power.	7.00.01
Management Problems Corrected in 8.42.04.0016	Introduced in Version:
Entering '!' in the command line followed by a special character (not alpha or numeric) may produce unexpected output echoed to the session.	1.07.19
The SSH protocol allows an SSH client to specify an optional command that is to be executed on the remote host. For example: <i>ssh <username>@<hostname> [<command> <arg1> <arg2> ...]</i>	7.00.01
The SSH server on EOS switches does not execute these commands. However, a requested command that contains 20 or greater arguments causes the switch to crash. The crash only occurs if user authentication succeeds, therefore unauthorized users cannot crash a switch in this manner.	
Memory is leaked when an SSH client attempts to connect to the switch and the Diffie-Hellman key exchange does not succeed. Numerous key exchange failures deplete memory and cause the switch to reset.	7.00.01
VSB Problems Corrected in 8.42.04.0016	Introduced in Version:
Rarely, a software VSB chassis may observe low throughput between VSB bonded chassis.	7.40.00

Problems Corrected in 8.42.03.0006

ACL Problems Corrected in 8.42.03.0006	Introduced in Version:
If an L2 access list has 3 or more rules, any rule specifying a destination MAC address is not matched correctly.	8.20.06
BFD Problems Corrected in 8.42.03.0006	Introduced in Version:
BFD sessions using a LAG port are not re-established after failover.	8.31.01
BFD removing a probe and re-adding the same probe does not always recover the session.	8.31.02

BFD Problems Corrected in 8.42.03.0006	Introduced in Version:
BFD session transitions to the DOWN state and causes the routing protocols to flap when the device modifies the clock for daylight savings time. The same situation occurs if you modify the clock using the <i>set time</i> command.	8.31.02
Host Services Problems Corrected in 8.42.03.0006	Introduced in Version:
The <i>show port transceiver</i> command does not display data for chassis 2 in a bonded K6 system.	8.31.01
Running CLI commands displaying router configuration might cause the host management to become unresponsive/locked.	8.20.02
ICMP Problems Corrected in 8.42.03.0006	Introduced in Version:
ICMP redirects are offered to hosts with different subnets. ARPs with the source and destination on the same interface, but with different subnets, send out ICMP redirects.	8.32.02
ICMP redirects are sent to incorrect VLAN/ES destination.	8.31.01
MultiAuth Problems Corrected in 8.42.03.0006	Introduced in Version:
MultiAuth port mode changes may cause authentication to stop working on LAG port(s).	8.21.01
Multicast Protocol Problems Corrected in 8.42.03.0006	Introduced in Version:
When a downstream interface is deleted first, and then an upstream interface is deleted, DVMRP may crash causing reset.	8.20.02
PIM-SM may build an incorrect join/prune message resulting in a source being pruned when it should be joined.	8.11.01
NLB Problems Corrected in 8.42.03.0006	Introduced in Version:
NLB traffic is dropped by host-access ACLs.	8.32.01
Tunneling Problems Corrected in 8.42.03.0006	Introduced in Version:
When a GRE tunnel is enabled, the blade that the tunnel port is on may reset.	8.20.02
VSF Problems Corrected in 8.42.03.0006	Introduced in Version:
If software chassis bonding is enabled on an SFP+ port of a K-Series fabric card, and an SFP module is present in a neighboring SFP+ port, then one or more of the SFP+ ports may fail to properly forward packets.	8.42.01

Features Enhancements 8.42.02.0012

Spanguard Enhancement in 8.42.02.0012	
The Spanguard feature is enhanced by the addition of a configurable setting (by CLI and SNMP) that controls the locking behavior on link loss. When enabled, link loss clears the lock. When disabled, link loss has no impact on the lock status.	

Problems Corrected in 8.42.02.0012

802.1D Filter Database Problems Corrected in 8.42.02.0012	Introduced in Version:
MAC addresses that should be learned in the filtering database are not learned and packets destined to those MAC addresses are flooded rather than forwarded using unicast.	7.01.04
Management Problems Corrected in 8.42.02.0012	Introduced in Version:
Transceiver information may not be updated for up to 10 minutes after link up or link down events on the port.	8.31.01
SSH sessions occasionally stop responding. After four sessions unresponsive sessions, the switch rejects all SSH and Telnet connection attempts. If this happens, you can only connect to the switch through the console port. Resetting the switch fixes this problem. Frequent SSH connections/disconnections by applications or users increases the occurrence of this problem.	7.00.01
Multicast Problems Corrected in 8.42.02.0012	Introduced in Version:
PIM-DM: After configuring pim dense-mode and rebooting, the PIM operating mode is restored as sparse-mode.	8.41.01
PIM-SM IPv4: After configuring <i>ip pim graceful-restart</i> and rebooting, the graceful-restart setting is not restored.	8.41.01
RADIUS Problems Corrected in 8.42.02.0012	Introduced in Version:
Receiving corrupted RADIUS frames may cause improper processing of future RADIUS requests.	7.00.01
VSB Problems Corrected in 8.42.02.0012	Introduced in Version:
If software VSB is enabled on a KK2008-0204 I/O module, SFP optics connected to one or more of the module's SFP+ ports prior to system initialization may fail to properly link up. To fix this problem, remove, and then reinsert, each SFP module after the system has booted.	8.41.01

Problems Corrected in 8.42.01.0007

Host Problems Corrected in 8.42.01.0007	Introduced in Version:
entPhySensorValue corresponding to ambient-temp-sensor-1 might not reflect current ambient temperature.	7.60.01
Management Problems Corrected in 8.42.01.0007	Introduced in Version:
MIB walks of the ctAliasMacAddressTable and ctAliasProtocolAddressTable may not return all present and active node and alias entries. NeSight Compass relies on the ctAliasMacAddressTable to display all node and alias entries, thus Compass may not function properly.	8.01.01

Management Problems Corrected in 8.42.01.0007	Introduced in Version:
When MIB walking ctAliasMIBAddress table, occasionally an IP entry with a invalid address of 0.0.0.0 might be returned.	8.01.01
Node and alias entries that correctly appear on the ports that they are received on also incorrectly appear as being received on host port (host.0.1).	7.00.01
Occasionally, when ports are disabled for node and alias processing, some entries still appear on those ports.	2.00.13

PoE Problems Corrected in 8.42.01.0007	Introduced in Version:
PoE controller might become inaccessible and not recover until module reset.	7.00.01

Spanning Tree Problems Corrected in 8.42.01.0007	Introduced in Version:
Bad BPDUs may be processed since CRC errors are not checked on BPDUs delivered to the Spanning Tree process.	7.00.01

Tunneling Problems Corrected in 8.42.01.0007	Introduced in Version:
You cannot configure more than 32 remote VXLAN VTEP IP addresses in aggregate on a single switch.	8.41.01

Features Enhancements 8.41.01.0004

VXLAN Encapsulation Enhancements in 8.41.01.0004
Support for VXLAN encapsulation has been added to the IP tunneling feature set. VXLAN encapsulation can be used as Layer 2 data center interconnect solution, or as a small-scale L2 fabric overlay.

Cryptography Enhancements in 8.41.01.0004
<p>The switch now allows AES CTR Ciphers.</p> <p>-----</p> <p>The allowed ciphers and allowed MACs lists used by the switch's SSH client and SSH server are hardcoded as follows:</p> <p>Ciphers: aes128-cbc,aes192-cbc,aes256-cbc,3des-cbc,none,blowfish-cbc, cast128-cbc,rijndael-cbc@lysator.liu.se</p> <p>MACs: hmac-sha1-96,hmac-sha1,hmac-md5,hmac-md5-96,hmac-ripemd160, hmac-ripemd160@openssh.com</p> <p>One (1) cipher has been removed from SSH: none ("none" cipher is used to bypass encryption)</p> <p>Three (3) new ciphers have been added to SSH: aes128-ctr AES in Counter mode, with 128-bit key aes192-ctr AES in Counter mode, with 192-bit key aes256-ctr AES in Counter mode, with 256-bit key</p> <p>Five (5) new Encrypt-then-MAC (ETM) MACs have been added to SSH:</p>

Cryptography Enhancements in 8.41.01.0004

hmac-sha1-etm@openssh.com:

SHA-1 with 20-byte digest and key length, encrypt-then-mac

hmac-md5-etm@openssh.com:

MD5 with 16-byte digest and key length, encrypt-then-mac

hmac-ripemd160-etm@openssh.com:

RIPEDM-160 algorithm with 20-byte digest length, encrypt-then-mac

hmac-sha1-96-etm@openssh.com:

SHA-1 with 20-byte key length and 12-byte digest length, encrypt-then-mac

hmac-md5-96-etm@openssh.com:

MD5 with 16-byte key length and 12-byte digest length, encrypt-then-mac

Additionally, both the allowed cipher list and allowed MAC list used by the SSH client and SSH server are now configurable using the CLI:

```
set ssh ciphers <cipher-list> (list is in order of precedence from high to low)
set ssh macs <macs-list> (list is in order of precedence from high to low)
clear ssh ciphers          (i.e., revert to default ciphers list)
clear ssh macs            (i.e., revert to default MACs list)
```

The default values for these lists contain all possible ciphers or MACs.
Names with an asterisk indicate not supported in FIPS mode:

Allowed Ciphers List (default):

```
aes128-ctr, aes192-ctr, aes256-ctr, aes128-cbc, aes192-cbc,
aes256-cbc, 3des-cbc, blowfish-cbc*, cast128-cbc*,
rijndael-cbc@lysator.liu.se*
```

Allowed MACs List (default):

```
hmac-sha1-etm@openssh.com, hmac-md5-etm@openssh.com*,
hmac-ripemd160-etm@openssh.com*, hmac-sha1-96-etm@openssh.com,
hmac-md5-96-etm@openssh.com*, hmac-sha1, hmac-md5*, hmac-ripemd160*,
hmac-ripemd160@openssh.com*, hmac-sha1-96, hmac-md5-96*
```

Netflow Enhancements in 8.41.01.0004

Additional support has been added for encapsulated traffic to include NetFlow support for multi-label MPLS and VXLAN-encapsulated traffic

SPB CLI Enhancements in 8.41.01.0004

CLI support has been added to configure hello interval and multiplier parameters per port:

```
set spb port <port-string> hello-interval
set spb port <port-string> hello-multiplier
```

“show support” CLI Enhancements in 8.41.01.0004

"show flowlimit stats", which shows flow stats per port, is now present in "show support" output.

VLAN CLI Enhancements in 8.41.01.0004

Added support for "show vlan fid <fid>" command.

Problems Corrected in 8.41.01.0004

ACL Problems Corrected in 8.41.01.0004	Introduced in Version:
If ACL logging is enabled on a policy ACL, it causes the policy ACL to be persisted as an extended ACL. On reboot, the ACL is restored as an extended ACL and the "set-dscp" action is missing. To recover from this, remove the ACL and re-create it. 8.32.01	8.32.01
When ACL logging is enabled on a Policy ACL, the Policy ACL specific field "set-dscp <value>" was not displayed in the log message.	8.32.01
ARP Problems Corrected in 8.41.01.0004	Introduced in Version:
For K-Series routers the ARP/ND limits are incorrect.	8.31.01
Distributed Services Problems Corrected in 8.41.01.0004	Introduced in Version:
Module might reset with messages similar to: "Chassis coherency timeout exceeded".	7.62.07
After a denial of service attack, in a multi-slot configuration, the 'dir' command only produces a list of the files on a single slot.	8.20.02
Chassis might experience stability/distribution issues during DoS LAN attack.	8.20.02
Denial of service (DoS) attack results in warning messages: "this server has been invalidated" printed to the console.	1.07.19
Flow Limiting Management Problems Corrected in 8.41.01.0004	Introduced in Version:
Flow limiting, limits, have actions applied when flow counts reach 1 less then configured limits.	1.07.19
GVRP Problems Corrected in 8.41.01.0004	Introduced in Version:
GVRP may fail to propagate dynamic VLANs on a LAG following a topology change. The result is that the switch on the remote side of the LAG fails to add the LAG to the tagged VLAN egress list. The only way to recover from this failure is to disable, and then re-enable, the LAG.	7.00.01
Hardware Problems Corrected in 8.41.01.0004	Introduced in Version:
Messages similar to the following might appear causing packets being dropped: - <163>Jan 29 13:06:59 100.10.10.22 Dune[1.dTcmTask]Petra[0] Received Interrupt PB_IHB_INVALID_DESTINATION_VALID instance 0, count 3, value=0x13deb - <3>Dune[1.dTcmTask]Petra[0] Received Interrupt PB_IHB_INVALID_DESTINATION_VALID instance 0, count 2159, value=0x1 - <165>Jun 5 11:32:19 100.10.10.22 Dune[16.tDuneErrM]Petra[0] Interrupt PB_IHB_INVALID_DESTINATION_VALID instance 0 still active - <165>Jun 5 11:32:29 100.10.10.22 Dune[11.tDuneErrM]Petra[0] Interrupt PB_IHB_INVALID_DESTINATION_VALID instance 0 is off	8.11.01
System might enter into reset loop and display message similar to: "Dune[11.tRootTask]Err_id=0x14c65001: error in petra_mgmt_all_ctrl_cells_enable_write() ExitPlace (45) Params(0,0,0,0,0)"	7.30.01

IGMP Problems Corrected in 8.41.01.0004	Introduced in Version:
When running in provider bridge mode, IGMP queries are not transmitted properly.	8.32.01
LLDP Problems Corrected in 8.41.01.0004	Introduced in Version:
LLDP sends incorrect requested and allocated power values in the 802.3 Power using MDI TLV.	Uknown
Management Problems Corrected in 8.41.01.0004	Introduced in Version:
EDR memory in free list error occurs while setting snmpTargetAddrTDomain to value other than snmpUDPDomain without changing snmpTargetAddrTAddress to match the domain type.	4.11.17
Updated CLI engine to make TAB-key function as '?' whenever command cannot be completed.	1.07.19
Multicast Problems Corrected in 8.41.01.0004	Introduced in Version:
If a module resets, or if a new module is inserted into a chassis, some egress ports on a static MAC multicast are removed.	1.07.19
When a Multicast Router/Querier port on a VLAN times out of the IGMP Multicast Router table, multicast flows on that VLAN may not be delivered correctly due to hardware being mis-programmed.	8.32.01
MVRP Problems Corrected in 8.41.01.0004	Introduced in Version:
VLAN egress registered dynamically by MVRP may bounce when the system is in a steady state.	7.91.01
The CPU utilization may spike up to 99% indefinitely due to MVRP. The system may crash or require manual intervention to force a reset.	8.31.01
MVRP may fail to propagate SPB Base VLANs on ports that are forwarding in the CIST context after disabling SPB on a device.	8.31.01
NetFlow Problems Corrected in 8.41.01.0004	Introduced in Version:
The help string for the "netflow set export-rate" command does not specify valid rate range.	8.01.01
The "clear netflow all" command does not clear non-default netflow export rate settings.	8.01.01
When exporting Netflow V9 records for switched flows that have a tunneled header encapped (for example, GRE), the records are incorrectly exported as routed flows.	8.22.01
For flows that have a tunneled header present, the L2, L3, and L4 fields in NetFlow exported records are not valid.	8.22.01
Non-IP flows that the switch encapsulates with a tunnel header have NetFlow records generated with invalid fields.	8.22.01
PIM-SM IPv4 Problems Corrected in 8.41.01.0004	Introduced in Version:
PIM may drop neighbor adjacencies when running with large number of PIM neighbors.	7.00.01
PIM bootstrap messages are sent out that slightly exceed the MTU, requiring unnecessary IP fragmentation.	7.00.01

PIM-SM IPv4 Problems Corrected in 8.41.01.0004	Introduced in Version:
When running PIM or DVMRP to route multicast traffic, errors similar to the following appear: RtrMc[1.tRMcEvt]Error deleting tmpFlow from TmpDb (2,723,1.1.1.1,225.1.1.1) = notFound[1.tRMcPkt]Hash find - flow vrflds don't match (0,2)	8.31.01
Platform Problems Corrected in 8.41.01.0004	Introduced in Version:
When tunneled bridge ports are active, infrequently, messages similar to the following appear: ===== ===== Message 5/241 Syslog Message 08.30.01.0033 08/02/2014 10:45:31 <3>PiMgr[16.tDispatch]piMgrBindSystemPortAndHwPort(0,0x3000):Port(s) are already bound. pimSystemPortToHwPort[0]=0x8000;pimHwPortToSystemPort[0x3000]=0x580 ===== ===== Message 24/173 Syslog Message 08.30.01.0033 03/11/2014 05:22:38 <3>chassis[1.tBcastStRx]remoteModuleInfoPowerUpdate(6,""):Unsupported board type found.	8.20.02
If a linecard is swapped with a different model of line card in K chassis, and the new card has an error during its power-up, with message such as: <163>Aug 6 11:18:38 0.0.0.0 NIM[11.tNismProc]Unable to power up NIM 3 with status 4010 6088 then the whole K chassis may be reset with this error: <163>Aug 6 11:18:43 0.0.0.0 Dune[11.tNismProc]assert:/fwbld/manhattan/08_30_01_0035/manhattan/dd/dune/src/petraC hip_dd.cxx:1554 12 > 2	7.00.01
Messages similar to the following may appear during line card initialization. These messages are harmless. <163>Aug 6 08:58:13 100.10.10.55 Fujii[11.tFujiRx]fujiRx: Invalid Rx 'inUse' info (Fuji 0, transId 0, seqNum 4, expected seqNum 10) <163>Aug 6 08:58:13 100.10.10.55 Fujii[11.tFujiRx]fujiRxTask: fujiRx() failed for block id	7.31.02
QoS/CoS Problems Corrected in 8.41.01.0004	Introduced in Version:
When switch is in Layer 2 mode, Layer 3 multicast protocols like VRRP and OSPF are not prioritized above user data.	7.00.01
Radius Problems Corrected in 8.41.01.0004	Introduced in Version:
Enabling and/or disabling the RADIUS accounting state over time may result in terminated network sessions (macauthentication, 802.1x, PWA, etc.) to continue to cause the transmission of RADIUS accounting requests after their termination.	4.00.50
Shortest Path Bridging Problems Corrected in 8.41.01.0004	Introduced in Version:
Traffic traversing an SPBV network does not egress out access ports. Filter database entries indicate traffic is not received on the correct internal ports. If the filter database is cleared, traffic correctly egresses out the access ports.	8.31.01
Port may not become internal to the region even though ISIS adjacency is indicated.	8.31.01

Shortest Path Bridging Problems Corrected in 8.41.01.0004	Introduced in Version:
MVRP may propagate SPBV Base-VID registrations on ports within the SPBV domain.	8.31.01
In Shortest Path Bridging, an SNMP query with a context of getNext on the ieee8021SpbTopNodeTable table causes the device to stop responding. The system ID index passed into the getNext query actually exists in the topology, which is the underlying problem. This effect may also occur when issuing SPB commands to show topology information, such as "show spb neighbors."	8.31.01
Configuring multiple bridges with different SPBV SPVID allocation modes can lead to high CPU utilization.	8.32.01
There is no user-evident notification that SPB ports go operationally down when setting those ports' spantree adminPathCost to a value greater than 16777213.	8.32.02
After clearing and re-creating a static multicast MAC, traffic destined through a Shortest Path Bridging network is dropped.	8.32.01
When backuproot is enabled for the CIST on a device that is part of an SPB region and the directly connected root bridge is external to the region, and backuproot is triggered for that device due to failure of the root bridge, the new topology resulting from the change in bridge priority is not affected. This results in a loss of connectivity. Spanning tree modifies the CIST bridge priority, but fails to convey the change to ISIS-SPB, which is responsible for calculating the topology within the SPB region.	8.32.01
SNTP Problems Corrected in 8.41.01.0004	Introduced in Version:
With SNTP unicast client configured, after 497 days, SNTP time requests may stop being sent.	4.05.08
Spanning Tree Problems Corrected in 8.41.01.0004	Introduced in Version:
Spanning tree consumes all packets with the destination address for the IEEE Bridge Group Address/Nearest Customer Bridge group address. This has two effects. First, other applications for which the PDU is intended are not received. Second, a PDU which is not a BPDU is processed by spanning tree and marked as an invalid BPDU.	7.30.01
Spanning tree debug counters are incorrect for RSTP.	8.20.02
Connecting an SPB device in customer bridge mode to a bridge running in provider mode can result in malformed adjacencies with other devices, leading to network instability and spanning tree ports in "listening" state.	8.31.01
On boot up, in a device with multiple connections to root, there may be an initial delay of up to 10 seconds for the root port to reach the forwarding state and pass traffic.	7.63.01
Syslog Problems Corrected in 8.41.01.0004	Introduced in Version:
Logging server list identifiers are translated incorrectly between releases causing logged messages to be directed to the incorrect logging server, console, file, or secure file.	7.40.00
Tunnel Manager Problems Corrected in 8.41.01.0004	Introduced in Version:
Non-IP packets entering an L2 tunnel may not be properly encapsulated.	7.40.01
Traffic may not be flooded to local ports if a tunnel bridge port is configured on the same VLANs egress list and there is no route to the remote IP address of the associated L2 tunnel.	8.31.01

VLAN Problems Corrected in 8.41.01.0004	Introduced in Version:
When a large number (thousands) of dynamic VLANs are deleted together (created by gvrp/mvrp) a core may be taken followed by a reset.	7.91.03
Clearing a VLAN created through "set vlan create" occasionally causes traffic destined to a GVRP- or MVRP-configured port to be lost on the cleared VLAN in a multi-card system.	7.91.01
VRRP Problems Corrected in 8.41.01.0004	Introduced in Version:
The router may stop responding or crash when adding an VRRP IP address to the existing configuration.	8.31.02
It is possible to enter an IPv6 address as a VRRP address when the VRID is a VRRPv3 IPv4 VRID. The address entered becomes a seemingly random IPv4 address in the configuration.	8.01.01
VSB Problems Corrected in 8.41.01.0004	Introduced in Version:
In software VSB chassis that have Shortest Path Bridging or Tunnels active, very infrequently, a module may reset with a message similar to: <0>Fuji[12.tNimIntr]Switch Chip 3 (Slot 12 Mainboard) detected fatal condition. (0x00ebfdb0 0x01827aa8 0x01821394 0x00443130 0x00470ea4 0x00472498 0x01b2de24 0xeeeeeeee) <1>MCNXMGR[7.tFujiAge]ASIC failed to write callback, transaction = 0, chip = 1 (0x00ec0078 0x019e8acc 0x0184b054 0x0184b3f0 0x01b2f3e4 0xeeeeeeee) <0>mazama[12.interrupt]Host buffer manager error. rxInterrupt error:0x8739 if:1 buffNum:1849 bufferUseLog:0x88883456 availCnt:0x35e2 actChain:0xa12fqData:0x7090400 fqHead:0x522 fqTail:0x3dcb btRdata:0 bmgIntf:0x2136b9 getBuf:0x84b4a9ca link0Buf:0xa9c6a9c6 link1Buf:0x84b084b0 walk0Buf:0 walk1Buf:0 freeBuf:0 freeChainPkt 0:0xd743d743 1:0xdee1dee1 2:0x6e406e40 (0x00ec0078 0x01858574 0x01856bbc 0x00470058 0x00000518 0x00000004 0x01b1a9f0 0x007116d0 0x0070096c 0x007009cc 0x03369dd4 0x006f1bb8 0x0070c94c 0x0070c8c4 0x0070d018 0x018611dc 0x018557d4 0x0185eb24 0x01b2f424 0xeeeeeeee) The following non-reset level messages may also be logged: <3>mazama[8.tDispatch]Host buffer already free. freeBuffer:10897 DIf that message is logged, there is vulnerability to the resets mentioned above." "System instability might be experienced with messages similar to ""Interhost Unit 1 no rx space in Net Pool"".D"	8.11.01
"System instability might be experienced with messages similar to ""Interhost Unit 1 no rx space in Net Pool"".D"	6.00.02

Features Enhancements 8.32.02.0008

CLI Enhancements in 8.32.02.0008
A CLI command has been added, "show flowlimit stats". The command shows flow stats per port and is included in D"show support" output.

Problems Corrected in 8.32.02.0008

ACL Problems Corrected in 8.32.02.0008	Introduced in Version:
When ACL logging is enabled on a Policy ACL, the Policy ACL specific field "set-dscp <value>" was not displayed in the log message.	8.32.01
If ACL logging is enabled on a policy ACL, it causes the policy ACL to be persisted as an extended ACL. On reboot, the ACL is restored as an extended ACL and the "set-dscp" action is missing.	8.32.01
IGMP Problems Corrected in 8.32.02.0008	Introduced in Version:
When running in provider bridge mode, IGMP queries will not be transmitted properly.	8.32.01
Multicast Problems Corrected in 8.32.02.0008	Introduced in Version:
After clearing and recreating a static multicast MAC, traffic destined through a Shortest Path Bridging network will be dropped.	8.32.01
PoE Problems Corrected in 8.32.02.0008	Introduced in Version:
CLI commands which access PoE information might hang/timeout during PoE microcode update.	7.00.01
Spanning Tree Problems Corrected in 8.32.02.0008	Introduced in Version:
On boot up, in a device with multiple connections to root, there may be an initial delay of up to 10 seconds for the root port to reach the forwarding state and pass traffic.	7.63.01

Features Enhancements 8.32.01.0021

Slip Horizon L2 Tunneling Support in 8.32.01.0021
Split Horizon L2 Tunneling: L2 Tunnel Enhancement providing for a loop free mesh topology without requiring a loop prevention protocol such as Spanning Tree. With Split Horizon configured on the switch, the switch will not forward packets between tunnel bridge ports if the associated tunnels bound to these tunnel bridge ports belong to the same Split Horizon group.
Layer 3 Policy ACL Support in 8.32.01.0021
Layer 3 Policy ACL: Policy ACLs allow the Administrator to specify an IPv4 packet signature and set the DSCP value for matching packets in order to prioritize relatively short duration connections between specific end points (such as VOIP traffic). Policy ACLs are intended to be used by an application capable of dynamically configuring the ACL to prioritize relatively short duration connections between specific end points. With external integration, entries in the policy ACL will be updated rapidly with an entry created for each new connection (VoIP call) and deleted when the connection terminates. The creation and application of policy ACLs do not persist after a system reset due to the transient nature of the connections to which they are applied.
IP Host Mobility and Fabric Routing Support in 8.32.01.0021
IP Host Mobility/Fabric Routing: The following behavioral enhancements have been made to Fabric Routing and IP Host Mobility features. There are no specific configuration requirements related to these behavioral enhancements. Virtual Subnet Support – This removes the 8.31 requirement for a layer 2 connection between sites (virtual or otherwise) for IP host mobility.

IP Host Mobility and Fabric Routing Support in 8.32.01.0021

Foreign Subnet Support – Support has been added to allow devices with foreign IP addresses (not belonging to the subnet) to utilize Fabric Routing infrastructure as well as become reachable with the IP Host Mobility feature with Proxy-arp enabled.

Problems Corrected in 8.32.01.0021

Distributed Services Problems Corrected in 8.32.01.0021	Introduced in Version:
Module might reset with messages similar to: "Chassis coherency timeout exceeded".	7.62.07
Denial of service (DOS) attack results in warning messages "this server has been invalidated" printed to the console.	1.07.19
Ethernet OAM Services Problems Corrected in 8.32.01.0021	Introduced in Version:
A CPU under heavy load may prevent transmission of OAMPDUs which can lead to a discovery timeout on an OAM peer.	8.31.01
Hardware Problems Corrected in 8.32.01.0021	Introduced in Version:
<p>Infrequently messages similar to: <163>Jan 29 13:06:59 100.10.10.22 Dune[1.dTcmTask]Petra[0] Received Interrupt PB_IHB_INVALID_DESTINATION_VALID instance 0, count 3, value= 0x13deb or <3>Dune[1.dTcmTask]Petra[0] Received Interrupt PB_IHB_INVALID_DESTINATION_VALID instance 0, count 2159, value= 0x1 or <165>Jun 5 11:32:19 100.10.10.22 Dune[16.tDuneErrM]Petra[0] Interrupt PB_IHB_INVALID_DESTINATION_VALID instance 0 still active <165>Jun 5 11:32:29 100.10.10.22 Dune[11.tDuneErrM]Petra[0] Interrupt PB_IHB_INVALID_DESTINATION_VALID instance 0 is off</p> <p>maybe displayed. Whenever one of these messages is displayed, a packet that should have been forwarded will be dropped.</p>	8.11.01
<p>If a linecard is swapped with a different model of line card in K chassis, and the new card has an error during its power-up with a message such as: <163>Aug 6 11:18:38 0.0.0.0 NIM[11.tNismProc]Unable to power up NIM 3 with status 4010 6088 then the whole K chassis may be reset with this error: <163>Aug 6 11:18:43 0.0.0.0 Dune[11.tNismProc]assert:/fwbld/manhattan/08_30_01_0035/manhattan/dd/dune/src/petraC hip_dd.cxx:1554 12 > 2</p>	7.00.01
<p>Messages similar to the following may be experienced during line card initialization. These messages are harmless. <163>Aug 6 08:58:13 100.10.10.55 Fujii[11.tFujiRx]fujiRx: Invalid Rx 'inUse' info (Fuji 0, transId 0, seqNum 4, expected seqNum 10). <163>Aug 6 08:58:13 100.10.10.55 Fujii[11.tFujiRx]fujiRxTask: fujiRx() failed for block id 0.</p>	7.31.02
A "Microcode load from SPI-EEPROM failed on PHY" error should cause system initialization to fail, but does not. Consequently, one or more 10G front-panel ports will likely fail to link up.	7.72.01

Host Mobility Problems Corrected in 8.32.01.0021	Introduced in Version:
The host-mobility aging feature has been removed in order to support host-mobility in a segmented network design whereby the router's VRRP interfaces are not connected.	8.11.01
Host-mobility is now supported in a segmented VRRP network design where the router's VRRP interfaces are not connected.	8.31.01

Host Services Problems Corrected in 8.32.01.0021	Introduced in Version:
When tunneled bridge ports are active, infrequently messages similar to: <pre> ===== Message 5/241 Syslog Message 08.30.01.0033 08/02/2014 10:45:31 <3>PiMgr[16.tDispatch]piMgrBindSystemPortAndHwPort(0,0x3000):Port(s) are already bound. pimSystemPortToHwPort[0]=0x8000;pimHwPortToSystemPort[0x 3000]=0x580 and/or ===== Message 24/173 Syslog Message 08.30.01.0033 03/11/2014 05:22:38 <3>chassis[1.tBcastStRx]remoteModuleInfoPowerUpdate(6,""):Unsupported board type found. may be logged. </pre>	8.20.02

IGMP Problems Corrected in 8.32.01.0021	Introduced in Version:
If a configuration is enabled for IGMP on a VLAN that becomes an SPVID, you cannot delete the config.	8.31.01
When using a BaseVid without spvid allocation due to an insufficient spvid pool or a lack of boundary egress, IGMP may not forward traffic.	8.31.01

IPStack Problems Corrected in 8.32.01.0021	Introduced in Version:
Within a network environment where DHCP clients are active, over time, could see an exhaustion of resources that prevent IP host communication and loss of device management.	7.91.01

Management Problems Corrected in 8.32.01.0021	Introduced in Version:
The K-Series chassis doesn't generate boardInsertion, boardRemoval, boardOperational, and boardNonOperational traps.	7.30.01

Management Problems Corrected in 8.32.01.0021	Introduced in Version:
<p>When syslog servers are configured, if any of the following cli commands are issued:</p> <ul style="list-style-type: none"> • show support • show config • show config logging <p>The switch will lose (leak) 144 bytes of memory. If commands are issued frequently enough the switch will reset, logging a message similar to:</p> <p>Message 3/30 EDR Record 07.62.05.0001H 07/27/2014 19:55:11 Severity/Facility: FATAL/KERNEL Task: tCLIO Injection Point: memPartLib.c:2498 Address: 0x00000000 memPartAlloc: block too big 84624 bytes (0x10 aligned) in partition 0x2234548</p>	7.11.01
Multicast Problems Corrected in 8.32.01.0021	Introduced in Version:
<p>When Shortest Path Bridging is enabled, upon clearing of a static Multicast Mac address, messages similar to following will be logged:</p> <pre><163>Oct 22 14:59:43 1.1.191.11 Fujij[7.tFujiAge]fujiAgeTransId: Fuji=0, adr=0x52580000, transId=4, seqNum=32966 write access timeout <162>Oct 22 14:59:43 1.1.191.11 RfrmrHw[7.tFujiAge]Error: FujiWriteFailCallBack() rib size:64 transid:4 blockId:0 index:0 vx ti tFujiRx</pre> <p>In addition, the chassis may also reset and log a message similar to:</p> <pre>Message 21/269 Syslog Message 08.31.02.0006 10/23/2014 10:09:07 <0>FtmLi[7.tHBChk]heartbeat task delayed for 41 seconds (0x00e18710 0x00e372dc 0x00e3af6c 0x00e3af94 0x016586e4 0xeeeeeeeee)</pre>	8.31.01
<p>Multicast cache entries show up in the router even without a multicast routing protocol enabled on an interface.</p>	8.31.01
<p>Multicast frames that are buffered and forwarded do not have TTL decremented.</p>	8.31.01
<p>IP Multicast is not forwarded correctly to local or remote ports after a port goes down that has a Router or Querier attached.</p>	8.31.01
MVRP Problems Corrected in 8.32.01.0021	Introduced in Version:
<p>VLANs that are either forbidden or mapped to the SPBV MST at bootup will not allow dynamic registration via MVRP or GVRP after the VLAN forbidden egress status or MST mapping is cleared.</p>	8.31.01
<p>The CPU utilization may spike up to 99% indefinitely due to MVRP. The system may crash or require manual intervention to force a reset.</p>	8.31.01

OSPF Problems Corrected in 8.32.01.0021	Introduced in Version:
If an OSPFv2 virtual link is configured with an invalid timer value of 0, the router will crash with the following syslog message, "sms_get timeout: oid=3e000001, tRtrPtcls state: running, last wakeup: 1 tics, IPS in use cnt: 1968, Bytes: 6527728"	7.00.01
PIM-DM Problems Corrected in 8.32.01.0021	Introduced in Version:
"show ip mcache" shows a corrupted/incorrect Source/Destination IP in the display output.	8.31.01
Platform Problems Corrected in 8.32.01.0021	Introduced in Version:
SFP sensors information may not be present for some option module ports when multiple option modules are installed. "Show port transceiver" CLI output may incorrectly indicate that sensor data is not available on these ports.	8.31.01
Unable to delete a file/image from the users directory if it has the same name as the current running image. You will get the following error return. (su)->delete slot1/myImage The active image cannot be removed. Failed to remove /slot1/myImage	7.30.01
Message similar to the following might be seen when bonding is disabled:<163>Feb 12 10:42:00 100.10.10.22 dot3Mgt[4.tEmanate10]dot3MgtDist::ifJackEntryGet():sendMessage(ackReq)!=kDs_good;sendMask=0x10000	7.70.00
RMON Problems Corrected in 8.32.01.0021	Introduced in Version:
"show rmon stats" report might fail to include a bond port. This problem is intermittent (all of the bond ports might show up on some reboots), and the omitted bond port could change from reboot to reboot.	7.91.01
Shortest Path Bridging Problems Corrected in 8.32.01.0021	Introduced in Version:
SPB devices may not agree topology agreement digest after changing master role.	8.31.01
Traffic may not recover after disable/re-enable SPB.	8.31.01
In a Shortest Path Bridging domain, when a device becomes the new regional root, designated ports on this new regional root go into listening state. Consequently, CIST traffic using this path is blocked. The issue is resolved by forcing a BPDU to be sent by the root port on the peer device.	8.31.01
In a Shortest Path Bridging-VLAN domain, when a device becomes the new regional root, customer traffic that ingresses the network on a base VID does not reach the intended destination endpoint(s). The associated SPVID lacks egress on some bridges throughout the SPBV network, and there is no clear indication of why this is so. The issue is resolved by forcing a BPDU to be sent by the root port on the peer device.	8.31.01
Occasionally on bootup, static layer 2 multicast traffic that runs through shortest path bridging will not recover.	8.31.01
In a Shortest Path Bridging VLAN (SPBV) domain, ports are incorrectly set to backup role and a state of blocking. The only ports affected are internal to the region and the consequence is limited network connectivity. Toggling the SPB configuration on the port may fix the problem, but not always.	8.31.01

Shortest Path Bridging Problems Corrected in 8.32.01.0021	Introduced in Version:
For Software Bonded flows, from SPB ports, the first 4 bytes of the Software Bond Header is not getting removed properly, causing loss of L2 multicast traffic.	8.31.01
The agreement protocol for Spanning Tree internal to the SPB region requires an exchange of BPDUs greater in number than what is required for rapid failover in RSTP or MSTP. Spanning Tree software rate limiters may cause a BPDU drop during this exchange causing the protocol to be interrupted for a HELLO period, two seconds by default, until the next periodic transmit of a BPDU. This will delay convergence when SPB has the digest convention configured for loopFreeBoth.	8.31.01
System crashes when rebooting one blade in a multi-blade system with message similar to: "<161>Oct 30 08:40:27 0.0.0.0 System[7]Chassis coherency timeout exceeded, resetting. delta:222000 curr:335186 nts:113186 nto:30000 hw:0x37000000 lnk:0x37000000 nv:0x37000000 img:0x37000000 max:0x37000000 (0x00e8535c 0x0071b18c 0x01ad4564 0xe8535353)".	8.31.01
Port state may be listening for SPB internal port due to neighbor transmitting BPDUs with the agreeDigestValid flag persistently false.	8.31.01
Traffic traversing a SPBV network does not egress out access ports. Filter database entries indicate traffic is not received on the correct internal ports. If the filter database is cleared, traffic correctly egresses out the access ports.	8.31.01
Occasionally when a port's operational state is changed, layer 2 static multicast traffic over Shortest Path Bridging is lost on that port.	8.31.01
Ports may become blocked when adding a BVLAN or SPVID and then immediately removing it. Spanning tree reinitializes the port topology information calculated by ISIS-SPB, but the information is not refreshed because the topology calculated by ISIS-SPB has not actually changed.	8.31.01
When Shortest Path Bridging is globally disabled, Layer 2 multicast traffic will not be forwarded across a Virtual Switch Bond when using a configured Shortest Path Bridging BaseVLAN.	8.31.02
When an SPB regional port becomes a boundary port and then reenters the region, ISIS-SPB and Spanning Tree may become out of sync with respect to the value the port is using for agreement digest. The value transmitted in an SPT BPDU may differ from the value transmitted in the SPB-Digest sub-TLV of the SPB Hello PDU. This may result in traffic loss due to agreement not being reached between the connected ports.	8.31.03
CIST root port may become stuck in the listening state when disabling and reenabling the global SPB status for all the nodes in an SPB region.	8.31.03
If Shortest Path Bridging is enabled, or enabled then disabled, a "show mac addr ..." command could take minutes or tens of minutes to complete. All matching Filter Database entries should still be returned.	8.31.01
SPB configurations using manual SPVID allocation mode without manually configured SPVIDs can lead to high CPU utilization and network instability.	8.31.01

Spanning Tree Problems Corrected in 8.32.01.0021	Introduced in Version:
A root or alternate port may get stuck in a state where it will not respond to a proposal BPDU with an agreement BPDU. This will cause port forwarding for the connected designated port to use timers rather than the rapid forwarding mechanism. Additionally, if the designated port is configured for lp (Loop Protect), it will detect a loop protect event and remain in the listening state.	7.60.01

Spanning Tree Problems Corrected in 8.32.01.0021	Introduced in Version:
The Multisource function detects multiple BPDU sources received on a point-to-point link and sets the point-to-point operational status to false. The point-to-point operational status is an input into the rapid transition to forwarding capability for rapid spanning tree. It is also a factor in the Loop Protection mechanism and in Shortest Path Bridging. A port that receives BPDUs from multiple sources, where those sources are exclusively different ports on the same transmitting bridge, will not be triggered for multisource and will remain operationally point-to-point.	8.31.01
FDB entry not removed for IST port in an SPB region during a topology change. This can cause traffic assigned to VLANS mapped to SID 0 to be directed out the wrong port until the FDB entry times out.	8.31.01
A port on the root bridge may select a backup role instead of a designated role if it receives a BPDU from another bridge where the role in the flags field indicates a designated role, the root identifier is the ID of the receiving bridge and the transmitting port ID is lower than the receiving port ID.	7.00.01
A temporary loop may be created when the root bridge relinquishes its root status and the direction of root in the network reverses, i.e. designated ports become root/alternate ports and root/alternate ports become designated.	7.00.01

Tunnel Manager Problems Corrected in 8.32.01.0021	Introduced in Version:
LLC packets might not be received at destination when sent across an L2 tunnel due to IPX being filtered inside tunnel flows.	8.11.01
Downgrading from a future version that requiring only a source tunnel endpoint to be configured on a tunnel, causes the tunnel to be deleted.	7.41.02

VLAN Problems Corrected in 8.32.01.0021	Introduced in Version:
VLAN egress registered dynamically by MVRP may bounce when the system is in a steady state.	7.91.01

Problems Corrected in 8.31.03.0001

Shortest Path Bridging Corrected in 8.31.03.0001	Introduced in Version:
In a large Shortest Path Bridging network, running the command "show spb path" will cause the Shortest Path network traffic to stop forwarding.	8.31.02
SPB Port configuration will be lost if hello parameters are configured and lower port numbers do not have hello parameters configured.	8.31.02

Spanning Tree Corrected in 8.31.03.0001	Introduced in Version:
A temporary loop may be created when the root bridge relinquishes its root status and the direction of root in the network reverses, i.e. designated ports become root/alternate ports and root/alternate ports become designated.	7.00.01

Problems Corrected in 8.31.02.0014

802.1d Filter Database Problems Corrected in 8.31.02.0014	Introduced in Version:
If Tunneled Bridge Ports or Network Address Translation are active, traffic may not be forwarded correctly.	8.31.01

BFD Problems Corrected in 8.31.02.0014	Introduced in Version:
Not all of the BFD sessions recover from a master blade failure. In the event there are unrecoverable BFD sessions, the user should completely remove the BFD configuration on both sides and then reconfigure.	8.31.01
Module might reset with message indicating DSI Exception in Thread Name: tTrackBfdS.	8.31.01
CoS Problems Corrected in 8.31.02.0014	Introduced in Version:
When the switch is in layer 2 mode, layer 3 multicast protocols (such as VRRP and OSPF), when switched, are not prioritized above user data.	7.00.01
Distributed Services Problems Corrected in 8.31.02.0014	Introduced in Version:
Module might reset with messages similar to: "DSI exception" and "Thread Name: tDSrecv4".	7.00.01
Ethernet OAM Problems Corrected in 8.31.02.0014	Introduced in Version:
A CPU under heavy load may prevent transmission of OAMPDUs which can lead to a discovery timeout on an OAM peer.	8.31.01
Hardware Problems Corrected in 8.31.02.0014	Introduced in Version:
A "Microcode load from SPI-EEPROM failed on PHY" error should cause system initialization to fail, but does not. Consequently, one or more 10G front-panel ports will likely fail to link up.	7.72.01
A blade may reset with a messages similar to: <1>MCNXMGR[12.tFujiAge]ASIC failed to write callback, transaction = 0, c hip = 2 or <3>Fuji[1.tNimIntr]Fuji LU RAM 2 MAIN intr: Fuji=0, Adr=0, Reg=0x00000100 or <3>mazama[8.tDispatch]Host buffer already free. freeBuffer:8715 These resets will occur more frequently if the switch has Tunnels or Shortest Path Bridging enabled.	8.21.01
IGMP Problems Corrected in 8.31.02.0014	Introduced in Version:
It is possible for IGMP to have a non general query refresh the other querier present timers. Causes no functional issues.	7.30.01
User is unable to disable or delete an IGMP configuration for a VLAN if the Vid becomes configured as a Spvid.	8.31.01
CLI Syslog may indicate that a failed IGMP configuration succeeded, when it did not.	7.00.01
If adding an SPB base Vid, before enabling IGMP, IGMP may not recognize the base Vid, resulting in traffic issues.	8.31.01
A user is able to enable IGMP query on an SPBV Spvid.	8.31.01
IP Multicast is not forwarded correctly to local or remote ports after a port goes down that has a Router or Querier attached.	8.31.01

Management Problems Corrected in 8.31.02.0014	Introduced in Version:
Loading configuration from a file on a Bonded K-Series may fail. A workaround is to power down or remove installed linecards before running configure command.	8.21.01
Multicast Problems Corrected in 8.31.02.0014	Introduced in Version:
mgmdStdMib InverseRouterCacheTable may not SNP walk properly.	7.60.01
Static layer 2 multicast traffic is not forwarded through a hardware VSB device that is using Shortest Path Bridging.	8.31.01
When Shortest Path Bridging is enabled, upon clearing of a static Multicast Mac address, messages similar to following will be logged: <163>Oct 22 14:59:43 1.1.191.11 Fujij[7.tFujiAge]fujiAgeTransId: Fuji=0, adr=0x52580000, transId=4, seqNum=32966 write access timeout <162>Oct 22 14:59:43 1.1.191.11 RfrmrHw[7.tFujiAge]Error: FujiWriteFailCallBack() rib size:64 transid:4 blockId:0 index:0 vx ti tFujiRx In addition, the chassis may also reset and log a message similar to: Message 21/269 Syslog Message 08.31.02.0006 10/23/2014 10:09:07 <0>FtmLi[7.tHBChk]heartbeat task delayed for 41 seconds (0x00e18710 0x00e372dc 0x00e3af6c 0x00e3af94 0x016586e4 0xeeeeeeee)	8.31.01
MVRP Problems Corrected in 8.31.02.0014	Introduced in Version:
VLANs that are either forbidden or mapped to the SPBV MST at bootup will not allow dynamic registration via MVRP or GVRP after the VLAN forbidden egress status or MST mapping is cleared.	8.31.01
NETFLOW Problems Corrected in 8.31.02.0014	Introduced in Version:
When clearing config (clear linecard X), for a linecard that is not inserted and active, netflow port enabled settings will not be cleared. If same, or new linecard is re-inserted, netflow will still be enabled on those ports.	7.03.01
OSPF Problems Corrected in 8.31.02.0014	Introduced in Version:
If an OSPFv2 virtual link is configured with an invalid timer value of 0, the router will crash with the following syslog message: "sms_get timeout: oid=3e000001, tRtrPtcls state: running, last wakeup: 1 tics, IPS in use cnt: 1968, Bytes: 6527728"	7.00.01
PIM-DM Problems Corrected in 8.31.02.0014	Introduced in Version:
Multicast frames that are buffered and forwarded do not have TTL decremented.	8.31.01
PIM-SM Problems Corrected in 8.31.02.0014	Introduced in Version:
IP Multicast flows may revert to a "register state" after PIM events such as neighbor loss, RP loss, etc.	8.31.01
Multicast cache entries show up in the router even without a multicast routing protocol enabled on an interface.	8.31.01

Platform Problems Corrected in 8.31.02.0014	Introduced in Version:
System instability might be experienced with messages similar to "Interhost Unit 1 no rx space in Net Pool".	6.00.02
When writing to a file on a remote blade, if the connection becomes unresponsive, the local blade could reset. An example would be running the following command from the master slot to a slot across a bond link: "show config all outfile slot13/showCfgAll.out" The log should have something similar to the following: Message 83/263 Exception PPC750 Info 08.30.01.0036 08/13/2014 08:54:27 Exc Vector: DSI exception (0x00000300) Thread Name: tCLI0"	7.00.01
Underlying transport errors will cause the messages "TIPC discarding incoming Ethernet message with destination <mac_address>" to be displayed resulting in internal network buffer loss and a segmentation of a slot in a chassis to stand alone mode.	8.31.01
When displaying debug CLI base information for some copper SFP cable assemblies, the output may incorrectly display the interface type as "40G Act Cbl" instead of "1000BASE-CX".	8.22.02
A performance reduction causes the throughput of new traffic processing to be reduced with default configuration.	8.31.01
PoE Problems Corrected in 8.31.02.0014	Introduced in Version:
PoE might occasionally stop delivering power to PDs.	7.00.01
RMON Problems Corrected in 8.31.02.0014	Introduced in Version:
"show rmon stats" report might fail to include a bond port. This problem is intermittent (all of the bond ports might show up on some reboots), and the omitted bond port could change from reboot to reboot.	7.91.01
Shortest Path Bridging Problems Corrected in 8.31.02.0014	Introduced in Version:
SPB convergence times may take longer than expected when region topology changes.	8.31.01
In a multi-slot or bonded chassis, LAG port egress may not be set properly for an SPVID on a non-switch master blade. There is a small timing window where the distributed spanning tree port state information is missed.	8.31.01
Insertion or removal of a module in a bonded system can cause poor network convergence times as well as a temporary loss of traffic.	8.31.01
SPB devices may not agree with topology agreement digest after changing master role.	8.31.01
Occasionally when a chassis blade is removed, Shortest Path Bridging traffic is temporarily lost even when no shortest paths pass through the blade.	8.31.01
When running spanning tree in SPB mode, traffic is lost when connected ports have differing configuration for SPB port status. One side sees the port as internal to the region while the other sees it as external. This results in a disputed BPDU status causing the port to remain in the listening state.	8.31.01
Changes in the topology of an SPB region result in convergence times above expectations. This is due to the number of BPDU transmit requests exceeding the txHoldCount value,	8.31.01

Shortest Path Bridging Problems Corrected in 8.31.02.0014	Introduced in Version:
including when that value is set to the maximum of 10. TxHoldCount is the number of BPDUs which may be sent immediately after which the transmit rate becomes one BPDU per second for the given port.	
Traffic may not recover after disable/re-enable SPB.	8.31.01
A new root port for an SPT may forward before the old root port on a remote blade disables forwarding, opening a transient loop.	8.31.01
When there is a change in the topology of the SPB region, ports might get stuck in the listening state.	8.31.01
Port may not become internal to the region even though ISIS adjacency is indicated.	8.31.01
In a Shortest Path Bridging-VLAN domain, when a device becomes the new regional root, customer traffic that ingresses the network on a base VID does not reach the intended destination endpoint(s). The associated SPVID lacks egress on some bridges throughout the SPBV network, and there is no clear indication of why this is so. The issue is resolved by forcing a BPDU to be sent by the root port on the peer device.	8.31.01
Occasionally on bootup static layer 2 multicast traffic that runs through shortest path bridging will not recover.	8.31.01
In a Shortest Path Bridging VLAN (SPBV) domain, ports are incorrectly set to backup role and a state of blocking. The only ports affected are internal to the region and the consequence is limited network connectivity. Toggling the SPB configuration on the port may fix the problem, but not always.	8.31.01
For Software Bonded flows, from SPB ports, the first 4 bytes of the Software Bond Header is not getting removed properly, causing loss of L2 multicast traffic.	8.31.01
MVRP may propagate SPBV Base-VID registrations on ports within the SPBV domain.	8.31.01
The agreement protocol for Spanning Tree internal to the SPB region requires an exchange of BPDUs greater in number than what is required for rapid failover in RSTP or MSTP. Spanning Tree software rate limiters may cause a BPDU drop during this exchange causing the protocol to be interrupted for a HELLO period, two seconds by default, until the next periodic transmit of a BPDU. This will delay convergence when SPB has the digest convention configured for loopFreeBoth.	8.31.01
System crashes when reboot one blade in a multi-blade system with message similar to: "<161>Oct 30 08:40:27 0.0.0.0 System[7]Chassis coherency timeout exceeded, resetting. delta:222000 curr:335186 nts:113186 nto:30000 hw:0x37000000 lnk:0x37000000 nv:0x37000000 img:0x37000000 max:0x37000000 (0x00e8535c 0x0071b18c 0x01ad4564 0xe8535353)".	8.31.01
Port state may be listening for SPB internal port due to neighbor transmitting BPDUs with the agreeDigestValid flag persistently false.	8.31.01

Spanning Tree Problems Corrected in 8.31.02.0014	Introduced in Version:
A root or alternate port may get stuck in a state where it will not respond to a proposal BPDU with an agreement BPDU. This will cause port forwarding for the connected designated port to use timers rather than the rapid forwarding mechanism. Additionally, if the designated port is configured for lp (Loop Protect), it will detect a loop protect event and remain in the listening state.	7.60.01
The Multisource function detects multiple BPDU sources received on a point-to-point link and sets the point-to-point operational status to false. The point-to-point operational status is an	8.31.01

Spanning Tree Problems Corrected in 8.31.02.0014	Introduced in Version:
input into the rapid transition to forwarding capability for rapid spanning tree. It is also a factor in the Loop Protection mechanism and in Shortest Path Bridging. A port that receives BPDUs from multiple sources where those sources are exclusively different ports on the same transmitting bridge will not be triggered for multisource and will remain operationally point-to-point.	
FDB entry not removed for IST port in an SPB region during a topology change. This can cause traffic assigned to VLANs mapped to SID 0 to be directed out the wrong port until the FDB entry times out.	8.31.01
A port on the root bridge may select a backup role instead of a designated role if it receives a BPDU from another bridge where the role in the flags field indicates a designated role, the root identifier is the id of the receiving bridge and the transmitting port id is lower than the receiving port id.	7.00.01

VSB Problems Corrected in 8.31.02.0014	Introduced in Version:
In a software VSB chassis, Precision Time Protocol frames (PTP), will not be forwarded to any ports that are on the chassis remote from one that packet was received on.	8.11.05

Features Enhancements 8.31.01.0006

Enhancements in 8.31.01.0006
Shortest Path Bridging - IEEE 802.1aq Shortest Path Bridging (SPB) is a protocol that provides data traffic a shortest cost path between any pair of switches in the SPB network. SPB features dynamic route calculation in a loop-free Layer-2 network and fast convergence time using IS-IS. The K-Series supports Shortest Path Bridging VLAN (SPBV).
Routing as a Service (RaaS) - Routing as a Service (RaaS), also known as Virtual Fabric Routing, is an integrated routing service providing a simple, scalable and efficient virtualized routing over any L2 network infrastructure that eliminates all routing protocols within the SPB domain of the network. The RaaS Fabric can scale from a single or pair of chassis to a collection of devices where all of the devices in the SPB domain of the network work as a single and collective layer 3 forwarding mechanism.
SFP/SFP+ Extended Information - Diagnostic information for supported transceivers is provided. In addition to serial number and model details, digital diagnostic information is displayed such as Temperature, Voltage, Transmit Current, Receive Power, Alarm State as well as High/Low thresholds.
Bi-Directional Forwarding Detection (BFD) Enhancements: <i>Shared Fate</i> – With Shared Fate, all routing protocols can be notified within a single BFD session, previous releases supported OSPF protocol only. <i>Graceful Re-start</i> – Support has been added to simultaneously use the routing protocol Graceful Re-start with BFD, in previous releases these features were mutually exclusive. <i>Local and Remote Echo</i> – Echo functionality allows the BFD feature set to test a neighboring routers forwarding plane.
Multicast Buffering Enhancement - Enhancement to support buffering of the initial packets of an IP Multicast flow that arrives prior to the Multicast Routing Protocol determining the proper route. Previous releases would drop the initial IP Multicast packets to be routed prior to the Multicast Routing Protocol determining the route.
ISIS Graceful Re-Start - Graceful Re-Start for the IS-IS protocol has been added. Graceful Re-Start provides for an IS-IS router to continue to forward existing traffic and remain on the forwarding path during a restart of the IS-IS software process.
IP Service Level Agreements Enhancements: This release adds two new types of UDP timing probes to the (IPSLA) feature suite.

Enhancements in 8.31.01.0006

UDP Timing Probe - Uses a variation of the UDP echo paradigm to contact a destination device to determine the round-trip-delay as well as packet delay variation (jitter). Packet delay variation requires both endpoints support the IPSLA feature and have their clocks synchronized.

DNS Timing Probe - Uses the DNS protocol to transmit a DNS query a destination device to determine the round-trip-delay of the DNS answer.

UDP/DNS State Probe - Provides the ability to verify data in a DNS resource record in the answer section of the DNS response packet.

K-Series ARP Capacity Increase – The K-Series ARP table capacity has been increased to 16k entries from 4k entries.

Problems Corrected in 8.31.01.0006

802.1d Filter Database Problems Corrected in 8.31.01.0006	Introduced in Version:
MAC addresses that should age out from filter database will fail to do so. The frequency of this will increase with lower mac age times.	1.07.19
802.1q Relay Problems Corrected in 8.31.01.0006	Introduced in Version:
"set port priority-queue x.x.x 0" may be added to configuration after replacing linecard with different type. This will not happen if the replaced linecard had default port priority-queue settings."	7.00.01
802.1x Problems Corrected in 8.31.01.0006	Introduced in Version:
When using EAP Authentication Methods that require passing certificates, if the packets for those certificates are greater than 1760 bytes, a portion of those packets may be transmitted with invalid data.	8.20.02
ACL Problems Corrected in 8.31.01.0006	Introduced in Version:
When a packet with a protocol other than IPv4 or IPv6 matches an L2 ACL, the L2 source and destination addresses will be displayed in place of the IPv4 and IPv6 addresses and the ethertype will be displayed as a hex value.	8.11.01
When an L2 ACL is applied to an interface, removed from an interface, or when an L2 ACL currently in use is modified, connections may not be removed. This can cause traffic to flow as it did before the change was made. Toggling the interface down then up will clear all connections and allow the L2 ACL to be correctly applied to traffic.	8.11.01
ARP Problems Corrected in 8.31.01.0006	Introduced in Version:
If the ARP table contains an entry for 0.0.0.0 or 255.255.255.255 then an SNMP MIB walk will result in a loop.	7.00.01
The router configured on a service provider switch may respond to ARPs received on a customer VLAN when the VLAN ID matches a router's interface VLAN ID. Conversely, the router configured on a customer switch may respond to ARPs received on a service provider VLAN when the VLAN ID matches a router's interface VLAN ID.	7.91.01
Using the command "clear arp <ipAddress>" may not function properly when clearing an ARP or ND entry in the stale state. If the host is still up a new ARP or ND entry will be added immediately after it is deleted.	7.00.01

CFM Problems Corrected in 8.31.01.0006	Introduced in Version:
CFM PDUs that contain the SenderID TLV will be improperly discarded as invalid frames.	8.21.01
Remote MEP states may be incorrect on CFM MEPs that are configured on LAG ports that span more than one module.	8.02.01
Remote MEP states may be incorrect on CFM MEPs that have no VLAN configuration ("Port MEPs").	8.21.01
The "-verbose" modifier for the "show cfm ... linktrace" CLI command has been removed.	7.91.01
Sending CFM Linktrace messages from MEPs residing on a bridge running in customer mode will not interoperate with CFM MPs residing on bridges running in provider mode.	8.01.01
GVRP Problems Corrected in 8.31.01.0006	Introduced in Version:
Module reset with a message similar to: "<1>DistServ[2.tDsBrdOk]serverWatchDog.3, client 96(GVRP) in rcv for 691" might occur in a system performing traffic flooding together with numerous dynamic VLAN registrations.	7.70.00
HostDos Problems Corrected in 8.31.01.0006	Introduced in Version:
Enabling the HostDoS portScan feature mistakenly filters inbound packets on port 22 when SSH is enabled. HostDoS should only filter these packets when SSH is disabled. This may render the switches SSH server inoperable, and the DoS attack detection logic may produce false positives. A workaround is to not enable HostDos portScan, or to enable it but with a relatively high portScan rate limit. Another workaround is to disable and then re-enable SSH (via a Telnet or console connection). However, the problem will return following a system reboot.	7.30.01
IGMP Problems Corrected in 8.31.01.0006	Introduced in Version:
IGMP may lose track of where a flow entered the system. It may cause flow Interruption due to bad internal hardware programming.	7.79.00
It is possible for IGMP to lose track of which port a flow comes in, and cause an IGMP verify failed, status:0x00020000 message.	7.79.00
When the command "set igmp flow-wait" has both oper-state and time set on the same line, only the oper-state is set.	8.11.01
IP Interface Manager Problems Corrected in 8.31.01.0006	Introduced in Version:
When removing a Layer-3 interface using the "no <interfaceName>" command you may receive a difficult to decipher error message if the interface does not exist.	7.41.02
IPSLA Problems Corrected in 8.31.01.0006	Introduced in Version:
The SLA scheduler sub-mode command 'reset' cannot be entered while the SLA entry is scheduled. In order to reset the attributes for the entry, the user must stop the SLA entry via the 'stop' command in the SLA scheduler sub-mode.	8.01.01
The user will see the following CLI error when attempting to configure an SLA entry that had been previously configured in another VRF: ' Error: Command failed - create IpSla Entry ' The user will either have to remove the SLA entry from VRF in which it is configured, or choose a different SLA entry to configure.	8.11.01

IP Stack Problems Corrected in 8.31.01.0006	Introduced in Version:
ICMP echo requests to IP interface addresses exceeding 100 per second will not all be answered.	8.20.02
IPv4 Forwarding Problems Corrected in 8.31.01.0006	Introduced in Version:
Host routes advertised from the host-mobility routers are installed in other host-mobility peers that direct frames to the core instead of the directly connected networks.	8.21.01
IPv6 Neighbor Discovery Problems Corrected in 8.31.01.0006	Introduced in Version:
ARP/ND entries may expire early if the host does not respond to periodic ARP/ND refresh attempts.	8.21.01
It is possible to configure a Static ND entry which uses the same IP address as an interface address or VRRP address if the static ND entry is created before the other address.	7.00.01
The configuration commands "arp" and "ipv6 neighbor" allow invalid VLAN interfaces such as vlan.0.4095.	7.00.01
Jumbo Problems Corrected in 8.31.01.0006	Introduced in Version:
Port Jumbo MTU settings allowed for values below 1519.	8.01.01
Management Problems Corrected in 8.31.01.0006	Introduced in Version:
Message "masterTrapSem time out, dropping trap" may appear in message log indicating an SNMP trap being dropped.	7.62.06
Entity mib "modelName" data corruption.	7.00.01
Multicast Problems Corrected in 8.31.01.0006	Introduced in Version:
It is possible for modules to reset with the following message: "Machine Check exception Thread Name: tlgmplnp", at boot time, and may also get stuck in a constant reboot loop.	8.11.01
MVRP Problems Corrected in 8.31.01.0006	Introduced in Version:
Dynamic VLANs that were registered by MVRP may still show up in "show vlan" when there are no longer any egress ports. This can happen if the egress was registered on a module port that has since joined a LAG.	7.91.01
The "show vlan" command may show that egress on a port unexpectedly continues to be seen on a VLAN that once was dynamically registered by MVRP if the VLAN is configured statically on that port and then subsequently removed.	7.91.01
Neighbor Discovery Problems Corrected in 8.31.01.0006	Introduced in Version:
The "age" column for the command "show ipv6 neighbors" displays the last time the ND entry was updated instead of the entry's age.	7.40.01

NETFLOW Problems Corrected in 8.31.01.0006	Introduced in Version:
When clearing config (clear linecard X), for a linecard that is not inserted and active, Netflow port enabled settings will not be cleared. If same, or new linecard is re-inserted, Netflow will still be enabled on those ports.	7.03.01
When running in Netflow Version 5 mode, records exported for routed flows may not have a valid Next Hop Router field.	7.20.01
Node Alias Problems Corrected in 8.31.01.0006	Introduced in Version:
If nodealias is disabled on a given port and the maxentries value is set to default, after upgrading to firmware version 8.11.01 or newer will cause the maxentries value to be set to the previous default value.	8.11.01
OSPF Problems Corrected in 8.31.01.0006	Introduced in Version:
If a config file saved prior to version 7.60 contains an OSPF passive interface, it will cause the box to hang if a 'configure' is executed on an upgrade. The config file can be edited to format vlan.0.# instead of vlan # to allow upgrade.	8.22.02
The "debug ip ospf packet" display for virtual interfaces reads "Interface not found for ifIndex 0".	8.11.01
When changing an OSPF network's area ID then failing over, the original area ID is running seen in "show ip ospf interface", though the config reflects the new area ID.	7.00.01
With the removal of passive-interface default, the no passive-interface commands are removed, but they return on reboot of the router. They have no adverse effect.	8.11.01
If OSPF is configured to use a non-existent track object for cost, it does not calculate the cost based on the configured reference bandwidth but leaves it at default.	8.21.01
Perisistence/NonVol Problems Corrected in 8.31.01.0006	Introduced in Version:
Blade may reset with the following log message after a configuration change: <1>NonVol[5.tNVolCUp]cleanup:Remove() of first file on store=0, fileIndex=0 majorId=162 failed retval=3.	8.20.02
Platform Problems Corrected in 8.31.01.0006	Introduced in Version:
"show system utilization slot <slot>" allows invalid slot numbers such as 0.	6.00.02
Module might reset with message similar to <1>DistServ[4.tDsBrdOk]serverWatchDog.1(Config), client 63(PEME) in recv for 6007 tics “(0x00d0f9e4 0x0067b420 0x006707ac 0x01683264 0x00000000)” while PoE Controller is being updated.	4.21.09
Setting the MAC age time to 10 seconds may cause the tNtpTmr task to use high amounts of CPU processing time.	8.21.01
Very infrequently, when flooding frames that are larger than 10,000 bytes, a message similar to logged: Fuji[3.tNimIntr]Fuji TXQ MAIN intr: Fuji=7, Adr=0, Reg=0x00000004 There are no negative effects, other than the message being logged.	8.21.01
The "show running-config all" command will not display the command "no mac-address" if the mac address of an interface has not changed from its default value.	7.00.01
After a reset, some K-10 fabric modules may fail to reboot in a state with no front panel LED's and no console access or messages. A solid blue LED may be visible on the card itself (not the front panel) indicating the power failure. Recovery from this state requires a hard reset.	7.60.01

Platform Problems Corrected in 8.31.01.0006	Introduced in Version:
When a K-Series system is initialized with SFP modules (operating at 1Gbps) inserted into multiple ports of a KK2008-0204 line card or KK2008-0204-Fxx fabric, one or more of these ports may fail to successfully link with its respective peer.	8.21.01
Policy Problems Corrected in 8.31.01.0006	Introduced in Version:
Policy MAC address rules may not be immediately applied to flows on Tunneled Bridge Ports.	8.21.01
Port Interface Manager Problems Corrected in 8.31.01.0006	Introduced in Version:
Debug syslog message generated when an attempt to create a layer 3 interface is made with an out of range value: PiMgr[1.tConsole]generateIflIndex():retval=0;owner(0);mediaType(7);mediaPos(4096).	7.00.01
Port Status/Control Problems Corrected in 8.31.01.0006	Introduced in Version:
A message similar to: "Fuji MAC MAIN intr: Fuji=4, Adr=0, Reg=0x00080000" may be logged if a tagged packet between 10244 and 10247 bytes (inclusive) in length is received on a jumbo-enabled port.	7.00.01
RIPng Problems Corrected in 8.31.01.0006	Introduced in Version:
If RIP is configured with passive interfaces and RIPng is configured, the passive-interfaces will function correctly but be displayed under RIPng.	7.30.01
When a RIPng interface is configured to be passive, the passive setting takes effect but it is not displayed in show running.	7.30.01
RMON Problems Corrected in 8.31.01.0006	Introduced in Version:
Changing the owner string within an RMON command will result in a small memory leak.	5.01.58
Spanning Tree Problems Corrected in 8.31.01.0006	Introduced in Version:
BPDU's are not processed when marked for discard by Policy. The port role and state will be designated forwarding. When the port is an inter-switch link and the attached port is designated forwarding, a loop will form if there is redundancy.	4.00.50
The "set spantree backuproot" command completes successfully but will not modify the value.	8.20.02
A reset may occur when the CLI command "clear linecard <n>" is executed and a port on that line card is operationally an internal port in an MST region.	8.20.02
SYSLOG Problems Corrected in 8.31.01.0006	Introduced in Version:
Failed to set -101" error is seen during logging configuration.	3.11.04
"show support" or "debug messageLog message" result in an exhaustion of memory and an error message: "memPartAlloc: block too big".	1.07.19
Pushing the "Offline/Reset" button on K-Series main board modules will not display any messages indicating it was pressed.	8.01.01

Tunneling Problems Corrected in 8.31.01.0006	Introduced in Version:
ICMP need fragmentation messages sent to a L2 tunnel source were not properly decoded and forwarded.	8.21.02
The TOS or TrafficClass value is not properly propagated from the inner IP header to the outer IP header when performing L2, 4 in 6, or 6 in 4 encapsulations.	7.41.02
When using the command "show port counters errors nonzero", error counters for Virtual Private Ethernet ports (tbp.0.*) may incorrectly show non-zero values when no errors have actually occurred.	8.21.03
The description command was missing from the tunnel interface CLI.	7.41.02
The description CLI command is unavailable on a tunnel interface.	7.42.01

VLAN Problems Corrected in 8.31.01.0006	Introduced in Version:
If the Ingress Priority and VLAN were both 0, the packet was being treated as untagged.	8.01.01

VRRP Problems Corrected in 8.31.01.0006	Introduced in Version:
When using VRRP fabric route mode, if a packet is sent to a host that is connected to the router that is in fabric-route mode (through the master router), the ARP response for that host will not make it back to the master router. This is because the ARP response will be consumed by the router in fabric route-mode.	7.60.01
IPv4 VRRP advertisements may be transmitted with a TTL of 64.	8.20.02
When a VRRP VRID is the master, the "show ip vrrp" command will show the default "Master Advertisement Interval" when the correct value should match "Advertisement Interval" of the VRID (since it is the master).	8.22.01
When removing a VRRP VRID from configuration the VIP may not be available to use on subsequent VRIDs if the command for the VIP address is negated just before the VRID is disabled.	8.21.01
A syslog messages may be displayed when VRRP is configured in accept mode and a blade resets. The message displayed will look similar to this: "Failed: ipAddressDelete 101.14.10.1 not create for accept mode on vlan 2100"	8.21.01

Problems Corrected in 8.22.03.0006

802.1x Problems Corrected in 8.22.03.0006	Introduced in Version:
When using EAP Authentication Methods that require passing certificates, if the packets for those certificates are greater than 1760 bytes, a portion of those packets may be transmitted with invalid data.	8.20.02

Host Problems Corrected in 8.22.03.0006	Introduced in Version:
ICMP echo requests to IP interface addresses exceeding 100 per second will not all be answered.	8.20.02

VRRP Problems Corrected in 8.22.03.0006	Introduced in Version:
IPv4 VRRP advertisements may be transmitted with a TTL of 64.	8.20.02

Problems Corrected in 8.22.02.0012

ARP/ND Problems Corrected in 8.22.02.0012	Introduced in Version:
The number ARP/ND packets dropped due to existing rate limiters is not accessible.	8.02.02
ARP Problems Corrected in 8.22.02.0012	Introduced in Version:
When the router receives a broadcast IP packet it may generate an ARP request to resolve 255.255.255.255.	8.21.01
Hardware Problems Corrected in 8.22.02.0012	Introduced in Version:
When in Provider Bridge Mode, after a chassis reboot, all ports on all linecards will not forward packets in accordance with relevant switching rules. When in Provider Bridge Mode, after individual linecards are reset, removed, and re-inserted, or inserted for first time, all ports on those linecards will not forward packets in accordance with relevant switching rules.	8.01.01
LLDP Problems Corrected in 8.22.02.0012	Introduced in Version:
"set lldp port tx-tlv poe" and "set lldp port tx-tlv med-poe" commands may be missing in a chassis that contains non-poe blades.	8.01.01
Persistence/Nonvol Problems Corrected in 8.22.02.0012	Introduced in Version:
When seeing the following error: <3>Default[1.tusrApplnit]moduleIsInSameLocation():Unable to open "/flash 2/moduleRecords/chassisSlotInfo.rec" for reading is usually due to a chkdisk repair of the DOS file system, any lost configuration is not recovered from other blades in the chassis.	8.21.01
Modules with corrupt file systems could get caught in a reboot loop, if the parent and subdirectory structure are bad. A message similar to below would be seen in the log: Message 8/346 Fatal Error 08.22.01.0020 04/07/2014 14:39:35 ERROR: file system check	8.21.01
Tunneling Problems Corrected in 8.22.02.0012	Introduced in Version:
When a tunnel becomes operationally up, proper forwarding to some tunnel destinations may not start or resume.	8.21.03
VLAN Problems Corrected in 8.22.02.0012	Introduced in Version:
When ports are configured to egress VLAN 1 tagged, after a system reset most ports are removed from VLAN 1 static egress. As a workaround, remove a linecard port from VLAN 1 static egress. By doing this, the correct configuration will be restored after system reset.	7.30.01
VSF Problems Corrected in 8.22.02.0012	Introduced in Version:
The 'clear bonding chassis' command does not clear an 'inactive' chassis after a system reset.	7.62.00

Feature Enhancements in 8.22.01.0022

Hardware Support Enhancements in 8.22.01.0022
Support for additional 10Gb active optical direct attach cable transceivers: 10GB-F10-SFPP 10Gb, Active optical direct attach cable with 2 integrated SFP+ transceivers, 10m 10GB-F20-SFPP 10Gb, Active optical direct attach cable with 2 integrated SFP+ transceivers, 20m
Captive Portal Re-direct Feature Enhancements in 8.22.01.0022
Captive Portal uses HTTP redirection to force a client's web browser to be redirected to a particular administrative web page. A network administrator can use this feature for authentication purposes (a user login and password), payment (i.e., at an airport hotspot), or usage-policy enforcement. This feature is an extension of the Policy infrastructure, where Policy Roles may be configured to force redirection of HTTP traffic.
OSPF Default Route Injection Feature Enhancements in 8.22.01.0022
Support for directly advertising a default route into OSPF has been added via the "default-information originate" command. There are two options available, advertise the default route into the OSPF domain, provided the advertising router already has a default route. Alternatively, advertise the default route regardless of whether the advertising router already has a default route. Option 2 is chosen by adding the "always" keyword to the "default-information originate" command.

Problems Corrected in 8.22.01.0022

802.1x Problems Corrected in 8.22.01.0022	Introduced in Version:
802.1x may not require an 802.1x supplicant to wait the configured quiet period (set dot1x auth-config quietperiod <period> <port-string>) to start a new authentication after a failed authentication.	8.21.01
ARP Problems Corrected in 8.22.01.0022	Introduced in Version:
If the system sends a packet to a remote IP address, an ARP request for the remote IP address may be transmitted on a configured interface.	8.21.01
The "show running-config" command may not display all static ARP/ND entries that are configured.	7.00.01
Auto-Negotiation Problems Corrected in 8.22.01.0022	Introduced in Version:
If "clear port advertise *.*.*" is executed on a system on which not all ports support auto-negotiation, the message "failed to set ifMauAutoNegCapAdvertisedBits on port x.y.z" will be displayed for each port that does not support auto-negotiation.	7.00.01
"Setting ifMauAutoNegRemoteFaultAdvertised (1.3.6.1.2.1.26.5.1.1.12) MIB value to offline(2) for a port brings the port down until reset, even if ifMauAutoNegRemoteFaultAdvertised value is changed to noError (1)."	5.11.21
Auto-Tracking Problems Corrected in 8.22.01.0022	Introduced in Version:
Auto-tracking radius-timeout-profile and radius-reject-profile per port configuration may allow profile ID configuration that is greater than allowed by the system.	8.01.01
Outputted log event from auto-tracking and quarantine-agent "Unable to set policy rule" port string is not user friendly.	8.01.01
If auto-tracking multiauth sessions are configured to be allowed on authentication required ports then unauthenticated traffic matching the auto-tracking multiauth session will be switched by the system.	8.01.01

Chassis Bonding Problems Corrected in 8.22.01.0022	Introduced in Version:
Configure from file fails when enabling bonding. When this error occurs a message similar to "<2>System[12]Detected missing or reset module, aborting configure" is logged.	8.21.03
CiscoDP Problems Corrected in 8.22.01.0022	Introduced in Version:
Cisco VTP packets are not forward when Cisco CDP is enabled.	7.91.01
Host Services Problems Corrected in 8.22.01.0022	Introduced in Version:
"show system utilization storage" will report inaccurate size and available size for USB drives greater than 2G.	7.60.01
Usually on a reboot after an uncontrolled reset (power-loss, board pull, exception, DSI, watchdog reset) you may see the following file system error during initialization: /flash2/ - disk check in progress ... "/flash2/usrroot/someFileName" too many clusters in file, adjusted. Errors detected. All corrections stored to disk and lost chains recovered.	7.30.01
Continuous poll of TCP or UDP MIBs may result in the exhaustion of memory resulting in an out of memory reset action on a specific slot.	7.40.00
Performing the "show vlan portinfo" CLI command under configurations where there are many VLANs in use may lead to the CLI becoming inoperable or the system to reset.	8.21.01
In the unexpected event where resources needed to transmit a routed L3 Multicast packet failed to be obtained, a blade will reset, and leave a message in log similar to: Message 9/333 Exception PPC750 Info 08.21.02.0002 12/21/2013 23:22:53 Exc Vector: DSI exception (0x00000300) Thread Name: tDispatch Exc Addr: 0x0191e77c Thread Stack: 0x06921000..0x06914000 Stack Pointer: 0x06920f40 Traceback Stack.....	7.00.01
Doing a set on a large range of data could cause a board reset. Example: cfm vlan-table primary 99 selector 1-98,100-4094. The syslog will show an error similar to below: <1>NonVol[1.tNVolCUp]cleanup:Remove() on store=0, fileIndex=2863311530 majorId=140 failed retval=8, write_file_num=50 (0x00d12590 0x00a79af4 0x 00a81504 0x01686324 0x00000000) A core file will be generated.	8.21.01
Infrequently, when switch is adding (encapng) tunnel headers, a message may be logged similar to: <163>Dec 5 15:11:28 100.10.10.22 PiMgr[16.tDispatch]piMgrBindSystemPortAndHwPort(0,0):Port(s) are already bound. pimSystemPortToHwPort[0]=0x8000;pimHwPortToSystemPort[0]=0x100 <163>Dec 5 15:11:28 100.10.10.22 PiMgr[16.tDispatch]piMgrBindSystemPortAndHwPort(0,0):Port(s) are already bound. pimSystemPortToHwPort[0]=0x8000;pimHwPortToSystemPort[0]=0x100 <165>Dec 5 15:11:28 100.10.10.22 PiMgr[16.tDispatch]piMgrHwPortRxlcpu (131072,2,63,0,0x7eb82188,1052):piMgrBindSystemPortAndHwPort(0,0) failed;hwPort=0;portCount=43;tmpBufLen=700.	7.40.00

IP SLA Problems Corrected in 8.22.01.0022	Introduced in Version:
The system may encounter a deadlock when retrieving statistics for SLA entries while SLA entries are at the beginning of a new test cycle.	8.01.01
The user configures SLA entries to run continuously using the 'recurrence' command, which is a static interval in seconds. The user can avoid the deadlock by refraining from retrieving the SLA data for several seconds while a new test cycle is starting. This applies to all SLA entries starting a new test cycle.	
IPv6 Forwarding Problems Corrected in 8.22.01.0022	Introduced in Version:
IPv6 packets destined to a remote subnet, whose route has a link-local nexthop address and deferred to neighbor discovery for MAC address resolution, may be transmitted with a destination MAC address of 00:00:00:00:00:00.	7.40.00
LAG Problems Corrected in 8.22.01.0022	Introduced in Version:
Prior to 8.11.1 the K-Series did not support LAGs with more than 16 ports. You could set them up, but Source Port Exclusion would not work. After 8.11.1 if you have LAGs with more than 16 ports, Source Port Exclusion may not work.	8.01.01
Mirroring Problems Corrected in 8.22.01.0022	Introduced in Version:
Configuring an Rx port mirror with a LAG destination fails to be set.	8.11.01
When the device acts as a Pseudowire tunneled endpoint, the de-capsulated packet would not egress out a software bond port.	8.21.01
The blade may be reset (and continuously reset) with the following messages if the LAG used by IDS mirror has more than 2 ports: <3>Dune[5.tSlac]Err_id=0x16a1d3af: error in fap21v_sch_is_subflow_valid() ExitPlace (40) Params(0,0,0,0,0) <0>Dune[5.tSlac]Err_id=0x16a1d3af: error in fap21v_sch_is_subflow_valid() ExitPlace (40) Params(0,0,0,0,0)	8.11.01
Multi-Auth Problems Corrected in 8.22.01.0022	Introduced in Version:
When multiauth sessions-unique-per-port is disabled and multiple multiauth agents are enabled, a failure of one agent may cause additional agents to fail, outputting the error message "Unable to set policy rule for mac XX-XX-XX-XX-XX-XX on system port 443".	8.01.01
Standardized multi-authentication session and idle timeout maximum values to be 172800 seconds.	8.01.01
With sessions-unique-per-port disabled and multiple authentication agents enabled and active, when a session moves from one slot to another, it may not session or idle timeout appropriately.	8.01.01
Multiauth Quarantine Agent sessions do not correctly apply policy if the policy mactable response is set to tunnel.	8.01.01

OSPF Problems Corrected in 8.22.01.0022	Introduced in Version:
If an OSPF area ID is changed while an interface is transitioning to the DOWN state, an assert may occur in thread tRtrPtcls with the following log: "SMS assert in qopmmim2.c at line 1958 : is one of if_cb->repl.row_data.oper_status 4 qopm_mib_if_product_data.oper_states.down or qopm_mib_if_product_data.oper_states.act_failed"	8.11.01
If multiple OSPF processes learn the same route, metrics are not compared between them. Both routes are installed in the route table as the administrative distance is the same and cannot be changed for an individual process.	7.00.01
An OSPF NSSA ABR configured as "transrole always" may not always be the translator.	8.01.01
OSPF log-adjacency cannot be removed with a no log-adjacency under router ospf <pid>.	8.01.01
If OSPF is configured to run BFD on a non-existent interface, the interface will not be displayed in show running. When the interface is created, the display will show, and BFD will run on that interface.	8.21.01
The display of an OSPF external LSA metric has the first byte truncated so the largest number displayed is 4095, though the real value may be up to 65535.	7.00.01
OSPFv3 Problems Corrected in 8.22.01.0022	Introduced in Version:
If debug logging is turned on for OSPF, and filter route-maps are in use, the route-src is seen as 0.0.0.0 for local routes from our router ID.	8.01.01
When an OSPFv3 NSSA translator is configured to always be translating, it will not always translate if a higher router ID is also eligible.	8.01.01
PIM-DM Problems Corrected in 8.22.01.0022	Introduced in Version:
Changing the route to sources may result in an assert similar to "<0>sms[2.tRtrPtcls]SMS assert in qptuwsn2.c at line 669 : (null) QPTM_USM_S_G_GET_JDES(s_g) 0 (null) 0 "	8.21.01
The use of IGMP V3 to PIM DM may cause crash.	8.21.01
The use of exclude mode in IGMPv3 may result in a PIM DM assert.	8.21.01
PIM-SM Problems Corrected in 8.22.01.0022	Introduced in Version:
IGMPv3/MLDv2 source-specific reporter state is missing from layer3/router.	7.30.01
The internal IGMP/MLD database may be inconsistent across all modules after a bonded system is segmented, then re-joined. This can lead to incorrect multicast operation and/or inconsistent aging of entries.	7.30.01

Platform Problems Corrected in 8.22.01.0022	Introduced in Version:
<p>If a macsource policy is applied, packet statistics from the following apps may not be valid: Smon stats Rmon Host/Matrix Router ACL Policy Routing Tunneling Policy</p> <p>In addition, if any SMON stats are enabled, messages similar to: SMON[6.tSmonCnt]getHwPrioStats(ge.6.3,0): packet count < previous 2/172401; detected 1 times, may be logged.</p>	8.21.01
<p>Repeated blade removals on the K-Series platform may cause the auto-tracking and quarantine agent components to run out of port memory resulting in this error message "<code><0>NIM[11.tNismProc]CNISM::checkStateTime():Time to process dataInited has expired NIM(2);Apps:32,33,</code>" which will reset the system.</p>	8.01.01
<p>On a bonded system, a file may be left in an improper state which is identified and corrected by the file system verification and recovery tool that is run at each boot up. If a file in this state is detected, a set of messages like the following will be displayed during boot up.</p> <pre>/flash2/ - disk check in progress ... "/flash2/usroot/foobar672" too many clusters in file, adjusted. Errors detected. All corrections stored to disk and lost chains recovered.</pre> <p>This state is recoverable and should have no effect on the normal operation of the file system.</p>	7.60.01
<p>A 10G port with a 1G SFP doesn't propagate its advertised speed to link partner.</p>	8.11.04
<p>System logs the message "<code>bcmStrat[11.tLcIntr]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000010, IP2_INTR_STATUS=0x00000000, XQPORT_INTR_STATUS.xgport0=0x00000000, XQPORT_INTR_STATUS.xgport1=0x00000000</code>" and resets.</p>	8.01.01
<p>System logs the message "<code>bcmStrat[11.tLcIntr]MEM_FAIL_INT_STAT=0x00200000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000</code>" and resets.</p>	8.01.01
<p>Doing a "dir" on a remote directory with a large number of files has a long delay before the output starts. Ex: A directory with 1000 files may take around 34 seconds before being displayed.</p>	7.91.01
<p>When doing a "dir" from CLI, if the directory is currently being modified (file being added/deleted) an incomplete listing can be returned.</p>	7.91.01
Policy Problems Corrected in 8.22.01.0022	Introduced in Version:
<p>Unable to clear all policy profiles with a single CLI command.</p>	1.07.19
<p>VLAN authorization commands allow for configuration and display of tunnel bridge ports although they are not supported port types for VLAN authorization.</p>	8.21.01
PWA Problems Corrected in 8.22.01.0022	Introduced in Version:
<p>The "<code>set pwa ipaddress <ip-address></code>" CLI command allows invalid values for the <code><ip-address></code> field.</p>	4.00.50

PWA Problems Corrected in 8.22.01.0022	Introduced in Version:
PWA occasionally becomes unresponsive under heavy load. Device resets with this message in the log: <0>PWA[1.tPwaHtWD]pwaHttpReadWatchDog expired!	4.00.50
RADIUS Problems Corrected in 8.22.01.0022	Introduced in Version:
RADIUS Server sticky sessions count may be inaccurate after session terminations.	8.11.01
If the radius algorithm is changed while multiauth sessions are active, incorrect sticky session counters may be both displayed and used by the system.	8.11.01
RADIUS Dynamic Authorization responses cannot be sent in response to disconnect or change of authorization RADIUS Dynamic Authorization requests, resulting in the error message "Unable to transmit the RADIUS frame" and retransmissions from the RADIUS server.	8.21.01
RMON Problems Corrected in 8.22.01.0022	Introduced in Version:
Swapping modules on a K-Series chassis may result in incorrect information being added to the output of the ""show config"" command consisting of multiple lines of "set rmon stats 0".	7.00.01
Security Problems Corrected in 8.22.01.0022	Introduced in Version:
PWA will discard HTTP GET requests with HTTP headers that exceed 2048 bytes.	4.00.50
K-Series linecards generate EAPOL frames with 00-00-00-00-00-00 source MAC address. This may cause switches to drop these frames.	8.21.01
Spanning Tree Problems Corrected in 8.22.01.0022	Introduced in Version:
When the root port of a bridge receives a value for remainingHops greater than 63, there will be overflow when storing the value. For example, if the remainingHops value is 100, it will be stored as 36. This is because the field width is six bits. This is enough to hold the standard defined maximum value of 40. This is true for both cistRemainingHops and remainingHops for any MSTI. This only has a practical effect within an MST region. These values are not used external to the region. Note that values greater than 40 are non-conformant as of 802.1Q-2005 so are not likely to be seen.	8.21.01
In a multi-blade chassis or stack, when setting Spanning Tree stpmode to the value _none_, the non-master blades will still operate as if the mode were _ieee8021_ until those blades are reset.	8.21.01
When a device in a multi-blade chassis or a bonded setup fails, and that device contained the spanning tree root port for the bridge, the new root port, if there is one, may not take on its root role, and therefore be stuck in a discarding state. If this does occur then a workaround for this is to disable the new root port (which will show a role of alternate port) and then re-enable the port.	8.21.01
Static Routes Problems Corrected in 8.22.01.0022	Introduced in Version:
Static route leaking between non-global VRFs does not work. The routes are not promoted to the FIB.	8.21.01

Tunneling Problems Corrected in 8.22.01.0022	Introduced in Version:
Host generated IPv6 packets that are encapsulated into an IP or GRE tunnel could have an incorrect DIP.	7.60.01
The software forwarding path was retrieving the GRE header when it was not part of the flow. This would sometimes cause the IP-in-IP to be translated as a L2 IP-in-IP flow.	7.62.02
When the device acts as a Pseudowire tunneled endpoint the de-capsulated packet would not egress out a software bond port.	8.21.01
For pseudo-wire tunnels, the soft forwarding path was not adding the Chassis Bond header when going across a software bond.	8.21.01
The egress point of a Tagged IPv6-GRE (with GRE Keyword) tunnel would not decrement the inner IPv4 TTL or change the TOS due to hardware limitations.	8.21.01
L2 Tunnels across a Software Bond was not updating the L2 IP's total length field when adding the GRE header and Chassis Bondheader to the egress packet.	8.21.01
Infrequently, when switch is adding (encapng) tunnel headers, a message similar to: <0>chassis[9.tBcastStRx]powerSupplyComputeModuleConsumedPower :Invalid uplink number 0x00 detected in remote info table, may be logged.	7.40.00
Infrequently, when switch is adding (encapng) tunnel headers, a message similar to: <3>chassis[1.tBcastStRx]remoteModuleInfoPowerUpdate(6,""):Unsupported board type found., may be logged.	7.40.00

VLAN Problems Corrected in 8.22.01.0022	Introduced in Version:
Ports may not be added to the default VLAN's static untagged egress list after the linecard initializes.	7.30.01

VRF Problems Corrected in 8.22.01.0022	Introduced in Version:
When clearing a VRF router config, "clear router vrf <NAME>" the error message "Error destroying BFD process 22185496: AMB_RC_NO_SUCH_OBJECT" is displayed, but has no adverse effect.	8.21.01

Problems Corrected in 8.21.03.0001

802.1x Problems Corrected in 8.21.03.0001	Introduced in Version:
K-Series linecards generate EAPOL frames with 00-00-00-00-00-00 source MAC address. This may cause switches to drop these frames.	8.21.01

Feature Enhancements in 8.21.02.0001

K-Series Chassis Bonding Enhancements in 8.21.02.0001
The ability to chassis bond two like K-Series chassis (K6/K6 or K10/K10) to form a single bonded unit. Requires the K-EOS-VSB license to be applied per chassis.

Virtual Private Ethernet Service Enhancements in 8.21.02.0001
L2VPN capability to connect Layer 2 networks transparently over a Switched or Routed IP core network using GRE or IP tunnels. With this feature, Layer 2 traffic within the switch (VLAN's) can be switched into and out of the encapsulated tunnel to be transmitted across the network.

Bi-directional Forwarding Detection (BFD) Enhancements in 8.21.01.0001
Support for BFD probe as a mechanism to detect a communications failure with an adjacent system forwarding plane. This version of BFD probe supports monitoring OSPF neighbors.

PIM Dense Mode Enhancements in 8.21.02.0001

Support for PIM-DM to allow dense mode multicast distribution utilizing PIM-DM flood and prune mechanism to build source distribution trees for multicast flows.

LAG Enhancements in 8.21.02.0001

The LAG capacity has been increased to 62 LAGs in K-Series chassis.

Tunnel Changes in 8.21.02.0001

The K-Series IP tunnel capacity has been decreased to 16 tunnels.

Remote Port Mirroring Enhancements in 8.21.02.0001

Remote port mirroring is now included in the base firmware and does not need a feature license.

CLI Enhancements in 8.21.02.0001

Show vlan portinfo CLI – CLI command has been added to display VLAN information regardless of forwarding state.

Added configuration to allow the UDP broadcast helper address to be configured to accept a classful network address. Global configuration mode: 'ip forward-protocol allow-classful'

'show ipv6 interface' list all multicast groups the VLAN has joined.

A command to disable DHCP server logging has been added.

'show support', now includes 'show linkflap' status.

Webview Enhancements in 8.21.02.0001

The left-hand WebView menu has been changed for better browser compatibility.

HOST Enhancements in 8.21.02.0001

Improved rate limiting and prioritization for Host traffic.

VLAN Enhancements in 8.21.02.0001

Support for 2 secondary VLANs per primary VLAN has been added.

Problems Corrected in 8.21.02.0001

802.1x Problems Corrected in 8.21.02.0001	Introduced in Version:
EAPOL frames may be switched when multiauth is in either forced-auth, auth-optional, or auth-required port mode.	7.00.01
802.1x global enable status may become enabled during a single board reset in multi-blade system.	8.11.01
Anti-Spoofing Problems Corrected in 8.21.02.0001	Introduced in Version:
'show config antispoof' may not display class names correctly.	8.01.01
IPv6 forwarding can be disabled on an interface that has IPv6 checkspoof configured.	7.31.02
Setting the antispoof notification interval to 0 and antispoof to enabled will consume all resources and cause the switch to be unresponsive.	8.01.01

Anti-Spoofing Problems Corrected in 8.21.02.0001	Introduced in Version:
Modifying the etsysAntiSpoofThresholdType MIB leaf to a value other than 1 (IPv4) will result in the following syslog: "Internal error: unknown remapping case (3) in make_error_pdu". The setting will not take effect as only the IPv4 Threshold Type is currently supported.	8.01.01
The "clear linecard" CLI command may not clear antispoof port configuration to default settings.	8.01.01
ARP Problems Corrected in 8.21.02.0001	Introduced in Version:
In very rare instances a module may complete its boot process with ARP/ND entries that are present on all other blades but missing from the blade that just booted.	7.00.01
Occasionally syslog messages may appear indicating that a MAC address for an existing ARP or ND entry has changed from: ec-c1-e5-ec-c1-e5 to a different MAC address. The MAC in question is a special purpose MAC address and the message does not indicate anything has gone wrong.	8.11.01
The commands "show arp" and "show ipv6 neighbors" will print "(null)" in the port column when the MAC address for the ARP/ND entry is a static multicast MAC address.	7.00.01
The Static and Dynamic ARP/ND limits incorrectly allow more entries than they should. The Dynamic limit should be 4000 and the static limit should be 512.	Unknown
ARP/ND Problems Corrected in 8.21.02.0001	Introduced in Version:
When populating the ARP/ND static ARP table (either via configuration or during the boot cycle) the router will display a message indicating the chassis is 50% full. The message implies that the dynamic ARP/ND entries are triggering the messages but the message actually refers to the static ARP/ND limit.	Unknown
Stale ARP/ND entries are not removed if a filter database entry exists for the MAC address of the ARP/ND entry.	7.71.02
Auto-Tracking Problems Corrected in 8.21.02.0001	Introduced in Version:
The Help string for auto-tracking port radius-reject-profile command is incorrect.	8.01.01
Bonding Problems Corrected in 8.21.02.0001	Introduced in Version:
If a VSB system is segmented, and the systems have different firmware versions, when the bond link is established between the two systems, a common image is not distributed and the system does not complete the bonding process.	Unknown
"clear linecard <slot>" for a line card that is present will generate SwitchChipFreeOrlResource errors.	7.31.02
Configuration of RMON stats and history options on existing default entries will be lost on reboot.	8.11.01
When attempting to enable more than 32 bonding ports a message with following format may be logged. "63>Oct 31 12:49:09 webview_test Bond[7.tBondCfg]clearDisableReason: Cannot add port=7d, max bonding ports are configured"	7.91.01
VSB license CLI commands treat the slot keyword the same as the chassis keyword.	7.40.00

Converged End Point (CEP) Problems Corrected in 8.21.02.0001	Introduced in Version:
Active Convergent End Point (CEP) entries will remain even if CEP is disabled globally or on a per-port basis.	6.02.04
CEP detection-id enabled/disabled state will not be displayed in 'show config' if set to disabled.	7.91.01
CFM Problems Corrected in 8.21.02.0001	Introduced in Version:
The CLI command "show cfm default-md VID <vid-number>" will display an incorrect selector type when attempting to display a single CFM Default MD.	7.91.03
CLI Problems Corrected in 8.21.02.0001	Introduced in Version:
If the "set system lockout port" is enabled and a user fails to login via SSH the maximum allowed attempts, the user login gets locked but the port lockout fails to get locked.	7.40.01
The "show config quarantine-agent" command may leak memory.	8.01.01
The "show config dot1x" command may leak memory.	8.11.01
The "show config auto-tracking" command may leak memory.	8.01.01
COM port for K-Series chassis mistakenly associated with linecard 1 instead of the chassis fabric module.	7.30.01
Issuing a "show config" or "show config pwa" will cause a small amount of memory to leak per iteration.	8.11.01
The traceroute command only executes once inside a CLI 'loop'.	7.00.01
Memory leak executing CLI command "show snmp counters".	4.05.08
COS Problems Corrected in 8.21.02.0001	Introduced in Version:
COS ORL actions may be applied to the equivalent port on the receiving blade if the egress port is on a remote blade.	7.00.01
"processCosPortConfig" message log entry may occur if removing and showing COS configuration at the same time.	7.00.01
DVMRP Problems Corrected in 8.21.02.0001	Introduced in Version:
DVMRP may crash when sending upstream prune after routes change.	7.60.01
DVMRP incorrectly requires the K-EOS-L3 license, while IPv6 PIM was incorrectly not requiring the K-EOS-L3 license.	8.11.01
ECMP Problems Corrected in 8.21.02.0001	Introduced in Version:
The CLI command to show the current setting of the IPv6 ECMP forwarding algorithm is missing.	7.00.01
Filter Data Base (FDB) Problems Corrected in 8.21.02.0001	Introduced in Version:
When multiple static mac address (unicast and/or multicast) are configured, at boot time messages similar to: "FilterDb[2.tusrApplnit]fast_add restore (local) failed 14,60968" may be logged. There are no negative consequences, other than the messages being logged.	8.11.03

Filter Data Base (FDB) Problems Corrected in 8.21.02.0001	Introduced in Version:
If the source port of a static unicast MAC address is changed without first deleting exiting entry and recreating it, messages similar to: "FilterDb[2.tusrApplnit]restored duplicate(60126112,1 - 26-00-01-02-03-04.5 on 2" may be displayed at boot time. In addition, after reboot an entry may not restore with correct source port, or a deleted entry may re-appear.	7.00.01
GVRP Problems Corrected in 8.21.02.0001	Introduced in Version:
The ctDot1qVlanGvrpRestrictedStatus MIB object cannot be set and the "set gvrp vlan" CLI command is ignored in provider bridge mode.	7.91.01
High Availability Upgrade (HAU) Problems Corrected in 8.21.02.0001	Introduced in Version:
CLI does not reject out of range slot lists when configuring HAU upgrade groups. For example, "set boot high-availability group 1 1-256" should result in a CLI error, but instead the command is accepted and slots 1-N (where N is the highest slot in the system) are assigned to group "1".	7.60.01
IGMP Problems Corrected in 8.21.02.0001	Introduced in Version:
When using SSM with IGMP, SSM packet drop counters may be incorrect.	7.30.01
After a chassis segments and reforms, message of the form: "Error: Mis-Matching MCI chain data tag:1 v6:1 for MCI:131 tag:1 v6:0" are displayed and the IGMP database may become corrupted.	7.00.01
IGMP will not correctly update the drop counter for leaves with a bad group address.	8.11.01
While running IGMP v3 with 'include' source-list, a module crashes with a message containing to: "Clgmp::GroupTableAddPortToGroupEntry Src port mismatch".	7.30.01
IGMP ignores reports immediately after booting until the connected interfaces are populated in the Route Table.	7.31.02
IGMP running in v1 mode will drop queries for missing Router Alert.	8.11.01
When loading a configuration from a file that contains IGMP config which has "set igmp disable <x>" where x is the VLAN, any command set after this will re-enable the IGMP config for this VLAN.	7.00.01
It is possible for flows to continue egressing out a port that was removed from an IGMP static configuration.	7.91.01
It is possible when using the "clear igmp static <x.x.x.x> <vlan> modify include <yyy> exclude <zzz>" command, that traffic will still flow to the specific port being removed.	8.01.01
IPv4 Forwarding Problems Corrected in 8.21.02.0001	Introduced in Version:
'ip checkspoof strict-mode' will no longer be applied to packets destined to host address configured on packet's ingress interface.	7.00.01
IPv6 Forwarding Problems Corrected in 8.21.02.0001	Introduced in Version:
Packets received on interfaces where IPv6 forwarding is disabled and destined to host address configured on a different interface are incorrectly delivered to the host.	7.00.01
An IPv6 address configured on a VLAN interface with a 128-bit mask is not reachable.	7.00.01

IPv6 Forwarding Problems Corrected in 8.21.02.0001	Introduced in Version:
IPv4-mapped IPv6 addresses and IPv4 compatible addresses are not supported but are accepted by the Command Line Interface. When entered an error occurs but the address in some cases appears to be valid when in fact it is not working.	8.01.01
IPv6 Neighbor Discovery Problems Corrected in 8.21.02.0001	Introduced in Version:
Attempts to send packets from the host to a directly connected IPv6 link-local address will not work because the incorrect MAC address will be used as the destination MAC address of the destination Link-Local address.	8.11.01
Jumbo Problems Corrected in 8.21.02.0001	Introduced in Version:
Invalid sized non-tagged packets of size 1519 to 1522 bytes, and tagged packets of size 1523 to 1526, received on non-jumbo enabled ports are correctly dropped. However, the SA MAC is incorrectly learned in MAC table.	7.00.01
Invalid sized non-tagged packets of size 10240 to 10243 bytes, and tagged packets of size 10244 to 10247, received on jumbo enabled ports are correctly dropped. However, the SA MAC is incorrectly learned in MAC table.	7.00.01
For some flows that require reframing, if any one of the first few packets in flow are jumbo sized, those packets could be dropped (and not forwarded).	7.60.01
LACP Problems Corrected in 8.21.02.0001	Introduced in Version:
LACP marker response not within frame rate limitation constraint for slow protocols.	1.07.19
In some instances, LACP is not setting collecting and distributing bits to false after a partner PDU change, resulting in the port not leaving the LAG as it should.	1.07.19
The 'clear linecard' command may not clear LACP port parameters to default values.	7.91.01
A set of a lag port attribute may fail without a message at the console.	1.07.19
Distribution of traffic over the ports in a LAG could vary over 10% port-to-port from a uniform distribution when an odd number of ports are in the LAG.	7.30.01
In rare instances, a port that joins a LAG briefly then stays down/is removed from the lag may still be considered an available egress port for a few percent of LAG traffic which would be undelivered. A subsequent change of state of any of the ports in the LAG or the addition/removal of a port in the LAG will clear the condition.	5.01.58
LLDP Problems Corrected in 8.21.02.0001	Introduced in Version:
MIB lldpStatsRemTablesAgeouts is not incremented when a neighbor ages out.	7.00.01
lldpStatsRxPortTLVsDiscardedTotal may not increment for non-support LLDP TLVs.	7.00.01
LLDP Management Address TLV has incorrect interface index.	7.00.01
LLDP Link Aggregation TLV was using a format that was deprecated in IEEE 802.1AB-2009.	7.00.01
The PoE TLV in a transmitted LLDP packet correctly shows a TLV length of 12, but the extended information shows an incorrect Type/Source/Priority (TSP) field, PD requested power value, and PSE allocated power value.	8.11.01
Occasionally while under heavy processing load, LLDP may cause the system to crash.	7.62.00

LLDP Problems Corrected in 8.21.02.0001	Introduced in Version:
Occasionally when clearing a linecard on a K-Series chassis, the system will crash due to a server watchdog of NDS. The system will display a log message like "serverWatchDog.1(Config), client 37(NDS) in rcv for 6659 tics".	7.30.01
MAC AUTH Problems Corrected in 8.21.02.0001	Introduced in Version:
Setting the authallocated macauthentication field ("set macauthentication authallocated <port string>") to a value of 0 does not correctly result in an outputted error although the value is not set.	5.01.58
Mirroring Problems Corrected in 8.21.02.0001	Introduced in Version:
The "clear port mirroring ori" command does not disable mirror outbound rate-limiting.	8.11.01
When mirroring, the physical loopback port does not go down when the tunnel goes operationally down.	8.11.03
The "clear port mirroring" and "set port mirroring [enable disable]" commands do not set the lower numbered destination ports if the destination port-string is in descending order (i.e. tg.4.3;tg.4.2). These commands function properly when the destination port-string lists the lowered ports first.	7.91.01
"Doing a "show config" may cause a DSI exception reset when a mirror index is configured with a large number of ports."	7.72.01
MULTI AUTH Problems Corrected in 8.21.02.0001	Introduced in Version:
Executing 'show multiauth session port <port-string>' might result in an error.	7.30.01
The 'show multiauth station port' command displays multiple entries for each provisioning agent type.	5.01.58
If 'multiauth sessions-unique-per-port' is disabled and CEP multiauth sessions are moving from one port to another RADIUS accounting data may be output inconsistently for that session.	8.11.01
Multiauth sessions that port roam may not session timeout at the expected time.	8.11.01
Clearing multiauthentication stations using the etsysMultiAuthStationClearUsers MIB leaf may cause the multiauthentication software to treat the clearing as a failure for both logging and trap purposes.	7.72.01
When multiauthentication traps for authentication success, authentication rejection, or port termination are enabled and are being sent, they result in duplicate notice level log events that indicate the same or similar information.	7.00.01
Networks utilizing multiauth session or idle timeouts greater than 65535 may have sessions that timeout inaccurately.	6.11.01
NETFLOW Problems Corrected in 8.21.02.0001	Introduced in Version:
When NetFlow is enabled, very infrequently, an error message similar to: "<3>netflow[4.tNetflow]netflow_record_processing_task - unexpected error taking semaphore" may be displayed. When that message is logged, a single frame, which can consist of anywhere from 1 to 30 NetFlow records, is dropped and will not be delivered to NetFlow collectors.	8.01.01

NETFLOW Problems Corrected in 8.21.02.0001	Introduced in Version:
When NetFlow export-data higher-layer is enabled, messages similar to: "PiMgr[7.tMcnxPer]generatelfIndex():retval=7;mediaType(0);mediaPos(8)" may be displayed. For each message generated, a NetFlow record with an invalid destination interface will be sent.	8.01.01
Very infrequently, when NetFlow export data higher layer is enabled, messages similar to: "PiMgr[7.tMcnxPer]generatelfIndex():retval=0;owner(1);mediaType(7);mediaPos(0)" may be logged. For every message logged, a NetFlow record would be generated with invalid source and/or destination interfaces.	8.01.01
OAM Problems Corrected in 8.21.02.0001	Introduced in Version:
Disabling OAM on a port does not clear the OAM or ULD operstatuscause.	7.30.01
OSPF Problems Corrected in 8.21.02.0001	Introduced in Version:
When running OSPF, and using the passive-interface default command, an assert could occur in thread tRtrPtcls with the following log, "SMS assert in qopmmim5.c at line 879 : (null) AVLL_IN_TREE(if_cb->active_if_tree_node) 0 (null) 0".	7.00.01
If an OSPF interface running over a tunnel is explicitly configured as point-to-point this is displayed in the config even though it is the default.	7.41.02
OSPFv2 will accept the configuration of an invalid nssa-range and display it incorrectly.	7.00.01
Configuring an OSPF cost metric outside the range results in an unclear message error.	8.11.01
When issuing a "clear ip ospf process" and multiple OSPF processes exist, the ambiguous message "Resetting the OSPF process" is seen multiple times.	7.00.01
When running OSPFv2 or v3 with auto-cost reference bandwidth and tracked objects, it is possible with multiple cost changes to have the router LSA not reflect the cost seen on the interface.	8.11.01
If OSPF logging is enabled and multiple OSPF processes are in use, an abundance of messages are seen about each process when the reference bandwidth is changed in a single process.	8.01.01
The wrong dead interval range was displayed in the help section of the cli for sham links.	8.11.01
'show running config' for the sham link authentication would not be displayed.	8.11.01
When looking at the debug syslog, sham-link interval mismatch messages do not decode ifindex to text strings.	8.11.01
The 'show ip ospf interface vlan.0.x' command may show additional space at the end if multiple addresses are configured on that interface that are not running OSPF.	8.01.01
"Running 'show support' on hardware that is not licensed for OSPF will result in the following error: Unknown: "ospf".	7.00.01
OSPF has no warning message when the calculated cost metric for an interface due to an auto-cost reference bandwidth change results in a too large metric.	8.11.01
OSPFv3 Problems Corrected in 8.21.02.0001	Introduced in Version:
If an OSPF vlan interface is configured to be POINT_TO_POINT, then the configuration is removed with "no ip ospf network point-to-point", the interface network type is POINT_TO_POINT instead of reverting to the default type of BROADCAST.	7.41.02
If an OSPF auto-cost reference bandwidth is configured that causes the interface to calculate a cost greater than the maximum, the cost remained based on the previous auto-cost reference bandwidth value.	8.11.01

PIM-SM Problems Corrected in 8.21.02.0001	Introduced in Version:
IGMP/Multicast in a bonded chassis appears to take longer for some events than an identically configured single chassis.	7.61.02
"ip pim multipath" configuration is not cleared after executing a "clear router vrf <vrfname>".	8.01.01
PIM configuration for ipv4 is accepted after removal of L3 license.	7.00.01
When using an access-list to define a group-list for use in PIM rp-candidate configuration, the number of groups must be no more than 255 or a system reset may occur.	7.40.00

PKI Problems Corrected in 8.21.02.0001	Introduced in Version:
When configuring an X.509 certificate via the "set pki certificate <pki-cert-list>" command a warning is displayed if the same certificate already exists on the list, and the user is prompted as to whether or not they want to accept the new certificate. The user can avoid this prompt (in order to avoid breaking automated scripts) by specifying the "no-confirm" option on the command line. The "no-confirm" option should suppress the duplicate certificate warning as well as suppressing the prompt.	8.11.01
If a configuration file which contains PKI data is modified by an external text editor and that editor adds control characters (such as '\r' 0x0D), then sourcing the modified config file may not restore very large certificates (on the order of 10K PEM characters, which is the maximum allowed by the device).	8.11.01

Platform Problems Corrected in 8.21.02.0001	Introduced in Version:
Reading a file from another blade (Ex: 'show file' or 'configure') could cause a DSI/reset, usually if the remote file is being updated, or remote connection goes away (other blade resets or bonding goes away).	7.00.01
Running "chkdsk repair" could cause a reset. This command is only available from debug, or during boot if filesystem corruption is detected.	7.00.01
Performing a configuration operation via the command line interface may result in the old configuration remaining due to file access errors.	7.70.01
1G port with 1G SFP plugged in incorrectly shows 100BASE advertised capabilities.	7.91.01
If a SFP+ Direct Attach cable assembly is used to connect two 1000Base-X ports (an unsupported configuration), the message "Incompatible pluggable module" will be logged on behalf of each port, but physical link will not be forced down.	7.00.01
SFP pluggable failure messages are not as user friendly as they should be.	8.01.01
Port advertisement settings are not persistent when auto negotiation is disabled.	7.91.01
An 100M SFP inserted into 1G port shows default speed and negotiation disable in 'show config'.	7.00.01
Core files might not be generated for defects which result in stack corruption whenever a DSI or ISI exception occurs, the system logs the original exception to NONVOL then attempts to generate a core file (i.e., /slot<x>/cores/<xxxx>.core.gz) which will include a stack trace of the offending task. If the stack is corrupted, then the process of printing the stack trace to the core file will itself create a new DSI. This new DSI prevents core file generation from completing and being saved to disk.	7.00.01
When a "clear linecard" command is executed for a K-Series line card on which 'macsource' and/or 'iptos' policy rules are configured, and these rules have associated COS drop-precedence values, an "Unable to remove flex-edge rule" message will be logged on behalf of each applicable port. These messages are innocuous and do not indicate an error condition.	7.30.01

Platform Problems Corrected in 8.21.02.0001	Introduced in Version:
At boot a board could get into reset loop with the following syslog output: 'NonVol[1.tusrApplnit]Nonvol reached max fileIdx 4080, storeNum 1, major 1'. Sets will be dropped until space is freed.	7.30.01
If a KK2008-0204 linecard has a port with an installed SFP, the port will not function properly after a linecard reset. Removing and reinserting the SFP will restore correct operation.	8.11.01
PoE Problems Corrected in 8.21.02.0001	Introduced in Version:
When 'show system hardware' is executed on system running PoE microcode update, the module might reset with message similar to: "<161>Mar 27 11:31:24 100.10.10.55 DistServ[1.tDsBrdOk]serverWatchDog.1(Config), client 85(sysHwDis) in rcv for 6581 tics (0x00c2dc94 0x005e8bf4 0x005d7ccc 0x014087a0 0xeeeeeeee)". PoE microcode update can be identified by message similar to: "<163>Sep 18 12:50:20 10.26.192.211 System[1]PoE controller 1 on slot 1 image is not up-to-date (exp:4.1.1.1, rcvd:4.1.0.2). Please do not power cycle the board or the PoE power supplies while image is being updated - downloading...".	7.00.01
When 'inlinepower psetrap' is enabled on more than one K linecard, it will not be possible to clear it.	7.30.01
When 'inlinepower threshold' is set on more than one K linecard, it will not be possible to clear it.	7.30.01
When 'set inlinepower psetrap' and/or 'set inlinepower threshold' is executed on multiple K line cards, only the last setting will be persistently stored.	7.30.01
Policy Problems Corrected in 8.21.02.0001	Introduced in Version:
Rules to drop GVRP or MVRP packets are ignored.	7.00.01
IP addresses in "set policy rule" would be treated as octal if a preceding "0" is present.	6.00.02
The CLI command 'show vlanauthorization' will not display the vlan authorization status of all ports in the system.	6.00.02
Multiauth failure traps may be output for port roaming sessions that roam to ports with insufficient per port multiauth number of users to support the new session.	7.72.01
PWA Problems Corrected in 8.21.02.0001	Introduced in Version:
PWA set portcontrol CLI commands do not output an error if wildcarding is used for a port string which contains no valid ports.	5.42.04
RADIUS Problems Corrected in 8.21.02.0001	Introduced in Version:
RADIUS authentication server max-sessions configuration is not output as part of "show config" or "show config all" commands.	8.11.01
RADIUS authentication server realm is not displayed as part of the "show config all" command if it is set to the default of any.	8.11.01
RADIUS-SNOOPING Problems Corrected in 8.21.02.0001	Introduced in Version:
If multiple CLI sessions are concurrently accessing RADIUS Snooping information, the system may crash or provide inaccurate results.	6.11.01

RADIUS-SNOOPING Problems Corrected in 8.21.02.0001	Introduced in Version:
Show config of the RADIUS Snooping, auto tracking, and quarantine provisioning agents displays default port parameters whenever at least one port field is set to a non-default setting.	6.11.01
RMON Problems Corrected in 8.21.02.0001	Introduced in Version:
Heavy use of RMON alarm and RMON event may result in a system reset and the log message "memPartFree: invalid block 0x3257c710 in partition 0x59a0a78 <memSysPartition>".	5.01.58
Configuration of RMON etherStats may return an incorrect value upon using an invalid index as input.	1.07.19
MIB leaf historyControlStatus can be set directly to under creation with non-existent index.	1.07.19
Configuration of an RMON function with an out of range index does not always return error.	5.01.58
"show rmon alarm" will show a negative value for alarm variables that are counters(unsigned), specifically for values between 2147483647 and 4294967294(rollover).	5.01.58
Routing Problems Corrected in 8.21.02.0001	Introduced in Version:
Negating interface checkspool setting without a keyword returns an error when checkspool loose-mode is configured.	7.00.01
SMON Problems Corrected in 8.21.02.0001	Introduced in Version:
Infrequently, a chassis module with port mirrors configured resets. On this failure a message similar to "setMirrorIndex(103201,122024): invalid mirror index transition 2->1" is logged.	7.00.01
SNMP Problems Corrected in 8.21.02.0001	Introduced in Version:
For SNMP view configuration, SNMP view mask values entered as single byte hexadecimal values (without a colon) that are less than 0x7f (and are printable ascii characters) appear as printable ascii characters instead of hexadecimal values, and result in missing configuration lines.	4.00.50
STP Problems Corrected in 8.21.02.0001	Introduced in Version:
Executing the CLI command "clear linecard <linecard #>" may cause a reset due to a failed port validation by spanning tree.	7.00.01
SYSLOG Problems Corrected in 8.21.02.0001	Introduced in Version:
Messages that should be logged to the console as part of the shutdown process are not seen.	7.80.01
Tracked Objects Problems Corrected in 8.21.02.0001	Introduced in Version:
Taking a tracked object out of service ('no inservice' sub-mode command) while a state change is in progress does not remove the state change action from the delay queue. If the tracked object is put back into service ('inservice' sub-mode command) prior to the state change action expiring from the delay queue, the new state change action is not queued. The new state change action is triggered when the old state change action expires from the delay queue.	7.60.01

Tunneling Problems Corrected in 8.21.02.0001	Introduced in Version:
When either a tunnel probe or the GRE keepalive is down, the tunnel is held down. This has been changed. If either the probe or keepalive is up or neither are configured, then the tunnel will be operationally up given other conditions are correct.	8.11.01
The range check on a tunnel keepalive period prevented the user from entering anything larger than 255.	8.11.01
A GRE keepalive nested within another GRE tunnel would be dropped.	8.11.01
HW connections may be incorrectly installed to drop virtual private port flows that include nested GRE packets with the protocol=0x6558	8.11.01
IPv6 encapsulated flow of an IPv4 flow was using the IP version from the Transformation. It now uses the IP Version from the Ingress Flow.	8.01.01

VRF Problems Corrected in 8.21.02.0001	Introduced in Version:
When using the maximum length VRF name, it insists on a context, but when one is specified, it takes the VRF name and discards the extra characters.	7.62.02
From device command line, a ping to device's address configured in another VRF fails even though VRF route leaking is provided by static routes.	8.11.01

VRRP Problems Corrected in 8.21.02.0001	Introduced in Version:
Host routes added by host mobility may age out during first age pass after they are added.	8.11.01

Feature Enhancements in 8.11.05.0006

Transceiver Enhancements in 8.11.05.0006
<p>CWDM support:</p> <p>10GB-LR271-SFPP - 10Gb, CWDM SM, 1271 nm, 10 km, LC SFP+</p> <p>10GB-LR291-SFPP - 10Gb, CWDM SM, 1291 nm, 10 km, LC SFP+</p> <p>10GB-LR311-SFPP - 10Gb, CWDM SM, 1311 nm, 10 km, LC SFP+</p> <p>10GB-LR331-SFPP - 10Gb, CWDM SM, 1331 nm, 10 km, LC SFP+</p>
<p>Additional DWDM support:</p> <p>10GB-ER21-SFPP - 10GB-ER, DWDM CH21 SFP+</p> <p>10GB-ER24-SFPP - 10GB-ER, DWDM CH24 SFP+</p> <p>10GB-ER31-SFPP - 10GB-ER, DWDM CH31 SFP+</p> <p>10GB-ER33-SFPP - 10GB-ER, DWDM CH33 SFP+</p>

Problems Corrected in 8.11.05.0006

ACLs Problems Corrected in 8.11.05.0006	Introduced in Version:
When the platform connection look-up level has been raised from L3 to L4 by application of an ACL, removing the ACL does not cause the look-up level to be reduced to L3.	7.40.01
When adding entries to an access-list, duplicates of existing entries are no longer accepted.	7.00.01
IGMP Problems Corrected in 8.11.05.0006	Introduced in Version:
When issuing a "show config" and reaching the MLD section, the config may get stuck in a loop and not allow the config to finish displaying.	7.30.01

IPv6 Neighbor Discovery Problems Corrected in 8.11.05.0006	Introduced in Version:
When inserting a new blade into the system, the new blade may end up with an interface in the "stalled" state which indicates that the IPv6 addresses have not passed Duplicate Address Detection. The interface will not forward IPv6 packets until the interface is bounced (the operational status goes down then back up).	7.41.02
LLDP Problems Corrected in 8.11.05.0006	Introduced in Version:
Every time the command "show config" or "show config all" is run, the system loses as much as 512Kb of memory. Enough memory losses eventually cause the system to reset.	8.11.01
Multiauth Problems Corrected in 8.11.05.0006	Introduced in Version:
Modification or removal of multi-authentication users may cause prolonged high CPU utilization and dropped traffic.	7.00.01
Multicast Problems Corrected in 8.11.05.0006	Introduced in Version:
Upper slots were limited by 16 bit mask. This has been updated to a 32 bit mask.	7.30.01
NETFLOW Problems Corrected in 8.11.05.0006	Introduced in Version:
If Netflow higher-layer export is enabled and the cache is disabled at a time when flows are actively being exported, and then later re-enabled, messages similar to: "PiMgr[3.tDispatch]generateIIndex() :retval=0;owner(3);mediaType(7);mediaPos(0)" may be generated. For each message generated, a single Netflow record with invalid data will be exported.	8.01.01
Node Alias Problems Corrected in 8.11.05.0006	Introduced in Version:
Under rare circumstances, the "ctAliasControlTable" will not return all valid entries.	7.91.01
If the switch is receiving MDNS or LLNMR or SSDP frames and Node, and Alias is not configured to have those protocols disabled (nor configured to have ports those frames are being received on disabled) and, in addition, one of the following is true: - Is also receiving IP Fragment packets - Receives at least one malformed MDNS, LLNMR, or SSDP frame One or more blades may get into a state where CPU usage is 100%. When in this state the "Switch Node & Alias" process will be shown as taking significant CPU for a "show system utilization". This will not affect packet forwarding or L2/L3 protocols, but will adversely affect all management. The only recovery method is to reset the individual blades that get into this state.	8.11.01
OSPF Problems Corrected in 8.11.05.0006	Introduced in Version:
An assertion failure and reset occurs and is recorded in message log as; "SMS assert in qoamlsts.c at line 1218"	7.00.01
When running OSPFv2 and flapping the passive value on an interface, an assert can occur in thread tRtrPtcls with the following message; "SMS assert in qopmmim5.c at line 879 : (null) AVLL_IN_TREE(if_cb->active_if_tree_node) 0 (null) 0 "	8.11.01

OSPF Problems Corrected in 8.11.05.0006	Introduced in Version:
When running OSPF a DSI can occur in thread tRtrPtcls, message displayed is: "SMS assert in ntlavll.c at line 644 : != AVL3_IN_TREE(*node) 0 0 0"	8.11.01
Platform Problems Corrected in 8.11.05.0006	Introduced in Version:
When a linecard of type KK2008-0204-F2 is taken offline, removed from its slot, and replaced with a linecard of a different type, the system will report a fatal error and restart.	8.11.01
It is possible for a 10G port to be incorrectly assigned 1G of bandwidth.	8.02.01
PWA Problems Corrected in 8.11.05.0006	Introduced in Version:
PWA is occasionally unable to respond to HTTP requests under heavy user login load. Related syslog message: "PWA[2.tLwipRecv]pwaTransmitPkt() transmit failed"	7.00.01
Spanning Tree Problems Corrected in 8.11.05.0006	Introduced in Version:
Reset could occur when (1) changing spantree operational mode between "ieee" and "none" or (2) when spantree version is "stpcompatible" and entering or leaving a topology change condition.	7.00.01
Switching Problems Corrected in 8.11.05.0006	Introduced in Version:
Precision Time Protocol (PTPv1) UDP broadcast port 139, when being forwarded through switch, may not function reliably.	1.07.19

Feature Enhancements in 8.11.04.0005

Transceiver Enhancements in 8.11.04.0005
Auto negotiation support for 1Gb SFP GBICs installed in SFP+ sockets.

Problems Corrected in 8.11.04.0005

CLI Problems Corrected in 8.11.04.0005	Introduced in Version:
Unable to read port type after line card is reset via "reset <slot>". When in this state "show port status" will display "unknown" type for affected ports.	8.11.01
Login banner configured via "set banner login <message>" is not displayed when logging in via SSH. The banner is displayed when logging in via Console or TELNET.	8.11.01
IGMP Problems Corrected in 8.11.04.0005	Introduced in Version:
The IGMP database can become corrupted leading to unpredictable multicast results and/or module crashes.	7.30.01
When using IGMP unknown-input-action setting "Flood To Routers", IGMP may not route these packets properly.	8.11.01
"IGMP may on board synchronization, or system reset, reset with the following message: IGMP[3.tDSSync2]ClgmpEtsc::DistGrpTbIRecvDistributedAdd Recv base index out of range baseidx:xxx flowldx:xxx"	8.11.01

NODE-ALIAS Problems Corrected in 8.11.04.0005	Introduced in Version:
Querying the ctAliasInterface table may not return all entries on a given interface.	8.11.01
NONVOL Problems Corrected in 8.11.04.0005	Introduced in Version:
The nonvol cleanup task can write incomplete files to the nonvol store that will not be detected until a reboot or the next time cleanup runs for that store and component: <3>NonVol[8.tNVolCUp]nvFilePtrMgr::verify(3) calcCsum() failed. store=5, fileIdx=10.51, udpSum=0x77e366a, sumCount=65534	3.00.33
The nonvol cleanup task can write incomplete files to the nonvol store that will not be detected until a reboot: NonVol[1.tusrAppInIt]nvFilePtrMgr::verify(0) checksum failure. store=4, fileIdx=0.37, udpSum=0x8f8dd5a, sumCount=65527	3.00.33
The nonvol cleanup task can cause a DSI reset: Exc Vector: DSI exception (0x00000300) Thread Name: tNVolCUp	3.00.33
The nonvol cleanup task can become stuck causing high system utilization: debug utilization show -i NAME TID PRI STATUS 5sec 1min 5min Got tid = 1 from successful call to getNextTaskId(). tNVolCUp 240412704 195 READY 99.37 99.28 99.27	3.00.33
PLATFORM Problems Corrected in 8.11.04.0005	Introduced in Version:
System logs the message: "bcmStrat[11.tFabDevInt]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000001, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000" and resets.	7.70.01
PoE Problems Corrected in 8.11.04.0005	Introduced in Version:
'set inlinepower management class' configuration might not be persistent.	8.01.01
RADIUS Problems Corrected in 8.11.04.0005	Introduced in Version:
RADIUS authentication servers created via SNMP without the etsysRadiusAuthClientServerStickyMaxSessions leaf present will default to a maximum sessions value of 0. This will effectively cause the sticky-round-robin RADIUS algorithm to work like the round-robin RADIUS algorithm.	8.11.01
SSH Problems Corrected in 8.11.04.0005	Introduced in Version:
"The SSH configuration parameter 'set ssh server allowed-auth password {enabled disabled}' was added in release 8.11. The default value for this new parameter should be 'enabled'. However, if upgrading from a pre-8.11 image to 8.11 the parameter may initialize as 'disabled'. This will prevent users from connecting to the device using SSH.	8.11.01
TACACS+ Problems Corrected in 8.11.04.0005	Introduced in Version:
If no attributes are passed back in an authorized TACACS+ response when performing TACACS+ command authorization, results may be non-deterministic resulting in some commands being authorized and others not. TACACS+ commands which fail authorization will correctly not be allowed.	6.11.01

VRRP Problems Corrected in 8.11.04.0005	Introduced in Version:
If IPv6 hosts are connected to a switch which is connected to a VRRP master and VRRP backup router is running host-mobility, the IPv6 hosts will periodically move from master to backup and back again to the master due to router advertisement being sent by backup using VRRP virtual MAC address.	8.11.01
Master VRRP router does not reply to ARP requests sent for the VIP's IP when fabric-router mode is enabled.	8.11.01

Feature Enhancements in 8.11.03.0005

Application Policy Feature Enhancement in 8.11.03.0005
A new Policy Classification rule type allows for control of additional application specific traffic. The Application Policy feature provides differentiation between requests and queries/announcements for common ZeroConf protocols to allow a simple granular policy assignment. These protocols include Apples Bonjour and Universal Plug and Play (UPnP).
Fabric Routing with IP Host Mobility Feature Enhancement in 8.11.03.0005
IP Host Mobility allows for optimized North/South traffic when deployed in a common route fabric environment. IP Host Mobility leverages host routing.
Isolated Private VLAN Feature Enhancement in 8.11.03.0005
This feature adds the ability for a secondary VLAN to share an IP interface assigned to a primary VLAN. Users within the secondary VLAN can be isolated from each other such that communication must flow through the router.
Tunneling, 'Virtual Private Port Service' Feature Enhancement in 8.11.03.0005
Layer 2 interconnect via GRE tunnel interface, allows for the encapsulation of all data entering a specified port for transport across the network infrastructure with a routable IP/GRE tunnel. (Licensed Feature)
Inter-VRF Access Control List Feature Enhancement in 8.11.03.0005
This feature adds Access Control List functionality for internal data traffic routed between multiple VRF instances running in the same device.
RADIUS / Policy Enhancements Feature Enhancements in 8.11.03.0005
Server Load Balancing – Adds support for RADIUS authentication server load balancing. Authentication Timeout Policy – Allows for the application of a specific RADIUS timeout policy profile to be applied during authentication timeout events. Authentication Failure Policy - Allows for the application of a specific RADIUS failure policy profile to be applied during authentication failure events. Re-Authentication Timeout Enhancement – Enhancement to allow for the use of the previous access level during a re-authentication timeout event. Accounting Enhancement – Accounting has been extended to allow for accounting of additional provisioning agents that previously were unaccounted. Including CEP, RADIUS snooping, AutoTracking and Quarantine.
SSH Public Key Authentication Feature Enhancement in 8.11.03.0005
SSH enhancement to support Public Key Authentication as an additional client authentication method.
RMON Stats and History Feature Enhancement in 8.11.03.0005
Enhancement to the operation of RMON EtherStats and History, allowing for the configuration of the direction of statistics collection; TX, RX or TX+RX.

Automated Deployment Feature Enhancement in 8.11.03.0005

This feature allows a newly installed device with no configuration (default configuration), to obtain the latest firmware revision and/or configuration automatically from the network. Leveraging DHCP, the device will obtain a temporary IP address and notify NetSight of its status on the network allowing NetSight to provide the specified changes to the device. A SNMP trap requesting configuration is sent to the SNMP server notifying it that the system is ready to be configured.

MAC Authentication Feature Enhancement in 8.11.03.0005

Allows the MAC Authentication password to use the configured password or the username as password.

IPv6 DHCP Server Feature Enhancement in 8.11.03.0005

DHCPv6 server support has been added. The DHCPv6 server can be used to configure DHCPv6 clients with IPv6 addresses, IP prefixes and other configuration required to operate in an IPv6 network.

Power over Ethernet LLDP advertisement update Feature Enhancement in 8.11.03.0005

IEEE amendment 802.3at-2009 update to "power via MDI" TLV is supported. This update includes three new fields: type/source/priority, PD requested power and PSE allocated power.

OSPF Reference Bandwidth Feature Enhancement in 8.11.03.0005

Enhancement to support configuring OSPF reference bandwidth, allowing for more granular auto-costing of OSPF links.

Neighbor Discovery Enhancement Feature Enhancement in 8.11.03.0005

Enhancement to detect and display configuration mismatches, duplex mode and speed settings, between endpoints using the various neighbor discovery methods.

Problems Corrected in 8.11.03.0005

Antispoofing Problems Corrected in 8.11.03.0005	Introduced in Version:
Issuing the CLI command "show antispoof binding" will result in a small amount of memory being leaked.	8.01.01
DHCP Problems Corrected in 8.11.03.0005	Introduced in Version:
'ipv6 dhcp relay source-interface' disappears when the master blade is reset in a chassis.	7.30.01
Flow Limiting Problems Corrected in 8.11.03.0005	Introduced in Version:
When flow limiting is enabled on a port, the flow event counter for that port will not be accurate.	8.01.01
Host Problems Corrected in 8.11.03.0005	Introduced in Version:
After issuing the traceroute command, the string "runTraceroute: ifindex <number>" is displayed before the results.	7.99.00
LLDP Problems Corrected in 8.11.03.0005	Introduced in Version:
Occasionally running the show neighbor command will display a neighbor multiple times.	7.91.01

Multi User Authentication Problems Corrected in 8.11.03.0005	Introduced in Version:
Executing the CLI command show multiauth session port <port-string>" might result in an error.	7.00.01
Neighbor Discovery Problems Corrected in 8.11.03.0005	Introduced in Version:
CLI output for the "show neighbors" command will infrequently exclude one or more neighbors from one or more modules.	7.31.02
OSPF Problems Corrected in 8.11.03.0005	Introduced in Version:
If OSPF passive interfaces are configured, upgrading from any 7.X release to an 8.x release could cause a DSI in thread tDsync5.	8.01.01
OSPFv3 Problems Corrected in 8.11.03.0005	Introduced in Version:
If an OSPFv3 interface is configured as passive under IPv6 router OSPF before it is enabled under the interface, and other OSPFv3 interface attributes had been applied, the passive interface would remain down.	8.01.01
Platform Problems Corrected in 8.11.03.0005	Introduced in Version:
"System logs the message ""bcmStrat[2.tNimIntr]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000010, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000"" and resets.	7.70.01
System logs the message "bcmStrat[5.tNimIntr]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000080, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000" and resets.	7.70.01
SCP Problems Corrected in 8.11.03.0005	Introduced in Version:
Secure Copy (scp) file transfers do not work. (i.e., "copy scp://<user>@<host>://<path>/<source-file> slot1/<destination-file>").	7.62.05

Feature Enhancements in 8.02.01.0012

HW Feature Enhancements in 8.02.01.0012
Support has been added for an 80Km SFP+ transceiver; 10GB-ZR-SFPP - 10 Gb, 10GBASE-ZR, SM, 1550 nm, 80 Km, LC SFP+
Support has been added for 100Mb copper SFP transceiver; MGBIC-100BT - 100 Mb, 100BASE-T Copper twisted pair, 100 m, RJ45 SFP

Feature Enhancements in 8.02.01.0012

IP Service Level Agreements Feature Enhancements in 8.02.01.0012
This feature (IPSLA) adds the ability to perform scheduled packet timing statistics gathering and analysis at the IP layer. This feature also adds round trip time measurements for network paths on a per hop basis.

Tracked Objects Feature Enhancements in 8.02.01.0012

Enhancement to existing feature to allow monitoring and actions on local physical interfaces. This feature also adds the ability to provide packet timing measurements for use with IPSLA feature.

User Tracking and Control Feature Enhancements in 8.02.01.0012

Additional features for tracking and control of user sessions. These features are leveraged by the Anti-Spoofing Suite.

Auto-Tracking – This feature tracks non-authenticated sessions to allow for visibility and policy control. Non-authenticated sessions were previously not tracked in the session table.

Quarantine agent – This feature provides the ability to provision sessions based on both their policy profile and the type of traffic they are sending. Policy rules will allow for a quarantine action which will allow for a quarantine policy profile to be defined that can trigger when traffic matches the traffic filter specification in the rule. The Anti-Spoofing suite will leverage this feature.

Anti-Spoofing Suite Feature Enhancements in 8.02.01.0012

A set of features to provide secure IP spoofing detection and prevention to the network dynamically through the use of a source MAC/IP binding table.

DHCP Snooping – tracks DHCP messaging and builds a binding table to enforce DHCP client/server access from specific locations in the network.

Dynamic Arp Inspection- utilizes the MAC to IP binding table to ensure that ARP packets have the proper MAC to IP binding

IP source guard –utilizes the MAC to IP binding table to limit/enforce a user's specific MAC and IP address access to the network.

DHCP Feature Enhancements in 8.02.01.0012

Relay Option 82 – The DHCP relay option 82 feature has been enhanced to allow circuit-ID (VLAN-ID) and Remote-ID (Chassis MAC) fields to be populated by default when relaying DHCP packets. Each of these fields can be manually overwritten with ASCII text.

Lease Capacity enhancement - The DHCP server lease capacity has been increased from 1,024 to 5,000.

Port Mirror Feature Enhancements in 8.02.01.0012 (K-EOS-L3 License is required)

N Packet 'Forensic' Port Mirror – This feature adds the ability to allow a specific flow to have a specified number of packets mirrored. The first "N" packets and only first N packets are mirrored.

Remote Port Mirror – The feature provides the ability to send port mirror traffic to a remote destination across the IP network. Traffic is encapsulated in a L2 GRE tunnel and can be routed across the network.

Multicast Feature Enhancements in 8.02.01.0012

PIM Graceful –This feature allows PIM sparse mode to continue to forward existing multicast streams during a graceful restart. This feature will also allow updates to occur during the restart but will not forward new streams until after the restart is complete.

PIM Multipath - This feature provides the ability to define the mechanism by which PIM chooses the next-hop for choosing the “reverse path” to a source. The user can optionally choose to use the highest next-hop, or use a SourceIP hash to choose a next-hop based on a hash of the source IP address. The feature allows PIM multicast load sharing over ECMP paths, as well as the ability to have a single deterministic next-hop for ECMP paths.

Multicast domains – This feature allows a PIM router to be a Border Router, as well as support MSDP (Multicast Source Discovery Protocol). MSDP interconnects multiple PIM sparse mode domains enabling PIM-SM to have Rendezvous Point (RP) redundancy where multicast sources can be known across domains allowing for inter-domain multicasting.

Multi-topology Multicast -This feature provides the ability to create a separate topology for use by PIM in routing multicast traffic. Routing protocols OSPF, OSPFv3 and IS-IS may be configured to support this separate multicast topology in an effort to contain multicast to a subset of the Enterprise.

IGMP input filters -This feature allows the user to configure input filters for a range of incoming multicast packets. The input filters provide the ability to define actions to allow, drop, or flood the protocol packets as well as the flow.

VLAN Provider Bridging (Q-in-Q) Feature Enhancements in 8.02.01.0012

This feature adds support for adding a second VLAN tag (S-tag) for transport of multiple customer VLANs across a common service provider infrastructure. The addition of the S-tag allows customer VLANs to be transported intact transparently across a layer 2 infrastructure.

MVRP - IEEE 802.1ak Feature Enhancements in 8.02.01.0012

Multiple VLAN Registration Protocol (MVRP) is the standardized replacement protocol for GVRP (GARP VLAN Registration Protocol), used to dynamically configure and distribute VLAN membership information throughout a network.

CFM - IEEE 802.1Q-2011 Feature Enhancements in 8.02.01.0012

Connectivity Fault Management (CFM) provides network operators a way to effectively monitor and troubleshoot services that may span single or multiple domain Ethernet networks. CFM supports mechanisms and diagnostics to insure devices along the path are configured properly, validate reachability and pinpoint connectivity loss.

Unidirectional Link Detection Feature Enhancements in 8.02.01.0012

This feature provides the ability to detect a single direction link where the ability to pass traffic over the link is not functioning in one direction. The feature also enables the ability to take a port out of service when a unidirectional link is detected through the use of Link Layer OAM.

Host Denial of Service ARP/ND Feature Enhancements in 8.02.01.0012

This enhancement, as part of the Host DOS feature, protects the CPU from receiving excessive Address Resolution Protocol (ARP) or Neighbor Discovery (ND) packets from the same host.

IPv6 Neighbor Discovery Feature Enhancements in 8.02.01.0012

Support for RFC 4191 and 6106 have been added to this release. RFC 4191 provides default router preferences and specific route priority information to IPv6 hosts through router advertisements via neighbor discovery. RFC 6106 provides options for distributing DNS server and suffix information to IPv6 hosts through router advertisements via neighbor discovery.

IPv6 Route table Capacity Feature Enhancements in 8.02.01.0012

The IPv6 route table capacity has been increased to 25,000 routes for the K-Series.

SSH Feature Enhancements in 8.02.01.0012

SSH CLI now supports configuration of keep alive count and interval. This may be used to reduce likelihood that ssh clients like 'putty' will cause a disconnect when they fail to maintain keep alive protocol. (Due to a bug in putty this protocol is not run while holding the putty scroll bar down or accessing the putty configuration screens.)

Problems Corrected in 8.02.01.0012

ARP Problems Corrected in 8.02.01.0012	Introduced in Version:
<p>When sending an ARP request to an interface address that exists on an interface other than the interface that received the ARP (proxy ARP), the MAC address of the interface that contains the destination IP address will be used in the ARP response instead of the MAC address of the interface that received the ARP request.</p> <p><i>For example:</i> If interface vlan.0.11 contains IP address 11.0.0.1/8 AND interface vlan.0.12 contains IP address 12.0.0.1/8 AND proxy ARP is enabled on interface vlan.0.11 AND interface vlan.0.11 receives an ARP request for IP address 12.0.0.1 THEN the ARP response will contain the MAC address of vlan.0.12 instead of vlan.0.11</p>	7.00.01
Hardware Problems Corrected in 8.02.01.0012	Introduced in Version:
Faulty I2C device may cause I2C access failures to other devices in the system.	7.00.01
HOSDOS Problems Corrected in 8.02.01.0012	Introduced in Version:
Default rate settings for hostDos threats icmpFlood and synFlood may disrupt protocol operation and/or further configuration of the device.	7.20.01
LLDP Problems Corrected in 8.02.01.0012	Introduced in Version:
The SNMP MIB IldpStatsRxPortAgeoutsTotal does not return the correct value.	5.42.xx
MTU Problems Corrected in 8.02.01.0012	Introduced in Version:
IP interfaces can exist with a Max Transit Unit (MTU) set to 0.	Unknown
OSPF Problems Corrected in 8.02.01.0012	Introduced in Version:
FIB may not be properly populated if routers with route entries pointing to loopback interfaces advertised by adjacent neighbors and virtual-link are being used, or the router across the virtual-link injects quite a few type-5 LSAs.	7.20.01
OSPF will reset and log a "SMS assert in qodmnssa.c" when user adds and all zeros NSSA route	7.00.01

Platform Problems Corrected in 8.02.01.0012	Introduced in Version:
Some types of failures in memory systems used by Switching ASICs lead to resets of chassis rather than shutdown of the line card that the Switching ASIC is on.	7.40.00
System may reset with Stats DMA error message. System should not reset when this condition occurs.	7.80.01
Policy Problems Corrected in 8.02.01.0012	Introduced in Version:
Some policy configuration may be missing after a reboot.	7.00.01
STP Problems Corrected in 8.02.01.0012	Introduced in Version:
Reset could occur when (1) changing spantree operational mode between "ieee" and "none" or (2) when spantree version is "stpcompatible" and entering or leaving a topology change condition.	7.00.01
SYSLOG Problems Corrected in 8.02.01.0012	Introduced in Version:
Messages sent to syslog servers could contain unprintable control characters in the middle of the messages.	7.11.01
VLAN Problems Corrected in 8.02.01.0012	Introduced in Version:
A VLAN interface based mirror will continue to mirror traffic after the VLAN interface is removed from the config with the clear command.	1.07.19
VRF Problems Corrected in 8.02.01.0012	Introduced in Version:
When doing a fail over, then a show running config, some limit commands will show up even though they were not set.	7.70.01

KNOWN RESTRICTIONS AND LIMITATION:

VSB
When using VSB several features are resized or restricted: IP tunnels including VXLAN and GRE Tunnels are not supported, (Remote Port Mirrors are supported) Port Mirroring: - IDS mirror is not supported - Frames can be the subject of one mirror only - The 10GB-ER-SFP (10 Gb, 10GBASE-ER, IEEE 802.3 SM, 1550 nm Long Wave Length, 40 Km, LC SFP+) is not supported as a VSB chassis interconnect.
RADIUS
Radius load balancing must be configured with CLI. When configured by Netsight using SNMP the system will not load balance. Scheduled to be fixed in a future release.
SFP+ Port Speed
When an SFP (1G) module is inserted or removed from an SFP+ (10G capable) port, all ports on the associated MAC chip are reset. This results in a momentary loss of link and traffic on affected ports and forces topology protocols to process a link bounce.

100Mb SFP

MGBIC-100BT doesn't support automatic detection of MDIX (Medium Dependent Interface Crossover) or Auto-negotiation.

Any problems other than those listed above should be reported to our Technical Support Staff.

IEFT STANDARDS SUPPORT:

RFC No.	Title
RFC0147	Definition of a socket
RFC0768	UDP
RFC0781	Specification of (IP) timestamp option
RFC0783	TFTP
RFC0791	Internet Protocol
RFC0792	ICMP
RFC0793	TCP
RFC0826	ARP
RFC0854	Telnet
RFC0894	Transmission of IP over Ethernet Networks
RFC0919	Broadcasting Internet Datagrams
RFC0922	Broadcasting IP datagrams over subnets
RFC0925	Multi-LAN Address Resolution
RFC0950	Internet Standard Subnetting Procedure
RFC0951	BOOTP
RFC0959	File Transfer Protocol
RFC1027	Proxy ARP
RFC1034	Domain Names - Concepts and Facilities
RFC1035	Domain Names - Implementation and Specification
RFC1071	Computing the Internet checksum
RFC1112	Host extensions for IP multicasting
RFC1122	Requirements for IP Hosts - Comm Layers
RFC1123	Requirements for IP Hosts - Application and Support
RFC1157	Simple Network Management Protocol
RFC1191	Path MTU discovery
RFC1195	Use of OSI IS-IS for Routing in TCP/IP
RFC1213	MIB-II
RFC1245	OSPF Protocol Analysis
RFC1246	Experience with the OSPF Protocol
RFC1323	TCP Extensions for High Performance
RFC1349	Type of Service in the Internet Protocol Suite
RFC1350	TFTP
RFC1387	RIPv2 Protocol Analysis
RFC1388	RIPv2 Carrying Additional Information
RFC1389	RIPv2 MIB Extension
RFC1492	TACAS+
RFC1493	BRIDGE- MIB
RFC1517	Implementation of CIDR

RFC No.	Title
RFC1518	CIDR Architecture
RFC1519	Classless Inter-Domain Routing (CIDR)
RFC1542	BootP: Clarifications and Extensions
RFC1624	IP Checksum via Incremental Update
RFC1659	RS-232-MIB
RFC1721	RIPv2 Protocol Analysis
RFC1722	RIPv2 Protocol Applicability Statement
RFC1723	RIPv2 with Equal Cost Multipath Load Balancing
RFC1724	RIPv2 MIB Extension
RFC1812	General Routing
RFC1850	OSPFv2 MIB
RFC1853	IP in IP Tunneling
RFC1886	DNS Extensions to support IP version 6
RFC1924	A Compact Representation of IPv6 Addresses
RFC1981	Path MTU Discovery for IPv6
RFC2001	TCP Slow Start
RFC2003	IP in IP Tunneling
RFC2012	TCP-MIB
RFC2013	UDP-MIB
RFC2018	TCP Selective Acknowledgment Options
RFC2030	SNTP
RFC2080	RIPng (IPv6 extensions)
RFC2082	RIP-II MD5 Authentication
RFC2096	IP Forwarding Table MIB
RFC2104	HMAC
RFC2113	IP Router Alert Option
RFC2117	PIM -SM Protocol Specification
RFC2131	Dynamic Host Configuration Protocol
RFC2132	DHCP Options and BOOTP Vendor Extensions
RFC2233	The Interfaces Group MIB using SMIv2
RFC2236	Internet Group Management Protocol, Version 2
RFC2328	OSPFv2
RFC2329	OSPF Standardization Report
RFC2338	VRRP
RFC2362	PIM-SM Protocol Specification
RFC2370	The OSPF Opaque LSA Option
RFC2373	RFC 2373 Address notation compression
RFC2374	IPv6 Aggregatable Global Unicast Address Format
RFC2375	IPv6 Multicast Address Assignments
RFC2401	Security Architecture for the Internet Protocol
RFC2404	The Use of HMAC-SHA-1-96 within ESP and AH
RFC2406	IP Encapsulating Security Payload (ESP)
RFC2407	The Internet IP Security Domain of Interpretation for ISAKMP
RFC2408	Internet Security Association and Key Management Protocol (ISAKMP)
RFC2409	The Internet Key Exchange (IKE)

RFC No.	Title
RFC2428	FTP Extensions for IPv6 and NATs
RFC2450	Proposed TLA and NLA Assignment Rule
RFC2453	RIPv2
RFC2460	IPv6 Specification
RFC2461	Neighbor Discovery for IPv6
RFC2462	IPv6 Stateless Address Auto configuration
RFC2463	ICMPv6
RFC2464	Transmission of IPv6 over Ethernet
RFC2473	Generic Packet Tunneling in IPv6 Specification
RFC2474	Definition of DS Field in the IPv4/v6 Headers
RFC2475	An Architecture for Differentiated Service
RFC2548	Microsoft Vendor-specific RADIUS Attributes
RFC2553	Basic Socket Interface Extensions for IPv6
RFC2577	FTP Security Considerations
RFC2578	SNMPv2-SMI
RFC2579	SNMPv2-TC
RFC2581	TCP Congestion Control
RFC2597	Assured Forwarding PHB Group
RFC2613	SMON-MIB
RFC2674	P/Q-BRIDGE- MIB
RFC2685	Virtual Private Networks Identifier
RFC2697	A Single Rate Three Color Marker
RFC2710	Multicast Listener Discovery (MLD) for IPv6
RFC2711	IPv6 Router Alert Option
RFC2715	Interop Rules for MCAST Routing Protocols
RFC2740	OSPF for IPv6
RFC2763	Dynamic Hostname Exchange Mechanism for IS-IS
RFC2784	GRE
RFC2787	VRRP MIB
RFC2819	RMON MIB
RFC2827	Network Ingress Filtering
RFC2863	IF-MIB
RFC2864	IF-INVERTED-STACK-MIB
RFC2865	RADIUS Authentication
RFC2866	RADIUS Accounting
RFC2869	RADIUS Extensions
RFC2890	Key and Sequence Number Extensions to GRE
RFC2893	Transition Mechanisms for IPv6 Hosts and Routers
RFC2894	RFC 2894 Router Renumbering
RFC2922	PTOPO-MIB
RFC2934	PIM MIB for IPv4
RFC2966	Prefix Distribution with Two-Level IS-IS
RFC2973	IS-IS Mesh Groups
RFC2991	Multipath Issues in Ucast & Mcast Next-Hop
RFC3056	Connection of IPv6 Domains via IPv4 Clouds

RFC No.	Title
RFC3101	The OSPF Not-So-Stubby Area (NSSA) Option
RFC3137	OSPF Stub Router Advertisement
RFC3162	RADIUS and IPv6
RFC3273	HC-RMON-MIB
RFC3291	INET-ADDRESS-MIB
RFC3315	DHCPv6
RFC3359	TLV Code points in IS-IS
RFC3373	Three-Way Handshake for IS-IS
RFC3376	Internet Group Management Protocol, Version 3
RFC3411	SNMP Architecture for Management Frameworks
RFC3412	Message Processing and Dispatching for SNMP
RFC3412	SNMP-MPD-MIB
RFC3413	SNMP Applications
RFC3413	SNMP-NOTIFICATIONS-MIB
RFC3413	SNMP-PROXY-MIB
RFC3413	SNMP-TARGET-MIB
RFC3414	SNMP-USER-BASED-SM-MIB
RFC3415	SNMP-VIEW-BASED-ACM-MIB
RFC3417	SNMPv2-TM
RFC3418	SNMPv2 MIB
RFC3446	Anycast RP mechanism using PIM and MSDP
RFC3484	Default Address Selection for IPv6
RFC3493	Basic Socket Interface Extensions for IPv6
RFC3509	Alternative Implementations of OSPF ABRs
RFC3513	RFC 3513 IPv6 Addressing Architecture
RFC3542	Advanced Sockets API for IPv6
RFC3567	IS-IS Cryptographic Authentication
RFC3579	RADIUS Support For Extensible Authentication Protocol (EAP)
RFC3584	SNMP-COMMUNITY-MIB
RFC3587	IPv6 Global Unicast Address Format
RFC3590	RFC 3590 MLD Multicast Listener Discovery
RFC3595	Textual Conventions for IPv6 Flow Label
RFC3596	DNS Extensions to Support IP Version 6
RFC3618	Multicast Source Discovery Protocol (MSDP)
RFC3621	POWER-ETHERNET-MIB
RFC3623	Graceful OSPF Restart
RFC3630	Traffic Engineering (TE) Extensions to OSPFv2
RFC3635	ETHERLIKE-MIB
RFC3678	Socket Interface Ext for Mcast Source Filters
RFC3704	Network Ingress Filtering
RFC3719	Recommendations for Interop Networks using IS-IS
RFC3768	VRRP
RFC3769	Requirements for IPv6 Prefix Delegation
RFC3787	Recommendations for Interop IS-IS IP Networks
RFC3809	Requirements for Provider Provisioned VPNs
RFC3810	MLDv2 for IPv6

RFC No.	Title
RFC3847	Restart signaling for IS-IS
RFC3879	Deprecating Site Local Addresses
RFC3956	Embedding the RP Address in IPv6 MCAST Address
RFC4007	IPv6 Scoped Address Architecture
RFC4022	MIB for the Transmission Control Protocol (TCP)
RFC4087	IP Tunnel MIB
RFC4109	Algorithms for IKEv1
RFC4113	MIB for the User Datagram Protocol (UDP)
RFC4133	ENTITY MIB
RFC4167	Graceful OSPF Restart Implementation Report
RFC4188	Bridge MIB
RFC4191	Default Router Prefs and More-Specific Routes
RFC4193	Unique Local IPv6 Unicast Addresses
RFC4213	Basic Transition Mechanisms for IPv6
RFC4222	Prioritized Treatment of OSPFv2 Packets
RFC4250	The Secure Shell (SSH) Protocol Assigned Numbers
RFC4251	The Secure Shell (SSH) Protocol Architecture
RFC4252	The Secure Shell (SSH) Authentication Protocol
RFC4253	The Secure Shell (SSH) Transport Layer Protocol (no support diffie-hellman-group14-sha1)
RFC4254	The Secure Shell (SSH) Connection Protocol
RFC4256	Generic Message Exchange Authentication for the Secure Shell Protocol (SSH)
RFC4265	Definition of Textual Conventions for VPN Mgt
RFC4268	ENTITY-STATE-MIB
RFC4268	ENTITY-STATE-TC-MIB
RFC4291	IP Version 6 Addressing Architecture
RFC4293	MIB for the Internet Protocol (IP)
RFC4294	IPv6 Node Requirements
RFC4295	Mobile IP Management MIB
RFC4301	Security Architecture for IP
RFC4302	IP Authentication Header
RFC4303	IP Encapsulating Security Payload (ESP)
RFC4305	Crypto Algorithm Requirements for ESP and AH
RFC4306	Internet Key Exchange (IKEv2) Protocol
RFC4307	Cryptographic Algorithms for Use in IKEv2
RFC4308	Cryptographic Suites for IPsec
RFC4419	Diffie-Hellman Group Exchange for the Secure Shell (SSH) Transport Layer Protocol (No support diffie-hellman-group-exchange-sha256)
RFC4443	ICMPv6 for IPv6
RFC4444	MIB for IS-IS
RFC4541	IGMP Snooping
RFC4541	MLD Snooping
RFC4552	Authentication/Confidentiality for OSPFv3
RFC4560	DISMAN-PING-MIB
RFC4560	DISMAN-TRACEROUTE-MIB
RFC4560	DISMAN-NSLOOKUP-MIB

RFC No.	Title
RFC4601	PIM-SM
RFC4602	PIM-SM IETF Proposed Std Req Analysis
RFC4604	IGMPv3 & MLDv2 & Source-Specific Multicast
RFC4607	Source-Specific Multicast for IP
RFC4608	PIM--SSM in 232/8
RFC4610	Anycast-RP Using PIM
RFC4611	MSDP Deployment Scenarios
RFC4624	MSDP MIB
RFC4632	Classless Inter-Domain Routing (CIDR)
RFC4668	RADIUS Client MIB
RFC4670	RADIUS Accounting MIB
RFC4673	RADIUS Dynamic Authorization Server MIB
RFC4716	The Secure Shell (SSH) Public Key File Format
RFC4750	OSPFv2 MIB
RFC4835	Crypto Algorithm Requirements for ESP and AH
RFC4836	MAU-MIB
RFC4836	IANA-MAU-MIB
RFC4861	Neighbor Discovery for IPv6
RFC4862	IPv6 Stateless Address Auto configuration
RFC4878	OAM Functions on Ethernet-Like Interfaces
RFC4878	DOT3-OAM-MIB
RFC4884	RFC 4884 Extended ICMP Multi-Part Messages
RFC4940	IANA Considerations for OSPF
RFC5059	Bootstrap Router (BSR) Mechanism for (PIM)
RFC5060	PIM MIB
RFC5095	Deprecation of Type 0 Routing Headers in IPv6
RFC5132	IP Multicast MIB
RFC5176	Dynamic Authorization Extension to RADIUS
RFC5186	IGMPv3/MLDv2/MCAST Routing Protocol Interaction
RFC5187	OSPFv3 Graceful Restart
RFC5240	PIM Bootstrap Router MIB
RFC5250	The OSPF Opaque LSA Option
RFC5294	Host Threats to PIM
RFC5301	Dynamic Hostname Exchange Mechanism for IS-IS
RFC5302	Domain-wide Prefix Distribution with IS-IS
RFC5303	3Way Handshake for IS-IS P2P Adjacencies
RFC5304	IS-IS Cryptographic Authentication
RFC5305	IS-IS extensions for Traffic Engineering
RFC5306	Restart Signaling for IS-IS
RFC5308	Routing IPv6 with IS-IS
RFC5309	P2P operation over LAN in link-state routing
RFC5310	IS-IS Generic Cryptographic Authentication
RFC5340	OSPF for IPv6
RFC5519	MGMD-STD-MIB
RFC5601	Pseudowire (PW) MIB
RFC5643	OSPFv3 MIB
RFC5798	Virtual Router Redundancy Protocol (VRRP) V3

RFC No.	Title
RFC6104	Rogue IPv6 RA Problem Statement
RFC6105	IPv6 Router Advertisement Guard
RFC6106	IPv6 RA Options for DNS Configuration
RFC6164	Using 127-Bit IPv6 Prefixes on Inter-Router Links
RFC6329	IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging
RFC6549	OSPFv2 Multi-Instance Extensions
RFC7348	Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks
Drafts	draft-ietf-idr-as-pathlimit
Drafts	draft-ietf-idr-mrai-dep (Partial Support)
Drafts	draft-ietf-isis-experimental-tlv (Partial Support)
Drafts	draft-ietf-isis-ipv6-te (Partial Support)
Drafts	draft-ietf-ospf-ospfv3-mib
Drafts	draft-ietf-ospf-te-node-addr
Drafts	draft-ietf-idmr-dvmrp-v3-11
Drafts	draft-ietf-vrrp-unified-spec-03.txt

ENTERASYS NETWORKS PRIVATE ENTERPRISE MIB SUPPORT:

Title	Title	Title
CT-BROADCAST-MIB	ENTERASYS-JUMBO-ETHERNET-FRAME-MIB	ENTERASYS-SPANNING-TREE-DIAGNOSTIC-MIB
CTIF-EXT-MIB	ENTERASYS-LICENSE-KEY-MIB	ENTERASYS-SYSLOG-CLIENT-MIB
CTRON-ALIAS-MIB	ENTERASYS-LICENSE-KEY-OIDS-MIB	ENTERASYS-TACACS-CLIENT-MIB
CTRON-BRIDGE-MIB	ENTERASYS-LINK-FLAP-MIB	ENTERASYS-UPN-TC-MIB
CTRON-CDP-MIB	ENTERASYS-MAC-AUTHENTICATION-MIB	ENTERASYS-VLAN-AUTHORIZATION-MIB
CTRON-CHASSIS-MIB	ENTERASYS-MAC-LOCKING-MIB	ENTERASYS-RADIUS-ACCT-CLIENT-EXT-MIB
CTRON-ENVIROMENTAL-MIB	ENTERASYS-MAU-MIB-EXT-MIB	ENTERASYS-RADIUS-AUTH-CLIENT-MIB
CTRON-MIB-NAMES	ENTERASYS-MGMT-AUTH-NOTIFICATION-MIB	ENTERASYS-VLAN-INTERFACE-MIB
CTRON-OIDS	ENTERASYS-MGMT-MIB	IANA-ADDRESS-FAMILY-NUMBERS-MIB
DVMRP-MIB	ENTERASYS-MIB-NAMES-DEFINITIONS	IEEE8021-PAE-MIB
CTRON-Q-BRIDGE-MIB-EXT	ENTERASYS-MIRROR-CONFIG	IEEE8023-LAG-MIB
CISCO-CDP-MIB	ENTERASYS-MSTP-MIB	IEEE8021-BRIDGE-MIB
CISCO-NETFLOW-MIB	ENTERASYS-MULTI-AUTH-MIB	IEEE8021-CFM-MIB
CISCO-TC	ENTERASYS-MULTI-TOPOLOGY-ROUTING-MIB	IEEE8021-CFM-V2-MIB
ENTERASYS-FLOW-LIMITING-MIB	ENTERASYS-MULTI-USER-8021X-MIB	IEEE8021-MSTP-MIB
ENTERASYS-AAA-POLICY-MIB	ENTERASYS-NETFLOW-MIB (v5 & v9)	IEEE8021-Q-BRIDGE-MIB

Title	Title	Title
ENTERASYS-CLASS-OF-SERVICE-MIB	ENTERASYS-OWS-MIB DEFINITIONS	IEEE8021-SPANNING-TREE-MIB
ENTERASYS-CONFIGURATION-MANAGEMENT-MIB	ENTERASYS-OSPF-EXT-MIB	IEEE8023-DOT3-LLDP-EXT-V2-MIB
ENTERASYS-CONVERGENCE-END-POINT-MIB	ENTERASYS-PFC-MIB-EXT-MIB	LLDP-MIB
ENTERASYS-DIAGNOSTIC-MESSAGE-MIB	ENTERASYS-PIM-EXT-MIB	LLDP-EXT-MED-MIB
ENTERASYS-DNS-RESOLVER-MIB	ENTERASYS-POLICY-PROFILE-MIB	LLDP-EXT-DOT1-MIB
ENTERASYS-DVMRP-EXT-MIB	ENTERASYS-POWER-ETHERNET-EXT-MIB	LLDP-EXT-DOT3-MIB
ENTERASYS-ETH-OAM-EXT-MIB	ENTERASYS-PTOPO-MIB-EXT-MIB	LLDP-EXT-DOT3-V2-MIB
ENTERASYS-IEEE8021-BRIDGE-MIB-EXT-MIB	ENTERASYS-PWA-MIB	LLDP-EXT-DOT3-V2-MIB (IEEE 802.3-2009) ETS Admin table read only
ENTERASYS-IEEE8021-SPANNING-TREE-MIB-EXT-MIB	ENTERASYS-RESOURCE-UTILIZATION-MIB	RSTP-MIB
ENTERASYS-IEEE8023-LAG-MIB-EXT-MIB	ENTERASYS-RIPv2-EXT-MIB	U-BRIDGE-MIB
ENTERASYS-IETF-BRIDGE-MIB-EXT-MIB	ENTERASYS-RMON-EXT-MIB	USM-TARGET-TAG-MIB
ENTERASYS-IETF-P-BRIDGE-MIB-EXT-MIB	VSB-SHARED-SECRET-MIB	ENTERASYS-VRRP-EXT-MIB DEFINITIONS
ENTERASYS-IF-MIB-EXT-MIB	ENTERASYS-SNTP-CLIENT-MIB	SNMP-RESEARCH-MIB
ENTERASYS-IP-SLA-MIB	ENTERASYS-ENTITY-SENSOR-MIB-EXT-MIB	

Extreme Networks Private Enterprise MIBs are available in ASN.1 format from the Extreme Networks web site at www.extremenetworks.com/support/policies/mibs/. Indexed MIB documentation is also available.

SNMP TRAP SUPPORT:

RFC No.	Title
RFC 1493	New Root Topology Change
RFC 1850	ospfIfStateChange ospfVirtIfStateChange ospfNbrStateChange ospfVirtNbrStateChange ospfIfConfigError ospfVirtIfConfigError ospfMaxAgeLsa ospfOriginatLsa
RFC 1907	Cold Start Warm Start Authentication Failure
RFC 4133	entConfigChange
RFC 2668	ifMauJabberTrap

RFC No.	Title
RFC 2819	risingAlarm fallingAlarm
RFC 2863	linkDown linkup
RFC 2922	ptopoConfigChange
RFC 2787	vrrpTrapNewMaster vrrpTrapAuthFailure
RFC 3621	pethPsePortOnOffNotification pethMainPowerUsageOnNotification pethMainPowerUsageOffNotification
RFC4268	entStateOperEnabled entStateOperDisabled
Enterasys-mac-locking-mib	etsysMACLockingMACViolation
Cabletron-Traps.txt	boardOperational boardNonOperational wgPsInstalled wgPsRemoved wgPsNormal wgPsFail wgPsRedundant wgPsNotRedundant fanFail fanNormal boardInsertion boardRemoval
	etsysPseChassisPowerRedundant etsysPseChassisPowerNonRedundant etsysPsePowerSupplyModuleStatusChange
Power-ethernet-mib	pethPsePortOnOffNotification pethMainPowerUsageOnNotification pethMainPowerUsageOffNotification
Enterasys-link-flap-mib	etsysLinkFlapViolation
Enterasys-ietf-bridge-mib-ext-mib	etsysIetfBridgeDot1qFdbNewAddrNotification etsysIetfBridgeDot1dSpanGuardPortBlocked etsysIetfBridgeDot1dBackupRootActivation etsysIetfBridgeDot1qFdbMovedAddrNotification etsysIetfBridgeDot1dCistLoopProtectEvent
Enterasys-flow-limiting-mib	etsysFlowLimitingFlowCountActionLimit1 etsysFlowLimitingFlowCountActionLimit2
Enterasys-notification-auth-mib	etsysMgmtAuthSuccessNotification etsysMgmtAuthFailNotification
Enterasys-multi-auth-mib	etsysMultiAuthSuccess etsysMultiAuthFailed etsysMultiAuthTerminated etsysMultiAuthMaxNumUsersReached etsysMultiAuthModuleMaxNumUsersReached etsysMultiAuthSystemMaxNumUsersReached
Enterasys-spanning-tree-diagnostic-mib	etsysMstpLoopProtectEvent etsysStpDiagCistDisputedBpduThresholdExceeded etsysStpDiagMstiDisputedBpduThresholdExceeded

RFC No.	Title
Lldp-mib	lldpNotificationPrefix (IEEE Std 802.1AB-2004)
Lldp-ext-med-mib	lldpXMedTopologyChangeDetected (ANSI/TIA-1057)
Enterasys-class-of-service-mib	etsysCosIrIExceededNotification
Enterasys-policy-profile-mib	etsysPolicyRulePortHitNotification
Enterasys-mstp-mib	etsysMstpLoopProtectEvent
Crtron-environment-mib	chEnvAmbientTemp chEnvAmbientStatus

RADIUS ATTRIBUTE SUPPORT:

This section describes the support of RADIUS attributes on the K-Series modules. RADIUS attributes are defined in [RFC 2865](#) and [RFC 3580](#) (IEEE 802.1X specific).

RADIUS AUTHENTICATION AND AUTHORIZATION ATTRIBUTES:

Attribute	RFC Source
Called-Station-Id	RFC 2865, RFC 3580
Calling-Station-Id	RFC 2865, RFC 3580
Class	RFC 2865
EAP-Message	RFC 3579
Filter-Id	RFC 2865, RFC 3580
Framed-MTU	RFC 2865, RFC 3580
Idle-Timeout	RFC 2865, RFC 3580
Message-Authenticator	RFC 3579
NAS-IP-Address	RFC 2865, RFC 3580
NAS-Port	RFC 2865, RFC 3580
NAS-Port-Id	RFC 2865, RFC 3580
NAS-Port-Type	RFC 2865, RFC 3580
NAS-Identifier	RFC 2865, RFC 3580
Service-Type	RFC 2865, RFC 3580
Session-Timeout	RFC 2865, RFC 3580
State	RFC 2865
Termination-Action	RFC 2865, RFC 3580
User-Name	RFC 2865, RFC 3580
User-Password	RFC 2865

RADIUS ACCOUNTING ATTRIBUTES:

Attribute	RFC Source
Acct-Authentic	RFC 2866
Acct-Delay-Time	RFC 2866
Acct-Interim-Interval	RFC 2866
Acct-Session-Id	RFC 2866
Acct-Session-Time	RFC 2866
Acct-Status-Type	RFC 2866
Acct-Terminate-Cause	RFC 2866
Calling-Station-ID	RFC 2865

GLOBAL SUPPORT:

By Phone: +1 800-998-2408 (toll-free in U.S. and Canada)

For the toll-free support number in your country:

www.extremenetworks.com/support/

By Email: support@extremenetworks.com

By Web: www.extremenetworks.com/support/

By Mail: Extreme Networks, Inc.

6480 Vía Del Oro

San Jose, CA 95119

+1 408-579-2800

For information regarding the latest software available, recent release note revisions, or if you require additional assistance, please visit the Extreme Networks Support web site.