

S-Series® and S-Series® Standalone Customer Release Note

Firmware Version 8.63.05.0004

Table of Contents

Customer Release Notes.....	4
Installing an I/O or I/O Fabric Module	6
Modules Minimum FW Version Required and HW Dependent Features:.....	6
Multislot Chassis—Minimum FW Version Required:.....	8
Option Module Behavior:	10
MAC Address Capacity	10
Multi-slot Chassis User Capacities.....	10
Maximum User Capacity:	11
S130 Class modules Multi-User Capacities.....	12
SSA User Capacities:.....	12
S130 Class SSA Multi-User Capabilities	12
S150 Class SSA Multi-User Capacities.....	12
S180 Class SSA Multi-User Capacities.....	13
SSA User Capacity Upgrade License.....	13
Chassis Multi-User Capacities Licensing.....	13
S-E OS-PPC - Port Capacities License	13
Port Mirroring.....	13
Class of Service:.....	14
Class of Service Support.....	14
Link Aggregation (LAG).....	14
Multi-User 802.1X	15
QSFP+ Mixed Port Speed Operational Overview	15
Compatibility mode:.....	16
Power over Ethernet Controller Code Upgrade	17
Features, Scale and Capacity	17
Router Capacities (Brief).....	17
DHCP Capacities	18
SDN Mode Capacities	19
License Features	19
Virtual Switch Bonding Licenses	21
100Mb Optics:.....	21
10Gb Optics:	22
QSFP-SFPP-ADPT transceiver support:.....	24
The 10GB-LW-SFPP adapter is not supported in the QSFP-SFPP-ADPT adapter.....	24
Capacity Reductions in 8.63.05.0004	28
Problems Corrected in 8.63.05.0004.....	28
Problems Corrected in 8.63.04.0003	30
New issue present in 08.63.04.0003 caused by fix for Router Protocols Problems	31

Problems Corrected in 8.63.03.0002.....	32
Problems Corrected in 8.63.02.0004.....	34
Problems Corrected in 8.63.01.0019.....	36
Problems Corrected in 8.62.04.0001.....	38
Problems Corrected in 8.62.03.0006.....	39
Problems Corrected in 8.62.02.0022.....	41
Problems Corrected in 8.61.02.0001.....	49
Features Enhancements 8.61.01.0018.....	49
Problems Corrected in 8.61.01.0018.....	50
Problems Corrected in 8.42.06.0001.....	52
Problems Corrected in 8.42.05.0003.....	52
Problems Corrected in 8.42.04.0016.....	53
Features Enhancements 8.42.02.0012.....	57
Problems Corrected in 8.42.02.0012.....	58
Feature Enhancements 8.42.01.0007.....	60
Problems Corrected in 8.42.01.0007.....	61
Features Enhancements 8.41.01.0004.....	62
Problems Corrected in 8.41.01.0004.....	65
Features Enhancements 8.32.02.0008.....	73
Features Enhancements 8.32.01.0021.....	75
Problems Corrected in 8.32.01.0021.....	76
Problems Corrected in 8.31.03.0001.....	85
Features Enhancements 8.31.01.0006.....	92
Problems Corrected in 8.22.03.0006.....	105
Problems Corrected in 8.22.01.0022.....	109
Feature Enhancements in 8.21.02.0001.....	121
Feature Enhancements in 8.11.05.0006.....	137
Feature Enhancements in 8.11.04.0005.....	139
Feature Enhancements in 8.11.03.0005.....	142
Feature Enhancements in 8.02.01.0012.....	150



6480 Via Del Oro
San Jose, CA 95119

+1 408-579-2800

Customer Release Notes

S-Series® and S-Series® Standalone

Firmware Version 8.63.05.0004

April 2019

INTRODUCTION:

This document provides specific information for version 08.63.05.0004 of firmware for the modular chassis and standalone versions of the S-Series including: S180, S140, S155, S150, and S130 class of S-Series Modules and the S-Series Standalone (SSA) 1RU chassis. The S180/S140/S155/S150 and S130 modules may be installed in the S8, S6, S4 and S1A chassis. The S140/S130 class I/O modules may also be installed in the S3 chassis. This version of firmware supports the following S-Series chassis and SSA switches:

S180 Class Modules and SSAs

SL8013-1206-F8	SL8013-1206-F8A	SK8008-1224-F8	SK8009-1224-F8
ST8206-0848-F8	ST8206-0848-F8A	SG8201-0848-F8	SK8208-0808-F8
SL8013-1206	SL8013-1206A	SK8008-1224	SK8009-1224
SKL8008-0810-F8	SSA-T8028-0652	SSA-G8018-0652	PV-FC-180

S140 Class Modules

ST2206-0848	ST2206-0848A	SG2201-0848	SK2008-0832
SK2009-0824	SGL2001-0850	STL2006-0850	

S155 Class Modules

SK5208-0808-F6	ST5206-0848-F6	SG5201-0848-F6	
----------------	----------------	----------------	--

S150 Class Modules and SSAs

SK1208-0808-F6	ST1206-0848-F6	SG1201-0848-F6	SK1008-0816
ST1206-0848	SG1201-0848	SSA-T1068-0652A	SSA-T1068-0652
SSA-G1018-0652			

S130 Class Modules and SSAs

ST4106-0248	SG4101-0248	ST4106-0348-F6	SSA-T4068-0252
-------------	-------------	----------------	----------------

Option Modules			
SOK1208-0102	SOK1208-0104	SOK1208-0204	SOG1201-0112
SOT1206-0112	SOK2208-0102	SOK2208-0104	SOK2208-0204
SOK2209-0204	SOG2201-0112	SOT2206-0112	SOGK2218-0212
SOTK2268-0212	SOV3208-0202	SOV3008-0404	

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	8.63.05.0004	Customer Release	April 2019
Previous Version	8.63.04.0003	Customer Release	February 2019
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	8.21.02.0001	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	August 2013
Previous Version	8.11.02.0001	Customer Release	July 2013

Previous Version	8.11.01.0014	Customer Release	June 2013
<p>Note: This image provides support for all S-Series HW classes in a single image. Prior to this version the S-Series FW was released as three separate images. (An image for the S140 I/O modules, SSA180/SSA150A and S130/S150/S155/SSA130/SSA150).</p> <p>Warning: The multicast capacity for the S130/S150/S155 and SSA130/SSA150 classes has been reduced in this image to allow mixed class compatibility. Please refer to the multicast capacities section found on page 10 of this note. An alternate image for S130/S150/S155 and SSA130/SSA150 classes only, with the previous established multicast capacity is available for download.</p>			

HIGH AVAILABILITY UPGRADE (HAU) FW COMPATIBILITY:

This version will be 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

When installing a new module to an existing system, the system's operating firmware image needs to be compatible with the new module. It is recommended that the system be upgraded prior to installation. If the system isn't upgraded prior to the installation, the new module may not complete initialization and join the rest of the chassis. It will remain in a halted state until the running chassis is upgraded to a compatible firmware version.

Modules Minimum FW Version Required and HW Dependent Features:

S180 Class Modules	Minimum FW Version	DCB	EEE	MACsec	MACsec Notes
SL8013-1206-F8A	8.31.01.0006	X	-	X	40G ports in 10G mode
SL8013-1206-F8	8.11.01.0014	X	-	-	
SK8008-1224-F8	8.11.01.0014	X	-	X	
SK8009-1224-F8	8.11.01.0014	X	X	X	
ST8206-0848-F8	8.22.01.0022/ (8.12.02.006)	X	-	-	
ST8206-0848-F8A	8.22.02.0012	X	X	X	
SG8201-0848-F8	8.22.01.0022/ (8.12.02.006)	X	-	-	

SK8208-0808-F8	8.22.01.0022/ (8.12.02.006)	X	-	-	
SL8013-1206A	8.31.01.0006	X	-	X	40G ports in 10G mode
SL8013-1206	8.11.01.0014	X	-	-	
SK8008-1224	8.11.01.0014	X	-	X	
SK8009-1224	8.11.01.0014	X	X	X	
SKL8008-0810-F8	8.32.01.0021	X	-	X	Only 40G ports running in 10G mode
S140 Class Modules					
ST2206-0848	8.01.01.0016	X	-	-	
S155 Class Modules					
ST2206-0848A	8.22.02.0012	X	X	X	
SG2201-0848	8.01.01.0016	X	-	-	
SK2008-0832	8.01.01.0016	X	-	X	
SK2009-0824	8.02.01.0012	X	X	X	
SGL2001-0850	8.32.01.0021	X		X	Only 40G ports running in 10G mode
STL2006-0850	8.32.01.0021	X	X	X	1G ports and 40G ports running in 10G mode
S150 Class Modules					
SK5208-0808-F6	7.21.02.0002	-	-	-	
ST5206-0848-F6	7.21.02.0002	-	-	-	
SG5201-0848-F6	7.21.02.0002	-	-	-	
S150 Class Modules	Minimum FW Version	DCB	EEE	MACsec	MACsec Notes
SK1208-0808-F6	7.01.01.000X	-	-	-	
ST1206-0848-F6	7.01.01.000X	-	-	-	
SG1201-0848-F6	7.01.01.000X	-	-	-	
SK1008-0816	7.01.01.000X	-	-	-	
ST1206-0848	7.01.01.000X	-	-	-	
SG1201-0848	7.01.01.000X	-	-	-	
S130 Class Modules					
ST4106-0248	7.02.02.0002	-	-	-	
SG4101-0248	7.02.02.0002	-	-	-	
ST4106-0348-F6	7.02.02.0002	-	-	-	
Option Modules (Compatible with S130/S150/S155 only)					
SOK1208-0102	7.01.01.000X	-	-	-	
SOK1208-0104	7.01.01.000X	-	-	-	
SOK1208-0204	7.01.01.000X	-	-	-	
SOG1201-0112	7.01.01.000X	-	-	-	
SOT1206-0112	7.01.01.000X	-	-	-	
Option Modules (Compatible with all classes)					
SOK2208-0102	7.72.01.0021	-	-	-	
SOK2208-0104	7.72.01.0021	-	-	-	

SOK2208-0204	7.72.01.0021	-	-	-	
SOK2209-0204	8.42.02.0012	-	-	X	
SOG2201-0112	7.72.01.0021	-	-	-	
SOT2206-0112	7.72.01.0021	-	-	-	
SOGK2218-0212	7.72.01.0021	-	-	-	
SOTK2268-0212	8.02.01.0012	-	-	X	10G ports only
Option Modules (Compatible S140/S180 only)					
SOV3208-0202	8.11.01.0014	-	-	-	

Expansion Modules (Compatible with S180, 10/40Gb I/O Modules)					
SOV3008-0404	8.11.01.0014	-	-	-	
SSA/Purview Models					
SSA-T1068-0652A	8.01.01.0016	-	-	X	10G ports only
SSA-T8028-0652	8.01.01.0016	-	-	X	10G ports only
SSA-G8018-0652	8.01.01.0016	-	-	X	10G ports only
SSA-G1018-0652	7.01.01.000X	-	-	-	
SSA-T1068-0652	7.01.01.000X	-	-	-	
SSA-T4068-0252	7.01.01.000X	-	-	-	
PV-FC-180	8.31.02.0014	-	-	X	10G ports only

Multislot Chassis—Minimum FW Version Required:

Multislot Chassis	
S8-Chassis	07.01.01.000X
S8-Chassis-POE4	
S8-Chassis-POE8	
S4-Chassis	
S4-Chassis-POE4	
S3-Chassis	
S3-Chassis-POE4	
S3-Chassis-A	07.72.01.0021
S3-Chassis-POEA	
S6-Chassis	07.22.01.0002
S6-Chassis-POE4	
S1-Chassis	07.73.01.0003
S1-Chassis-A	08.11.01.0014

System Behavior

The S-Series I/O modules when combined in a chassis, will select a master module to control the overall management of the system. All information that the master module controls is distributed to all modules in the chassis. In the event that the master module is unable to continue the management task, another module will automatically assume responsibility for answering management queries and distributing system information.

If a new module is inserted into the system, the new module will inherit all system parameters and all firmware files stored on each module in the system. Any firmware files stored on the new device, which are not common to the system, will be automatically removed. If the new module does not have a copy of the current

system's boot image, it will automatically be upgraded, and then the module will re-initialize and join the system.

NOTE:

If the new module requires a newer firmware image than the image running in the chassis, the master module **MUST** be upgraded to the newer firmware before inserting the new module. If the system isn't upgraded prior to the installation, the new module will not complete initialization and join the rest of the chassis. It will remain in a halted state until the running chassis is upgraded to a compatible firmware version.

The system will treat the following conditions as if a new module (I/O or I/O fabric module) has been installed:

- Moving module from one slot to another
- Moving module to another chassis
- If an Option Module is added or removed from a blade* (See Option Module Behavior table below)

Configuration may be cleared for other reasons including (but not limited to):

- DIP switch 7,
- CLI command,
- MIB manipulation

If a module needs to be replaced, it will inherit all the configuration settings of the previous module as long as the new module is an exact replacement of model number, slot number and Option Module (if one was previously installed). Any configuration files that were stored in the file system of the newly inserted module will not be deleted and will be available to reconfigure the system.

Option Module Behavior:

Original HW Config	New HW Config	Resulting Action
No Option Module	Option Module	No config change
Option Module	No Option Module	No config change
Option Module Rev. X	Option Module Rev. Y	No config change
Option Module Type A	Option Module Type B	Option Module config cleared

If configuration exists for an Option Module (or its ports) that config will remain after the Option Module is removed until such time as one of the above clearing events takes place. This means an Option Module could be removed, RMA-ed, and replaced with a like type and the configuration for those ports will be restored even if the board it used without the Option Module in the interim.

MAC Address Capacity

128K MAC addresses are supported.

Multi-slot Chassis User Capacities

Each of the empty S-Series chassis (S1(A)/S3(A)/S4/S6/S8 and the POE variants) has a user capacity entitlement of 1024 users. The chassis will combine its user capacity with the user capacity of the blades installed in the chassis to derive the total user capacity for the populated chassis.

Maximum User Capacity:

Chassis Type	Maximum User Capacity
S8-Chassis	9,216 (9K)
S8-Chassis-POE4	
S8-Chassis-POE8	
S6-Chassis	6,122 (6K)
S6-Chassis-POE4	
S4-Chassis	5,120 (5K)
S4-Chassis-POE4	
S3-Chassis	4,096 (4K)

Chassis Type	Maximum User Capacity
S3-Chassis-POE4	(S140 Class)
S3-Chassis-A	
S3-Chassis-POEA	
S1-Chassis	8,192 (8K) *Licensed
S1-Chassis-A	

S180/150/S155/S140 Class modules Multi-User Capacities

Each module contributes 1024 users to the overall chassis capacity. Each module has unrestricted access to the entire system user capacity. This allows for up to the populated system's user capacity to be consumed on a single port.

Module Class	Blade Contribution	Restrictions (if applicable)
S180/S140	1024 Users	None
S150/S155		
S130	512 Users	8 Users/port

S130 Class modules Multi-User Capacities

Each of the S130 modules contributes 512 users to the overall chassis capacity. Each S130 class module has restricted access to the user capacity based on port type.

Each S130 high density 10/100/1000Mb copper port supports up to 8 authenticated users per port. This applies to the ST4106-0248 module and SOT1206-0112 option module. Each S130 high density SFP port supports up to 8 authenticated users per port. This applies to the SG4101-0248 module.

Uplink ports installed on the S130 modules, defined as modular SFP, 10 Gbps, and 100Mb FX ports, support up to 128 authenticated users per port. This includes modules:

Series 1 Option Modules	Series 2 Option Modules
SOK1208-0102	SOK2208-0102
SOK1208-0104	SOK2208-0104
SOK1208-0204	SOK2208-0204
SOG1201-0112	SOG2201-0112
	SOGK2218-0212
	SOTK2268-0212
	SOK2209-0204

SSA User Capacities:

Chassis Type	Class	Default User Capacity	Licensed User Capacity
SSA-T4068-0252	SSA130	512	1K
SSA-T1068-0652	SSA150	2K	4K
SSA-G1018-0652	SSA150	2K	4K
SSA-T1068-0652A	SSA150	2K	4K

Chassis Type	Class	Default User Capacity	Licensed User Capacity
SSA-T8028-0652	SSA180	4K	8K
SSA-G8018-0652	SSA180	4K	8K

S130 Class SSA Multi-User Capabilities

The S130 SSA supports a total capacity of 512 users. The S130 SSA has restricted access to the user capacity based on port type. The S130 high density 10/100/1000Mb copper port supports up to 8 authenticated users per port. Uplink SFP+ ports on the S130 SSA support up to 128 authenticated users per port. 802.3 LAG ports support 128 users. This applies to model number SSA-T4068-0252.

An 'S-EOS-PPC' license can be used to remove the per port restrictions, allowing unrestricted access to the total 512 user capacity.

S150 Class SSA Multi-User Capacities

Each of the S150 SSAs supports a total capacity of 2048 users. Each S150 SSA has unrestricted access to the entire user capacity. This allows for up to the entire system's user capacity to be consumed on a single port. This applies to model numbers, SSA-T1068-0652, SSA-T1068-0652A and SSA-G1018-0652.

S180 Class SSA Multi-User Capacities

Each of the S180 SSAs supports a total capacity of 4096 users. Each S180 SSA has unrestricted access to the entire user capacity. This allows for up to the entire system's user capacity to be consumed on a single port. This applies to model numbers, SSA-T8028-0652 and SSA-G8018-0652.

SSA User Capacity Upgrade License

An optional user capacity upgrade license is available for the SSA. The SSA-EOS-2XUSER license doubles the user capacity of the SSA it is installed on.

- In the SSA180 class the default will be increased from 4096 to 8192 users per SSA
- In the SSA150 class the default will be increased from 2048 to 4096 users per SSA.
- In the SSA130 class the default capacity will be increased from 512 to 1024 user per SSA.

The license, when applied to the SSAS130 class, also removes the per port user restrictions, allowing for the entire capacity of the device to be authenticated on a single port.

Chassis Multi-User Capacities Licensing

An optional license for the S1-Chassis and S1-ChassisA is available. The S1-EOS-USER User Capacity license can be applied to the S1 and allows the chassis to support up to 8k users in total.

An optional license for the S130 Class is available. The S-EOS-PPC license removes the per port user capacity restriction, allowing access to the entire system capacity. The S-EOS-PPC license is applied to a module and is required, if default port user capacities on that module are to be exceeded.

S-EOS-PPC - Port Capacities License

A license is required for each S130 module requiring additional port user capacity.

The license removes the per port restriction of 8 or 128 users per port for a specified module.

Users per port increase to a maximum value of the system capacity, with a default value of 256 users/port.

When present, the PPC license defaults the user capacity at 256 users per port. This value can be overridden using the CLI command 'set multiauth port numusers' and increased to the maximum allowable by the system.

Port Mirroring

The S-Series device provides support for 15 mirrors. A mirror could be a:

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

For the "one-to-many" there is no limit to the amount of destination ports.

For the "many-to-one" there is no limit to the amount of source ports.

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

The port and VLAN mirror function does not mirror error frames.

This allows support for an IDS mirror and 14 other active mirrors.)

Note that the examples above are provided to illustrate the number and types of mirrors we support, as well as how they can be used concurrently. The mirror configurations are not limited to these examples.

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 a L2 GRE tunnel and can be routed across the network.

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
- Queues
 - Support for Strict, WFQ and Hybrid Arbitration
 - All queues support rate-shaping
 - S130/S150 Classes, 12 Transmit Queues per port (1 reserved for control-plane traffic)
 - SSA130/SSA150 Classes, 12 Transmit Queues per port (1 reserved for control-plane traffic)
 - S155/S180/S140 Classes, 16 Transmit Queues per port (1 reserved for control-plane traffic)
 - SSA180 Class, 16 Transmit Queues per port (1 reserved for control-plane traffic)
- Rate Limiters
 - 32 Inbound-Rate-Limiters per port (SSA130/S130-class 10/100/1000 ports support 24)
 - 16 Outbound-Rate-Limiters per port (SSA130/S130-class 10/100/1000 ports support 4)
- Support for Flood-Limiting controls for Broadcast, Multicast, and Unknown Unicast per port.
- Management
 - Support for Enterasys CoS MIB

Link Aggregation (LAG)

The S-Series chassis, S1/S3/S4/S8, supports a total of 190 LAGs per chassis with up to 64 ports per LAG.

The SSA products support up to 62 LAGs per SSA with up to 64 ports per LAG.

* Support for no more than 1 IDS mirror. An IDS mirror can have up to 10 destination ports in it. (Note the major change from 6.X series FW on the N-Series – an IDS mirror now takes only one mirror resource.)

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

Multi-User 802.1X

Authentication of multiple 802.1X clients on a single port is supported. This feature will only operate 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 will result in the 802.1X authenticator missing multicast destined users' logoff and login messages. Systems used by multiple consecutive users will remain authenticated as the original user until the re-authentication period has expired.

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 will be required to display and manage multiple users (stations) on a single port.

QSFP+ Mixed Port Speed Operational Overview

Each 40Gb QSFP+ port supports operation as (1) 40G port or (4) 10G ports. Groups of 2 QSFP+ ports must operate in same mode, fg.x.1-2, fg.x.3-4 and fg.x.5-6. The grouped ports will be referred to as a "port speed group". The system always presents all possible 40Gb and 10Gb ports, fg.x.1-6 and tg.x.1-24 and ports not associated with their active operating speed display as 'oper-status not-present'. Example – If first speed configuration group is operating in 40G mode then tg.slot.1-8 will convey an 'oper-status not-pres'.

Port speed may be changed using one of the following methods:

1. Via the CLI
2. Via the standard mib port speed attributes
3. Insertion of a 40Gb or 10Gb transceiver that is not in conflict with other members in the port speed group. Conflict is defined as a transceiver that requires a different speed than is currently operating and there is a QSFP+ installed in the other port of the port speed group that is compatible with the current operating speed. If there is conflict then the system reports "conflict".

Note: In some cases the module must reset to transition the ports to new operating speeds.

A new operating speed can be selected by using the CLI command 'set port speed fg.x.y 10000'. The command requests 10G operating mode for all ports in the port speed group that fg.x.y is a member of. The CLI command 'set port speed tg.x.y 40000' requests 40Gb operating mode for all ports in the port speed group that fg.x.y is a member of.

Currently these commands are only supported for ports that are "present", meaning you can't "recall" a speed setting without first rebooting the blade. Exception: You can leverage a QSFP+ reinsertion to revert the speed change – see Note below.

Warning: If you attempt to retract a speed change using a not-present port, the system will appear to accept the retraction (including syslog to the effect indicating the change will happen on reset) but the first setting will be taken upon reset.

A future release of FW will permit recalling a speed setting change via CLI.

Note: Currently there is only one way to “recall” a requested speed change. You must insert or remove/reinsert a QSFP+ in the port speed group that can operate at the original speed. After the removal /reinsertion the ports will no longer be held ‘oper-status down’ for “self” and will return to normal operation immediately. The speed change scheduled for the next reset will be canceled. The CLI command ‘show port speed’ will convey the reverted state.

When a new operating speed is selected:

- The system reports a syslog message indicating the blade must be reset to adjust to requested speed.
- The ports in the speed configuration group associated with the new operating speed remain “not-pres” until blade resets.
- The ports in the speed configuration group not associated with the new desired operating speed go ‘oper- status down’ with oper-status cause “self”.
- The blade must be reset to complete the speed transition.

Many QSFP+ devices support operation at both 10Gb and 40Gb speeds. These include QSFP+ assemblies with fixed cable assemblies that have QSFP terminations at both ends of the assembly, such as Direct Attached Cables (DAC). At the time of this writing only the QSFP+ to 4x SFP+ “hydra” cable assemblies which terminate one end with a SFP+ and QSFP+ to single SFP+ adapters must operate in 10Gb.

Summary: To establish an operational QSFP+ port two conditions must be adhered to:

- The port speed and transceiver desired must not conflict with the existing members of the port speed group.
- The QSFP+ transceiver must be compatible with the provisioned operating speed for the port.

Compatibility mode:

Compatibility mode establishes the type of signaling that will be used on the backplane between modules. It affects the way the S180 fabric operates (fixed vs variable cells and no bonding header vs bonding header). Compatibility mode version one (v1) must be used whenever the chassis has a legacy S130/S150/S155 card installed. Compatibility mode version two (v2) should be used when all of the modules are 140/S180 class.

By default compatibility mode is automatically established upon first boot up of 7.99.06 or greater 7.99 images (factory images) and 8.11 FW and newer, or any time the configuration is lost (clear config, switch 7 on all fabrics, all new fabrics) or the following commands are issued clear chassis compatibility, ‘set chassis compatibility auto [chassis-id]’.

The automatic assignment occurs once at boot time and when established the operational compatibility mode will be sticky and persist through various HW changes, or until the configuration is manually changed or cleared. (Chassis-id may be omitted on systems with bonding disabled and will default to chassis1 on bonding systems.)

There are several reasons a compatibility setting would need to be manipulated.

1. If an existing S180 class chassis has an S130 or S150 I/O module added, the chassis compatibility mode will need to change from v2 to v1. The S130/S150 I/O module will be prevented from joining the system until the compatibility mode is set to v1 for the chassis.
2. If a combined HW class chassis has all of the legacy S130/S150 class HW removed the compatibility setting should be manually changed to v2. When configured in v2 mode the fabrics run with different signaling. When possible the HW should be configured in v2 mode.
3. *VSB considerations:* Each physical chassis operates with its own compatibility setting. When selecting the appropriate compatibility mode setting you must consider the HW population of the individual physical chassis participating in the bond. (not the logical combined bonded chassis)
 - a. If a S140/S180 only chassis is to be bonded to a S150/S155/S130 chassis the S140/S180 should have a compatibility setting of v2 and the S150/S155/S130 class chassis will have a v1 setting.
 - b. A similar consideration must be made when a S3-S130 class chassis is to bond to a S3-S140 chassis. The compatibility setting for the S3-S140 should be v2 and the S130-S3 will use the v1 setting.

Power over Ethernet Controller Code Upgrade

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

Features, Scale and Capacity

Each release of S-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 HW with its current licensing. Please use this command to get a complete list of capacities for this release.

Router Capacities (Brief)

	S180/S140 /S155	S150	S130	SSA180	SSA150 / SSA150 A	SSA130
ARP Entries (per router / per chassis)	65,000					
Static ARP Entries	1,024					
IPv4: Route Table Entries (RIB)	1.6M	100,000	100,000	1.6M	100,000	100,000
IPv4: Route Table Entries (FIB)	800K					
IPv6: Route Table Entries (RIB)	128K	25,000	25,000	128K	25,000	25,000
IPv6: Route Table Entries (FIB)	100K					
IPv4: Router interfaces	1,024					
IPv6: Router interfaces	256					
OSPF Areas	16					
OSPF LSA(s)	50,000					
OSPF Neighbors	60					
Static Routes	2,048					
RIP Routes	3,000					
Configured RIP Nets	300					
VRRP Interfaces	1,024					
Routed Interfaces	1,024					
ACLs	1,000					
-Access Rules	5,000					
-Access Rules – Per ACL	5,000					
Policy Based Routing Entries	100					
ECMP Paths	8					
Static VRFs	256	128	128 *License d	256	128/256	128 *License d
Dynamic VRFs	128	64	16 *License d	128	64/128	16 *License d
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

	S180/S140 /S155	S150	S130	SSA180	SSA150/ SSA150A	SSA130
IGMP/MLD Static Entries	64					
IGMP/MLD *,G and S,G Groups ¹	64K					
IGMP/MLD Snooping Flow Capacity	5K	5K	5K	5K	5K	5K
Multicast Routing (PIM/DVMRP flows)	5K	5K	5K	5K	5K	5K
Multicast Routing (PIM/DVMRP flows) When Virtual Switch Bonded in a S3/S4/S6 or S8 chassis	5K	5K	5K	-	-	-
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

	S180/S140 /S155	S150	S130	SSA180	SSA150/ SSA150 A	SSA130
DHCP Server Leases	5,000					
DHCP Pools	100					

TWCB Capacities

	S180/S140 /S155	S150	S130	SSA180	SSA150/ SSA150 A	SSA130
Bindings	131,072	131,072	131,072	131,072	131,072	131,072
Caches	500					
Servers Farms	50					
WebCaches	50					

LSNAT Capacities

	S180/S140 /S155	S150	S130	SSA180	SSA150/ SSA150 A	SSA130
LSNAT Bindings	131,072	131,072	-	131,072	131,072	-
SLB Real Server IPs	500	500	-	640	640	-
SLB Server Farms	320					
VIP Addresses	1,000					
SLB Virtual Servers	500					
Sticky Entries	131,072	131,072	-	131,072	131,072	-

NAT Capacities

	S180/S140/S155	S150	S130	SSA180	SSA150/SSA150A	SSA130
Bindings	131,072	131,072	-	131,072	131,072	-
IP Addresses (Dynamic/Static)	2,000					
Source List Rules	10					
Address Pools	10					
Dynamic Port Mapped Addresses	20					
Static Translation Rules	1,000					
Translation Protocol Rules	50					

Shortest Path Bridging

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

Tunneling Capacities

	S180/S140/S155	S150	S130	SSA180	SSA150/SSA150A	SSA130
Total Number of Tunnels	62	62 *Licensed	62 *Licensed	62	62 *Licensed	62 *Licensed

SDN Mode Capacities

OpenFlow Capacities	S180 or S155 Fabric in an S1-CHASSIS, SSA180, or SSA150A only
OpenFlow Tables	256 – each for general purpose
Maximum Number of Flows	500,000 [‡]
Changes to Scale of Traditional Forwarding Methods when Using SDN Mode	
IPv4: Route Table Entries (RIB)	16K
IPv4: Route Table Entries (FIB)	16K
IPv6: Route Table Entries (RIB)	16K
IPv6: Route Table Entries (FIB)	16K
ARP Entries (per router / per chassis)	16K
LS-NAT Bindings	1K
NAT Bindings	1K

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

License Features

The S-EOS-L3-S130 license adds VRF, BGP and tunneling features to the S130 class of HW.

A single license will be required per chassis or SSA. The license is applicable to:

- S130 class SSA,
- S3 chassis (using S130 I/O modules),

- The S1, S4, S6 and S8 chassis using the S130 Class fabrics or a combination of S150 and S130 Class fabrics (The VRF and BGP functionality in the S150 class is included without the need for a license.)

The S-EOS-L3-S150 license adds tunneling support to the S150 Class of HW. This license will be extended in the future to add additional tunneling options. The S155 class supports these features without the need for the license.

SSA-EOS-2XUSER license doubles the default user capacity of the SSA. In the S130 class the default capacity will be increased from 512 to 1024 users/SSA and per port restrictions will be removed allowing for the entire user capacity to be consumed on a single port. In an SSA150 class the default will be increased from 2048 to 4096 users per SSA. In an SSA180 class the default will be increased from 4096 to 8192 users per SSA.

S1-EOS-USER S1/S1A User Capacity License - User Capacity license allows support up to 8k users. Used for single fabric systems installed in S1/S1A chassis.

S-EOS-Flow - Flow capacity license for SSA/Purview appliance 10G ports to allow 1M flows per CoreFlow2 ASIC. Applicable to 10G ports only, see port to ASIC mapping.

The MACsec licenses are applied a on a per module basis (not per chassis) and require a unique license per module when using MACsec. There is a HW dependency required to support MACsec. Please see the “Modules Minimum FW Version Required and HW Dependent Features” for a list of capable modules.

S-EOS-MACSEC – MACsec license for 1Gb S-Series modules and modules with MACsec capable uplinks.

S-EOS-KMACSEC – Macsec licenses for 10Gb S-Series modules and 40Gb S-Series modules.

‡ Flow installation times will vary. Please speak with your Extreme Networks representative to set appropriate expectations for OpenFlow capacities based on your application.

§ MACsec may only be enabled on MACsec capable 40Gb ports when they are run in 10Gb Ethernet mode.

Virtual Switch Bonding Licenses

SSA-EOS-VSB	S-Series SSA Virtual Switch Bonding License Upgrade, (For use on SSA Only)
S-EOS-VSB	S-Series Multi-slot chassis Virtual Switch Bonding License Upgrade, (For use on S130/S150 Class Modules and S140/S180 Class Modules without HW VSB ports)
S1-EOS-VSB	S-Series S1 Chassis Virtual Switch Bonding License Upgrade, (For use on S1-Chassis Only)

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. Please review the software release notes for your specific network.

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

SFP	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:

SFP	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**	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)

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

SFP	Description
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+
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, 10km 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+

40Gb Optics:

QSFP+ Optics	Description
40GB-SR4-QSFP	40Gb, 40GBASE-SR4, MM 100 m OM3, MPO QSFP+ Transceiver
40GB-ESR4-QSFP	40Gb, Extended Reach SR4, MM, 300m OM3, MPO QSFP+
40GB-LR4-QSFP	40Gb, 40GBASE-LR4, SM 10 km LC QSFP+ Transceiver
10326	40Gb, QSFP+ Parallel Single Mode (PSM), MPO connector, 10km SMF
10327	MPO to 4xLC SMF 10m patch cord (for use with 10326)
10329	40Gb, Bidirectional MMF, 100m OM3, QSFP+, LC
10334	40Gb, 40GBASE-LM4, MM OM3 140m, OM4 160m /SM 1 km LC QSFP+ Transceiver
10335	40Gb, 40GBASE-ER4, SM 40 km LC QSFP+ Transceiver
QSFP+ DAC	Description
40GB-C0.5-QSFP	40Gb, Copper DAC with integrated QSFP+ transceivers, 0.5 m
40GB-C01-QSFP	40Gb, Copper DAC with integrated QSFP+ transceivers, 1 m
40GB-C03-QSFP	40Gb, Copper DAC with integrated QSFP+ transceivers, 3 m
40GB-C07-QSFP	40Gb, Copper DAC with integrated QSFP+ transceivers, 7 m
40GB-F10-QSFP	40Gb, Active Optical DAC with integrated QSFP+ transceivers, 10 m
40GB-F20-QSFP	40Gb, Active Optical DAC with integrated QSFP+ transceivers, 20 m
10318	40Gb, Active Optical DAC with integrated QSFP+ transceivers, 100 m
10GB-4-C03-QSFP	10Gb, Copper DAC Fan out, 4xSFP+ to QSFP+, 3m
Adapters/Cables	
QSFP-SFPP-ADPT	QSFP+ to SFP+ Adapter (Note: The 10GB-LRM-SFPP transceiver is not supported and only MGBIC-LC01 and MGBIC-LC09 1Gb transceivers are supported.)

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)

See the Pluggable Transceivers data sheet for detailed specifications of supported transceivers.

NOTE:

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

NOTE:

Only the above listed Extreme Networks 40 Gigabit optical transceivers are supported by Extreme. Use of any other optical transceiver types results in a warning message.

Example Message for 40G cables that are unrecognized or unauthenticated

- System[1]port fg.1.4 contains an unauthenticated pluggable module('manufacturer'/'part no.')

Example message for unauthenticated 40G optical transceiver

- System[1]port fg.1.4 contains an unauthenticated pluggable module('manufacturer'/'part no.')

The S-Series will recognize a 10GB-4-xxx-QSFP cable when inserted in a QSFP+ port and reconfigure a QSFP+ port to 4 x 10 Gigabit Ethernet. A system reset is required for the port speed change to take effect

Example messages if the device installed in the QSFP+ port does not match the current configured mode:

- System[1]port tg.1.49 contains a 40GB MAU but is currently in 4x10GB mode and will remain down until system is reset
- System[1]port fg.1.1 contains a 4x10GB MAU but is currently in 40GB mode and will remain down until system is reset

QSFP-SFPP-ADPT transceiver support:

The 10GB-LRM-SFPP transceiver is not supported when plugged into a QSFP+ port via a QSFP-SFPP-ADPT. If an attempt is made to operate the transceiver the following error is logged:

```
port <port-name> will remain down because the pluggable module('<vendor>'/'<part-number>') is not supported and the port will remain operationally down.
```

The 10GB-LW-SFPP adapter is not supported in the QSFP-SFPP-ADPT adapter.

Using QSFP+ copper passive direct attach cables to interconnect S-Series/7100-Series and Summit/BlackDiamond systems:

When using any QSFP+ copper passive direct attach cable to connect S-Series/7100-Series QSFP+ ports to Summit/BlackDiamond QSFP+ ports, link will not come up unless auto-negotiation is disabled on the S-Series/7100-Series QSFP+ port.

Disable auto-negotiation on the S-Series/7100-Series 40Gb port: `set port negotiation fg.x.y disable`

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
DVMRPv3	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

Features		
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	Distributed Chassis Management (Single IP Address)	RIP ECMP, CIDR configuration
802.1D – 1998	SNMP v1/v2c/v3	Virtual Router Redundancy Protocol (VRRP)
802.1Q – Virtual Bridged Local Area Networking	Port Mirroring/Remote Port Mirroring	ICMP
GARP VLAN Registration Protocol (GVRP)	Flow Setup Throttling	Protocol Independent Multicast - Sparse Mode (PIM-SM)
802.1p – Traffic Class Expediting	MAC locking (Static/Dynamic)	Proxy ARP
802.1w – Rapid Reconfiguration of Spanning Tree	Node/Alias table	Basic Access Control Lists
802.1s – Multiple Spanning Trees	Policy-Based Routing	Extended ACLs
802.1t – Path Cost Amendment to 802.1D	SSH v2	Auto MDI-X Media Dependent Interface Crossover Detect (Enhanced for non auto negotiating ports)
802.3 – 2002	OSPF NSSA, equal cost multi-path	DHCP Server
802.3ad – Link Aggregation (128 users)	Audit trail logging	DHCP Relay w/option 82
802.3x – Flow Control	RADIUS Client	Jumbo Frame support
Load Share Network Address Translation (LSNAT)	FTP/TFTP Client	Directed Broadcast
Static Multicast Configuration	Telnet – Inbound/Outbound	Cisco CDP v1/2
Broadcast Suppression	Configuration File Upload/Download	CLI Management
Inbound and Outbound Rate Limiting	Text-based Configuration Files	DFE CPU and task Debugging
Strict and Weighted Round Robin Queuing	Syslog	RADIUS (Accounting, Snooping)
IGMP v1/v2/v3 and Querier support	Span Guard	Split RADIUS management and authentication
SMON Port and VLAN Redirect	RAD (Remote Address Discovery)	Link Flap detection
Spanning Tree Loop Protection	Cabletron Discovery Protocol (CDP)	Daylight Savings Time
TACACS+	NetFlow v5/v9	RFC 3580 with Policy support
Type of Service (ToS) Re-write	LLDP and LLDP-MED	Flex-Edge
NAT(Network Address Translation)	TWCB (Transparent Web Cache Balancing)	eBGP
iBGP	BGP Route Reflector	BGP 4 byte AS number
BGP Graceful Restart	BGP Route Refresh	BGP Extended Communities
Multi-VRF (IPv4/IPv6)	VRF-Aware NAT	VRF-Aware LSNAT
VRF-Aware TWCB	VRF-Aware Policy Based Routing	VRF-Aware DHCP Relay
VRF Static Route Leaking (IPv4/IPv6)	IPv6 Static Routing	IPv6 ACLs
IPv6 Policy Based Routing	IPv6 DHCP Relay	PIM-SSM
PIM-SSM v6	PIM-SM v6	RIPng
MLDv1/MLDv2	IPsec support for OSPFv3	IPv6 Node Alias Support
802.1Qaz ETS, (Data Center Bridging – Enhanced Transmission Selection)	802.3-2008 Clause 57 (Ethernet OAM – Link Layer OAM)	Virtual Switch Bonding (Like Chassis)

Features		
High Availability FW Upgrades	Fabric routing/ Fabric Routing with Host Mobility	IP Service Level Agreements
Tracked Objects	L3VPN over GRE	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	IP Source Guard	RADIUS Server Load Balancing
Routing as a Service (RaaS)	802.1aq-2012 Shortest Path Bridging (SPBv)	802.1Qau - Data Center Bridging, Congestion Notification (CN)
802.3az Energy Efficient Ethernet	MACsec IEEE802.1AE-2006 and 802.1X-2010	IEEE 802.1Qbb Priority Flow Control
VXLAN (RFC 7348)	Open Flow v1.3	

FIRMWARE CHANGES AND ENHANCEMENTS:

NOTICE: Minimum Firmware Revision Support Change

The SOK2209-0204 option module minimum firmware version has changed to 8.42.02.0012. You should upgrade to this level or later when using this module.

The ST8206-0848-F8, SG8201-0848-F8 and SK8208-0808-F8 fabric modules minimum firmware version has changed. You may need to upgrade to a supported firmware image to continue to use these modules.

Two High Availability Upgrade (HAU) firmware tracks have been updated to provide support:

- 8.22.01.0022 or later versions support the fabrics.
- 8.12.02.0006 or later 8.12.XX versions support the fabrics.

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-F11M	N 8.41.01

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 8.63.05.0004

Management Problems Corrected in 8.63.05.0004	Introduced in Version:
<p>Previous versions of firmware only support the following SSH Key Exchange algorithms (1024-bit keys):</p> <pre>diffie-hellman-group1-sha1 diffie-hellman-group-exchange-sha1</pre> <p>This release adds the following algorithm, which uses a 2048-bit key.</p> <pre>diffie-hellman-group14-sha1</pre> <p>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).</p>	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

MACsec Problems Corrected in 8.63.05.0004	Introduced in Version:
<p>When MACsec is enabled and Pre-Shared-Key (PSK) is configured with a "raw" CKN (i.e., octets rather than ASCII characters) and the CKN contains a 0x00 octet, then the port will use a truncated version of the CKN.</p> <p>This could result in non-matching CKNs, which would cause the MACsec connection to remain in the PENDING (i.e., blocking) state rather than reaching the SECURE (i.e., encrypting) state.</p> <p>The problem does not occur if the CKN consists entirely of printable ASCII characters, or if the</p>	8.41.01

CKN does not contain.	
Examples:	
"set macsec pre-shared-key port ge.1.1 ckn MyKey00 cak ..." is not affected (not raw)	
"set macsec pre-shared-key port ge.1.1 ckn raw 11223344 cak ..." is not affected (no zero octet)	
"set macsec pre-shared-key port ge.1.1 ckn raw 11220044 cak ..." is affected (3rd octet is zero)	

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

Hardware Problems Corrected in 8.63.05.0004	Introduced in Version:
Any I/O blades that have 1G of DRAM and 8 Switch Chips, will go into a reboot loop, leaving a message similar to the following in the message log:	8.63.04
<pre> Message 8/159 EDR Record 08.63.04.0003. 01/31/2019 11:12:55 Full Version: 08.63.04.0003 Severity/Facility: FATAL/KERNEL Task: tErfTask Injection Point: memPartLib.c:2853 Address: 0x00000000 memPartAlloc: block too big 1048576 bytes (0x1000 aligned) in partition 0x72dff30 Traceback Stack: 0x007e3c0c 0x007e2688 0x00424ab8 0x01be81c8 0x01be8368 0x01be845c 0x01088bc8 0x01c89ccc 0x01be64f4 0x01c1d45c 0x01c1da9c 0x01c2d308 </pre>	

<pre> 0x01c18e3c 0x01c19028 General Registers: msr : 0x0000b032 lr : 0x01be8148 ctr : 0x00000000 pc : 0x01be8148 cr : 0x88000288 xer : 0x20000000 r[0]:0xffffffff r[1]:0x075c4c10 r[2]:0x046413b0 r[3]:0x075c4c90 r[4]:0x01be8cf0 r[5]:0x00000000 r[6]:0x00000000 r[7]:0x00100000 r[8]:0x00001000 r[9]:0x00000000 r[10]:0x00000000 r[11]:0x075c4c8c r[12]:0x03769c60 r[13]:0x06155400 r[14]:0x00000000 r[15]:0x00000000 r[16]:0x00000000 r[17]:0x00000000 r[18]:0x00000000 r[19]:0x00000000 r[20]:0x00000000 r[21]:0x00000000 r[22]:0x00000000 r[23]:0x00000005 r[24]:0x00110001 r[25]:0x00000000 r[26]:0x00001000 r[27]:0x00100000 r[28]:0x00000000 r[29]:0x00000000 r[30]:0x01be8cf0 r[31]:0x072dff30 </pre>	
<p>When blades that have only 1G of RAM and 8 Switch Chips are being managed (SSH, telnet, SNMP, etc.), very infrequently one of the blades may reset and leave a message similar to the following in the message log:</p> <pre> ===== Message 5/316 EDR Record 08.63.03.0002 01/27/2019 05:11:50 Severity/Facility: INFO/REBOOT Task: tExcTask Injection Point: rebootLib.c:250 Address: 0x00000000 reboot(BOOT_NORMAL) - countdown and then reboot Traceback Stack: 0x007e3bfc 0x007e2678 0x00424ab8 0x01bd1040 0x01bcfaf0 0xe3e3e3e3e3e3e3e3 0xe3e3e3e3e3e3e3e3 </pre>	<p>8.63.03</p>

Problems Corrected in 8.63.04.0003

<p>IPV6 Forwarding Problems Corrected in 8.63.04.0003</p>	<p>Introduced in Version:</p>
<p>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.</p>	<p>7.00.01</p>

<p>Router Protocols Problems Corrected in 8.63.04.0003</p>	<p>Introduced in Version:</p>
-------------------------------------------------------------------	--------------------------------------

Syslog message like "Chassis memory limit of 498073600 bytes has been reached for routing protocols" is reported with multiple BGP peers, preventing BGP from reaching advertised route limits.	8.00.01
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------

New issue present in 08.63.04.0003 caused by fix for Router Protocols Problems

I/O blades reboot issue	Introduced in Version:
If the 8.63.04.003 image is loaded on a chassis that has I/O blades with 1G of SDRAM memory, and 8 switch chips present, those blades will go into a reboot loop caused by running out of SDRAM memory. To get out of the reboot loop, the running image will need to be back-revved. Any chassis with I/O blades should not have 08.63.04.0003 loaded on it.	8.63.04.003



Note: This issue does not exist for any Fabric blades regardless of SDRAM size and switch chip count.

To find out SDRAM size:

A "show system hardware" CLI command will display the SDRAM size and amount of Switch Chips present on the blade it is executed on.

Example:

The below output shows a SDRAM size of 1G:

```
SLOT 1 (Chassis 1 Physical Slot 1)
  Model:                ST4106-0248
  Part Number:          9404313
  Serial Number:        09395786636B
  Vendor ID:            1
  Base MAC Address:     00-1F-45-5C-17-F3
  MAC Address Count:    50
  Uptime:               0008,19:08:21 (d,h:m:s)
  Style:                2
  Hardware Version:     0
  Firmware Version:     08.63.05.0001.valk5C45EFCB
  BootCode Version:     01.01.00
  BootPROM Version:     01.01.05
  CPU Version:          28674 (PPC 750GX)
  SDRAM:             1024 MB
  NVRAM:                32 KB
```

The amount of switch chips present is a sum of switch chips on all NIM's:

```
NIM[1]:
  Chassis 1 location:   lower right
  Description:          24 Port 10/100/1000 RJ45, 1X, Double-Wide
  Bottom, PoE+ Capable
  Board Revision:       10      (0xA)
  PLD Revision:         15      (0xF)
  FRU:                  no
  PoE[1]:
```

Software Revision:	Unavailable
Device Id:	Unavailable
SWITCH CHIP[0]:	
Type:	ASIC
Revision:	3327 (0xCFF)
Id:	2
SWITCH CHIP[1]:	
Type:	ASIC

Mirroring Problems Corrected in 8.63.04.0003	Introduced in Version:
On a PV-FC-180, when a Tunnel Bridge Port is used to forward packets across a Remote Tunnel Mirror, non-IP frames that are mirrored will have an incorrect Total Length field in the IP header portion of the GRE encapsulation. This may lead to intermediate routers or the Tunnel destination dropping the frame as an invalid frame.	8.20.02

Problems Corrected in 8.63.03.0002

Platform Problems Corrected in 8.63.03.0002	Introduced in Version:
In VSB chassis, port out discard counters may be incremented for incorrect port.	8.11.01

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</pre>	7.00.01

[3] 0x01be6044

[4] 0x01087080

[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:

mshr: 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
r[16]:0x2fb4d9c0 r[17]:0x0381e158 r[18]:0x20dce048 r[19]:0x1f951818
r[20]:0x0000005e r[21]:0x01cf2ed0 r[22]:0x00000000 r[23]:0x00000003
r[24]:0x00000020 r[25]:0x00000010 r[26]:0x00000000 r[27]:0x01d02328
r[28]:0x6b40a010 r[29]:0x00000d00 r[30]:0x6b40a010 r[31]:0x072ce6c0

```

In the message log.

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:
IPv4 DHCP Relay packet sent to server with UDP SRC port of 68 instead of 67. Added VRF and interface cli command to set the SRC port to 67: ip dhcp relay information option port-67	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

Hardware Problems Corrected in 8.63.02.0004	Introduced in Version:

On SOGK2218-0212 NIM, when a swap of transceivers occurs in ports 1 or 2 that requires a change of speed between 1G and 10G, the remaining ports may require a disable/enable cycle to re-establish link.	7.72.01
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------

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

MACsec Problems Corrected in 8.63.02.0004	Introduced in Version:
Valid ingress MACsec data packets will be dropped as "Not Valid" if the MACsec peer's Secure Channel Indicator (SCI) contains a Port Id higher than "15". Egress MACsec traffic is not affected. Extreme Networks S-Series and ToR7100 switches always set their SCI Port Id to "1", so they are not affected by this defect. However, Extreme Networks X460-G2 switches use Port Id's other than "1". Also, 3rd party MACsec-capable equipment might also set Port Id to something other than "1".	8.41.01
Due to an ambiguity in IEEE802.1X-2010, vendors have implemented a certain aspect of the MACsec Key Agreement protocol (MKA) in two different ways. The author of the spec addressed the issue in September 2017: http://www.ieee802.org/1/files/public/docs2017/xck-seaman-mka-pn-exhaustion-0917-v1.pdf The S-Series and ToR7100-Series put the most recent key information in the 'Latest Key' fields of the MKPDU SAK Use parameter set, while some other vendors put the most recent key information in the 'Old Key' fields. The S and ToR are incompatible with this other version of the SAK Use parameter set, which manifests itself as the MKA connect status perpetually bouncing between SECURE and PENDING. As suggested in the above document, the S and ToR will continue to encode the most recent key information in the 'Latest Key' fields but they will now be "permissive on receive". This means S and ToR will accept latest key information in either the 'Latest Key' or 'Old Key' fields.	8.41.01
S-Series and ToR7100-Series were sending Live Peer List and Potential Peer List parameter sets in every MACsec PDU (MKPDU), even when the lists were empty. For compatibility with other vendor's MKA implementations and for compliance with IEEE802.1X-2010, the S and ToR will no longer send empty Peer List parameter sets. This change is backwards compatible with older S and ToR images.	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

Netflow Problems Corrected in 8.63.01.0019	Introduced in Version:
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: - 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. 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.	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.	8.41.01

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).	
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

Chassis Problems Corrected in 8.63.01.0019	Introduced in Version:
"Show system hardware"™ might not display power supply fan speed for one of the bonded S180 SSAs.	7.80.00

Mirroring Problems Corrected in 8.63.01.0019	Introduced in Version:
On a PVFC-180 device, if more then the supported 2 Tunneled Bridge Ports are attempted to be used as a mirror destination, a message similar to: 3>Mirror[1.tMirrMgr]commitToHw: Failed to get available relkup port for tunnel tbp.0.3 will be displayed, and then later the device may reset leaving a message similar to: 1>DistServ[1.tDsBrdOk]serverWatchDog.1(Config), client 88(mirrorMgr) inrecv for 6456 tics (0x00ff8d14 0x0084b2b0 0x00840478 0x01c85d64 0xeeeeee) in the message log. Now the following message will be displayed: Unable to set all requested mirror destination ports successfully And the device will not reset.	8.41.01

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

Packet Dispatch Problems Corrected in 8.62.03.0006	Introduced in Version:
Some protocol frames may use a different rate-limiter depending if the corresponding IP Interface is operUp or operDown.	8.31.02
If there are virtual routers configured, the router may not respond to any more than 100 ARP requests a second with ARP replies. The 100 a second would be the sum across all virtual routers.	7.30.01

Boot Config Problems Corrected in 8.62.03.0006	Introduced in Version:
<p>While syncing of modules with different configurations of "ip helper-address" or "ipv6 dhcp relay destination" the resulting configuration may result in the elimination of one of the configuration lines or in a corrupted configuration that displays something like this:</p> <pre> ip helper-address vrf </pre> <p>or</p> <pre> ipv6 dhcp relay destination 2001:67c:2d44:1110::547:2 vrf </pre>	8.11.01

Routing Protocols Problems Corrected in 8.62.03.0006	Introduced in Version:
When OSPF is configured with cryptographic authentication and a small hello timer interval, OSPF may drop hello packets due to cryptographic out-of-sequence errors.	8.22.01

IP Stack Problems Corrected in 8.62.03.0006	Introduced in Version:
When the router generates an ICMP Unreachable, Time Exceeded, Redirect, or Address Mask Request message, very infrequently the IP addresses in the generated packet will not be the ones that should be used.	7.00.01

Problems Corrected in 8.62.02.0022

IPv6 Forwarding Problems Corrected in 8.62.02.0022	Introduced in Version:
If DHCP for IPv6 is configured on an interface and VRRP is also configured with an IPv6 address, DHCP relay will now use the interface IPv6 address as the source IP in the IPv6 relay packet.	8.11.01

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

Routing Protocols Problems Corrected in 8.62.02.0022	Introduced in Version:
At times NSSA ABR's translating Type 7 to Type 5 LSA loses the Forwarding address.	7.00.01

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:
<p>Add a new command to display the underlying port of a lag :</p> <pre>show lacp underlyingPort <lagPortString> { da-sa { <da-mac> <sa-mac> } dip-sip { <dip> <sip> }}</pre> <p>Example:</p> <pre>show lacp underlyingPort lag.0.20 dip-sip 172.22.71.2 169.254.2.2</pre> <p>Port [ge.2.20] will be used for that flow.</p> <p>If on a bonded system each chassis is analyzed for the underlying Port:</p> <pre>show lacp underlyingPort lag.0.20 dip-sip 172.22.71.2 169.254.2.2</pre> <p>Port [ge.2.20] will be used for that flow if ingressing chassis 1.</p> <p>Port [ge.2.21] will be used for that flow if ingressing chassis 2.</p>	7.00.01
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

PoE Problems Corrected in 8.62.02.0022	Introduced in Version:
PoE power budget might vary and messages similar to the following might occasionally be displayed: System[4]PoE power supply 1 has been removed. System[4]PoE power supply 1 has been installed. System[4]PoE power supply 1 is operational.	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

UDP Forwarding Problems Corrected in 8.62.02.0022	Introduced in Version:
If UDP forwarding is configured on an interface that is also the destination for an IP Directed Broadcast UDP packet, UDP forwarding will now occur.	8.42.04

Hardware Problems Corrected in 8.62.02.0022	Introduced in Version:
MGBIC-02 will not link when configuring from 1G to 100Mb on certain supplier part number (FINISAR FCLF-8521-3-EN).	8.61.02

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
BFD Problems Corrected in 8.62.01.0034	Introduced in Version:
A failover of the host master results in the BFD control packets not being transmitted to the peer after a new host master is selected. The BFD sessions on the device where the host master had a failover show the BFD sessions up. However, the peer shows the BFD sessions down.	8.31.02
The slave blades were not updating the BFD Min Rx Interval, nor indicating that the echo-mode was activated by the master. Therefore, when a slave blade becomes the master during a failover, the BFD timing information is out-of-sync with the peer, and either the peer or the new master may cause the BFD session to transition to the DOWN state. This causes the routing protocols relying on BFD to bounce their neighbor.	8.31.02
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

Hardware Problems Corrected in 8.62.01.0034	Introduced in Version:
<p>Some blades, that have only 1G of RAM, may go to reset loop, logging a message similar to the following:</p> <pre>===== Message 125/208 EDR Record 08.61.01.0019 06/09/2016 12:17:38 Severity/Facility: FATAL/KERNEL Task: tDot1x Injection Point: memPartLib.c:2853 Address: 0x00000000 memPartAlloc: block too big 65536 bytes (0x10 aligned) in partition 0x72c73e0</pre>	8.61.01
ICMP Problems Corrected in 8.62.01.0034	Introduced in Version:
The help for ping and traceroute (interface) options are not descriptive.	7.01.02
IPv4 Forwarding Problems Corrected in 8.62.01.0034	Introduced in Version:
Recognizing an IP address as both source and destination address may incorrectly produce an ICMP error.	8.32.01
Link Aggregation Problems Corrected in 8.62.01.0034	Introduced in Version:
A LAG with 10 gig ports used as 1 gig may go down and affect higher end protocols under the situation where the lowest numbered port of a LAG rejoins the LAG after a down event.	1.07.19
MACsec Problems Corrected in 8.62.01.0034	Introduced in Version:
The 'show macsec' command can take up to a minute to complete on a large switch, such as bonded S8's. Command execution time has been reduced, and is now on the order of 10 seconds for bonded S8's. The command's syntax and output remain the same.	8.41.01
<p>Adding support for the following IEEE802-SECY-MIB (i.e., MACsec) counters:</p> <pre>SecyRxSCStatsEntry secyRxSCStatsUnusedSAPkts Counter64 secyRxSCStatsNoUsingSAPkts Counter64 secyRxSCStatsNotValidPkts Counter64 secyRxSCStatsInvalidPkts Counter64 secyRxSCStatsOKPkts Counter64 SecyTxSCStatsEntry secyTxSCStatsProtectedPkts Counter64 secyTxSCStatsEncryptedPkts Counter64</pre> <p>These newly supported counters are accessible via SNMP or via the CLI commands "show macsec secy stats <port-string>" and "show macsec all <port-string>".</p>	8.41.01
MACsec Secure Connectivity Association Keys (CAKs) can be configured via CLI as either a raw value (i.e., a 128-bit secret key) or as a secret passphrase (which is hashed into a 128-bit secret key).	8.41.01
If the CLI user chooses to enter a secret passphrase in interactive non-echo mode, the CLI will	

MACsec Problems Corrected in 8.62.01.0034	Introduced in Version:
now prompt the user to re-enter the passphrase and will continue with the command only if both passphrases match.	
MACsec Secure Connectivity Association Keys (CAKs) can be configured via CLI as either a raw value (i.e., a 128-bit secret key) or as a secret passphrase (which is hashed into a 128-bit secret key).	8.41.01
The maximum passphrase length allowed by CLI has been increased from 16 to 64 characters. Longer passphrases increase the entropy of passphrase generated CAKs.	

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:
S-140 and S-180 class blades, in hardware VSB or non-VSB chassis, exhibit the following ingress port mirror issues: - For a many to one port mirror, if mirroring is disabled on any port or any of the source mirror ports go operationally down, ingress port mirroring will cease to function on all ports in the mirror - For a mirror with a LAG as a source port, if any underlying port in the LAG detaches from the LAG, ingress port mirroring on the entire LAG will cease to function.	8.11.01
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

MSDP Problems Corrected in 8.62.01.0034	Introduced in Version:
Misconfiguring MSDP peers may cause the system to reboot.	8.01.01

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
OpenFlow Problems Corrected in 8.62.01.0034	Introduced in Version:
Core may result if more than 50 groups of type all are created.	8.61.01
A system limit for the number of available groups of type all was not available.	8.61.01
The debug command "rtr flow debug table all print" would result in a reset. The work-around is to specify the tables individually instead of using the "all" keyword.	8.61.01
Certain distributions of values in match fields could result in excessive search times.	8.61.01
PoE Problems Corrected in 8.62.01.0034	Introduced in Version:
When a PoE-capable SSA is bonded with an SSA that is not PoE-capable (or PoE-capable S1 is bonded with S1 that is not PoE-capable), the "inlinepower" configuration may be missing from the command.	7.40.01
Port Status/Control Problems Corrected in 8.62.01.0034	Introduced in Version:
<p>Messages similar to the following may unexpectedly be logged on S180 I/O Fabric Module and S180 I/O Module assemblies:</p> <pre><164>Jun 6 14:07:27 10.50.65.155 NIM[6.tNimErrMon]MAC 12 port 6 transmitted 699995 malformed packet(s) in last 60 seconds (total = 699995) <164>Jun 6 14:07:27 10.50.65.155 NIM[6.tNimErrMon]MAC 12 port 8 transmitted 904537 malformed packet(s) in last 60 seconds (total = 904537) <164>Jun 6 14:08:27 10.50.65.155 NIM[6.tNimErrMon]MAC 12 port 6 transmitted 1099510927781 malformed packet(s) in last 60 seconds (total = 1099511627776) <164>Jun 6 14:08:27 10.50.65.155 NIM[6.tNimErrMon]MAC 12 port 8 transmitted 1099510723239 malformed packet(s) in last 60 seconds (total = 1099511627776)</pre> <p>The messages indicate that an internal error may have occurred when reading counter values from switch hardware. Normal packet switching/routing is not affected.</p>	8.32.01
When MGBIC-02 is installed into a disabled port it will not go operational when enabled.	8.61.01
If a SOV3208-0202 option module is installed on a SK8208-0808-F8 fabric module, and one or more of the 10G ports on the fabric module contain 1G SFPs, then link may not be properly established on the ports with SFPs following system initialization.	8.42.01

Radius Problems Corrected in 8.62.01.0034	Introduced in Version:
The RADIUS client may fail initialization at boot time. The following message appears in the log: <3>Radius[2.tusrApplnit]Unable to initialize master socket database	8.41.01
Routing Problems Corrected in 8.62.01.0034	Introduced in Version:
L3VPN between directly connected PE routers without a tunnel in the core is no longer supported.	8.21.01
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
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: 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] 0xecececece 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	8.31.01

VRF Problems Corrected in 8.62.01.0034	Introduced in Version:
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 -----	

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

Platform Problems Corrected in 8.61.02.0001	Introduced in Version:
Blades with 1GB of memory may go into a reset loop.	8.61.01

Features Enhancements 8.61.01.0018

OpenFlow Enhancement in 8.61.01.0018
Support for OpenFlow 1.3 has been added to S150A and S180 class SSAs, as well as S1 chassis with an S180 or S155 class fabric I/O module. OpenFlow is used as a southbound protocol in some Software Defined Networking (SDN) architectures.
Running OpenFlow requires the S-Series to run in a new operational mode which limits the scale of traditional protocols and forwarding behavior to free space for OpenFlow to run. This operational mode uses “hybrid mode” forwarding behavior where OpenFlow takes precedence for forwarding decisions. For each packet, the switch first looks for a match in the OpenFlow rule set, and if one is not present, traditional forwarding behavior is used.

Transceiver Enhancements in 8.61.01.0018
Added support for 10G tunable DWDM single mode SFP+ transceiver (part number 10325).
Added support for QSFP+ BiDi transceiver (part number 10329).
Added support for QSFP+ LM4 transceiver (part number 10334).
Added support for QSFP+ ER4 transceiver (part number 10335).

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 ip access-group 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
The time required to apply ACLs with append is unexpectedly long.	8.32.01

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

Hardware Problems Corrected in 8.61.01.0018	Introduced in Version:
Automatic upgrades of the power controllers on SL8013-1206-F8A and SL8013-1206A is not supported for all hardware versions. Warning messages similar to the following may appear during initialization: "<165>Mar 24 09:50:25 0.0.0.0 NIM[2.tRootTask]NIM 3 power controller cannot be read or is out of date [0x37551b00]"	8.41.01

LS-NAT Problems Corrected in 8.61.01.0018	Introduced in Version:
Increased Reframer Database to support NAT 128k bindings.	8.42.01

MACsec Problems Corrected in 8.61.01.0018	Introduced in Version:
Executing <i>show port status</i> or <i>show port operstatus</i> on MACsec ports (*U.* or *C.*) may falsely report operational when the common port has no link or is admin down.	8.41.01
You can enable IEEE802.1X-2010 Message Key Agreement (MKA) protocol on a MACsec-capable port and then disable encryption by setting the IEEE8021X-PAE-MIB object 'ieee8021XKayMacSecDesired' to FALSE. Although MKA without MACsec encryption is a valid mode of operation per IEEE802.1X-2010, the intended usage for the MacSecDesired object is to be read-only and always TRUE. This ensures that data protection and integrity cannot be compromised by a mis-configuration.	8.41.01
The MACsec set <i>macsec secy window <num-packets> <port-string></i> command was limiting the maximum window size to 65,536 packets. The command now accept the entire window range (0–4,294,967,295) as defined by 'secyIfReplayProtectWindow' in IEEE8021-SECY-MIB.	8.41.01
The IEEE8021X-PAE-MIB object 'ieee8021XPaeSysAnnouncements' was accepting SNMP set TRUE operations; however, announcements are not supported, so SNMP writes to this object are now rejected with reason 'notWritable'.	8.41.01
The IEEE8021X-PAE-MIB object 'ieee8021XPaeSysAccessControl' was accepting SNMP set TRUE operations; however, global enabling and disabling of MKA and MACsec is not	8.41.01

MACsec Problems Corrected in 8.61.01.0018	Introduced in Version:
supported. MKA and MACsec can only be enabled and disabled on a per-port basis. SNMP writes to the SysAccessControl object are now rejected with reason 'notWritable'.	

Management Problems Corrected in 8.61.01.0018	Introduced in Version:
Over time there may be a loss of available Telnet or SSH connections.	7.00.01
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 etsysConifgMgmtChangeDelayTime 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: <pre>ssh <username>@<hostname> [<command> <arg1> <arg2> ...]</pre> <p>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.</p>	7.00.01

Multicast Problems Corrected in 8.61.01.0018	Introduced in Version:
Some SSA-G8018-0652 units may exhibit high CPU utilization for no apparent reason, hindering access to the serial console.	8.01.01

NAT Problems Corrected in 8.61.01.0018	Introduced in Version:
At times a NAT binding may become 'stuck' causing NAT to report "duplicate bindings" and producing binding create failures, and causing high CPU utilization. To allow processing to continue, when this issue is detected, force a delete of the 'stuck' binding.	7.03.06

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

Tunneling Problems Corrected in 8.61.01.0018	Introduced in Version:
A memory leak occurs when displaying tunnel probe configuration.	8.32.01

Transceiver Problems Corrected in 8.61.01.0018	Introduced in Version:
The command <code>debug sfp show baseinfo</code> displays the 10GBASE-LR as 10GBASE-ER.	8.32.01H1
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 <code>"show port egress -v"</code> .	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
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:	8.31.01

VRF Problems Corrected in 8.42.05.0003	Introduced in Version:
<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] 0xe0000000 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>	

Problems Corrected in 8.42.04.0016

ACL Problems Corrected in 8.42.04.0016	Introduced in Version:
There is a delayed response to applying ACL configuration.	8.32.01
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:
Subnet broadcast packets that also have a Layer 2 broadcast destination address are not forwarded.	7.63.01
Hardware Problems Corrected in 8.42.04.0016	Introduced in Version:
Automatic upgrades of the power controllers on SL8013-1206-F8A and SL8013-1206A is not supported for all hardware versions.	8.41.01

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-CONFIGURATION-MANAGEMENT-MIB's <i>etsysConifgMgmtChangeDelayTime</i> object) cause this problem.	7.00.01
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
MACsec Problems Corrected in 8.42.04.0016	Introduced in Version:
A <i>show port status</i> or <i>show port operstatus</i> command on MACsec ports (*U.* or *C.*) may falsely report operational when the common port has no link or is administratively disabled.	8.41.01
Management Problems Corrected in 8.42.04.0016	Introduced in Version:
Gradually, there may be a loss of available Telnet or SSH connections causing connections to be denied.	7.00.01
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> 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.	7.00.01
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
MSDP Problems Corrected in 8.42.04.0016	Introduced in Version:
Misconfiguring MSDP peers may cause the system to reboot.	8.01.01

Port Status/Control Problems Corrected in 8.42.04.0016	Introduced in Version:
<p>Messages similar to the following may unexpectedly be logged on S180 I/O Fabric Module and S180 I/O Module assemblies:</p> <pre><164>Jun 6 14:07:27 10.50.65.155 NIM[6.tNimErrMon]MAC 12 port 6 transmitted 699995 malformed packet(s) in last 60 seconds (total = 699995) <164>Jun 6 14:07:27 10.50.65.155 NIM[6.tNimErrMon]MAC 12 port 8 transmitted 904537 malformed packet(s) in last 60 seconds (total = 904537) <164>Jun 6 14:08:27 10.50.65.155 NIM[6.tNimErrMon]MAC 12 port 6 transmitted 1099510927781 malformed packet(s) in last 60 seconds (total = 1099511627776) <164>Jun 6 14:08:27 10.50.65.155 NIM[6.tNimErrMon]MAC 12 port 8 transmitted 1099510723239 malformed packet(s) in last 60 seconds (total = 1099511627776)</pre> <p>The messages indicate that an internal error may have occurred when reading counter values from switch hardware. Normal packet switching/routing is not affected.</p>	8.32.01
<p>If a SOV3208-0202 option module is installed on a SK8208-0808-F8 fabric module, and one or more of the 10G ports on the fabric module contain 1G SFPs, then link may not be properly established on the ports with SFPs following system initialization.</p>	8.42.01

PoE Problems Corrected in 8.42.04.0016	Introduced in Version:
<p>When a PoE-capable SSA is bonded with an SSA that is not PoE-capable (or PoE-capable S1 is bonded with S1 that is not PoE-capable), the "inlinepower" configuration may be missing from the command.</p>	7.40.01

RADIUS Problems Corrected in 8.42.04.0016	Introduced in Version:
<p>The RADIUS client may fail initialization at boot time. The following message appears in the log:</p> <pre><3>Radius[2.tusrApplnit]Unable to initialize master socket database</pre>	8.41.01

VSB Problems Corrected in 8.42.04.0016	Introduced in Version:
<p>Rarely, a software VSB chassis may have low throughput between VSB bonded chassis.</p>	7.40.00

Problems Corrected in 8.42.03.0006

802.1Q Relay Problems Corrected in 8.42.03.0006	Introduced in Version:
<p>A large portion of flooded frames are dropped when running 8.40 or later on S140 and S180 hardware.</p>	8.41.01

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 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 <code>set time</code> command.	8.31.02

Host Services Problems Corrected in 8.42.03.0006	Introduced in Version:
Running CLI commands displaying router configuration might cause the host management to become unresponsive/locked.	8.20.02
When transferring a new image using FTP, CLI prompt become unresponsive and "Waiting for chassis to distribute image to compatible modules..." message appears.	8.42.02

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 wrong VLAN / ES destination.	8.31.01

LSNAT Problems Corrected in 8.42.03.0006	Introduced in Version:
Occasionally, when processing NAT-controlled FTP control packets, the system deadlocks and creates a watchdog timer event.	8.22.03

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 noin/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

VSB Problems Corrected in 8.42.03.0006	Introduced in Version:
If software chassis bonding is enabled on an SFP+ port of an option module in an S140/S180 assembly, 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
A blade may reset on a software VSB chassis shortly after boot.	7.40.00
When soft VSB bonding is used while receiving an MPLS flow, and the MPLS label is swapped and sent out the software bonding port, the software bonded header is not added.	8.41.01

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
File Management Problems Corrected in 8.42.02.0012	Introduced in Version:
On larger bonded systems (greater than 8 blades) image distribution to some of the blades may fail, producing the following log messages: <164>Sep 1 08:05:43 100.10.10.22 FileMgr[16.tlimageSync]Data connection error (4) for (13:./images/03) <163>Sep 1 08:05:42 100.10.10.22 FileMgr[7.tlimageSync]downloadImage: Failed checksum <163>Sep 1 08:05:42 100.10.10.22 FileMgr[12.tlimageSync]downloadImage: Failed remote open(12:./images/03); errno = 0x00380003 [S_dosFsLib_FILE_NOT_FOUND]	7.60.01
MACsec Problems Corrected in 8.42.02.0012	Introduced in Version:
If MACsec is configured and enabled on a 10GBase-T port operating at 10Gbps, the secure MACsec connection with the peer may go down if the port speed is reduced to 1Gbps. To use MACsec on such a port, the corresponding blade assembly must be initialized with the port operating at 1Gbps.	8.41.01
When MACsec mode is enabled, MKPDU protocol packets are occasionally lost. If two or more consecutive protocol packets are lost, the secure connection temporarily disconnects as indicated by the following messages: <166>Aug 27 14:19:30 192.168.1.90 MACsec[4][tg.4.204] KaY: Rx MKPDU: MKA Peer List (1) contains my MI but wrong MN (Acutal:6675, Expected:6677) <166>Aug 27 14:19:30 192.168.1.90 MACsec[4][tg.4.204] KaY: Discarding Rx MKPDU: A live peer did not include me in their PEER-LIST <166>Aug 27 14:19:32 192.168.1.90 Last message repeated 1 time <165>Aug 27 14:19:34 192.168.1.90 MACsec[4][tg.4.204] KaY: Live peer removed (no valid MKPDUs rcv'd in last 6.16 seconds) <166>Aug 27 14:19:34 192.168.1.90 MACsec[4][tg.4.204] Connecting PENDING: auth(0) secure(0) fail(0), unauthallowed(never) unsecuredallowed(mkaServer) <165>Aug 27 14:19:34 192.168.1.90 LinkTrap[4]Interface tgC.4.204 is Down. <165>Aug 27 14:19:34 192.168.1.90 LinkTrap[4]Interface tg.4.204 is Down.	8.41.01
If MACsec is being used on a 10GBase-T port and the port is subjected to continual reconfiguration (enabling/disabling MACsec), a message in the form "KaY: Rx MKPDU: MKA Peer List (2) contains my MI but wrong MN" may appear and the secure connection with the peer is dropped. You must restart of the port's blade assembly to restore proper operation.	8.41.01
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

Management Problems Corrected in 8.42.02.0012	Introduced in Version:
<p>The following blades: ST1206-0848 SG1201-0848 ST1206-0848-F6 SG1201-0848-F6 when installed with following option modules: SOK1208-0204 SOK2208-0204 SOK2209-0204 SOGK2218-0212 SOTK2268-0212 are vulnerable to running out of memory and resetting with a message similar to:</p> <pre> ===== Message 10/267 EDR Record 08.42.01.0008 12/07/2015 16:59:13 Severity/Facility: FATAL/KERNEL Task: tSshM1 Injection Point: memPartLib.c:2853 Address: 0x00000000 memPartAlloc: block too big 65536 bytes (0x10 aligned) in partition 0x68b03c0 Traceback Stack: 0x006ca5cc 0x006c9048 0x00424ab0 0x01ab11fc 0x01ab139c 0x01ab19a8 0x00f579a8 0x01a7c9a8 0x01a7c940 0x0078424c 0x00784cf8 0x005c5164 0x005c5930 0x027c1020 General Registers: msr : 0x0000b032 lr : 0x01ab117c ctr : 0x00000000 pc : 0x01ab117c cr : 0x88008288 xer : 0x20000000 r[0]:0xffffffff r[1]:0x3df397f0 r[2]:0x0436f2b0 r[3]:0x3df39870 r[4]:0x01ab1d24 r[5]:0x00000000 r[6]:0x00000000 r[7]:0x00010000 r[8]:0x00000010 r[9]:0x00000000 r[10]:0x00000000 r[11]:0x3df3986c r[12]:0x03521580 r[13]:0x05e75450 r[14]:0x00000002 r[15]:0x1fd13cc0 r[16]:0x00000001 r[17]:0x035dbcf4 r[18]:0x00000000 r[19]:0x1abad010 r[20]:0x3df9f060 r[21]:0x1fd13cc0 r[22]:0x00000000 r[23]:0x1c0b7ce0 r[24]:0x00110001 r[25]:0x00000000 r[26]:0x00000010 r[27]:0x00010000 r[28]:0x00000000 r[29]:0x00000000 r[30]:0x01ab1d24 r[31]:0x068b03c0 ===== </pre>	<p>8.42.01</p>

Management Problems Corrected in 8.42.02.0012	Introduced in Version:
They are especially vulnerable to this after issuing the commands <i>show config</i> and <i>show support</i> .	

Multicast Problems Corrected in 8.42.02.0012	Introduced in Version:
PIM-DM: After configuring <i>pim dense-mode</i> and rebooting, the PIM operating mode is restored as <i>sparse-mode</i> .	8.41.01
PIM-SM IPv4: After configuring <i>ip pim graceful-restart</i> and rebooting, the <i>graceful-restart</i> 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

Routing Problems Corrected in 8.42.02.0012	Introduced in Version:
BGP: After configuring <i>bgp maxas-limit</i> and rebooting, the <i>maxas-limit</i> setting is not restored.	8.41.01

Feature Enhancements 8.42.01.0007

NAT/LSNAT/TWCB Bindings Enhancement in 8.42.01.0007
The number of global NAT bindings available on S-Series, including LSNAT and TWCB, increases to 128K from 64K.

Third-Party Optics Enhancement in 8.42.01.0007
Use of third-party 40Gb optical transceivers is no longer restricted. A message is generated identifying ports where unsupported transceivers are in use. For examples of messages, see System Behavior section.

Port Speed Enhancement in 8.42.01.0007
100Mb operation is supported on MGBIC-02 transceivers purchased after January 1, 2013, and labeled "Finisar."
If you attempt to set 100Mb speed on an older version of the MGBIC-02 that does not support 100Mb, the following message appears:
S3-A Chassis(su)->set port spe ge.1.7 100 default speed 100 mbps not supported on port ge.1.7.

Port Usage Enhancement in 8.42.01.0007
A physical port is no longer reserved when configuring a remote port mirror on the Purview Appliance (PV-FC-180).

Problems Corrected in 8.42.01.0007

File Management Problems Corrected in 8.42.01.0007	Introduced in Version:
On multi-module chassis or bonded systems, remote modules may boot in a state where they do not respond to remote procedure calls. This effects commands such as "dir" or "show file" from remote modules.	8.41.01

Host Problems Corrected in 8.42.01.0007	Introduced in Version:
On multi-module chassis or bonded systems, remote modules may boot in a state where they do not respond to remote copy requests. The following message appears in the log: <165>Jul 23 02:34:13 1.1.1.1 System[1]Requesting a copy of the non-volatile store for slot 2. <164>Jul 23 02:34:56 1.1.1.1 Default[2.tNvBulk]dfeNfsMountNonvol: Unable to communicate with remote slot(1); ip=127.0.3.1, exporting=/flash1/nonvol, path=/nvNfsRem.001; errno=3155732	8.41.01
entPhySensorValue corresponding to ambient-temp-sensor-1 might not reflect current ambient temperature.	7.60.01

MACsec Problems Corrected in 8.42.01.0007	Introduced in Version:
When MACsec is enabled, chassis might experience coherency issues and error messages might appear. These messages might initiate a system reboot.	8.41.01

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. NetSight Compass relies on the ctAliasMacAddressTable to display all node and alias entries, thus Compass may not function properly.	8.01.01
When MIB walking ctAliasMIBAddress table, occasionally IP entries with the invalid address of 0.0.0.0 might be returned.	8.01.01
Node and alias entries that correctly appear on ports that they are received on also appear incorrectly 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

OSPF Problems Corrected in 8.42.01.0007	Introduced in Version:
An OSPF network that is an exact match for the range configured in the following command, "area <areaid> range <network> <mask> not-advertise", is summarized into other areas.	7.00.01

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 because 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:
Large RIP/RIPng packets do not cause a "too big" message to be sent to the source when sending over L2GRE or VXLAN tunnels.	8.21.01
You cannot configure more than 62 remote VXLAN VTEP IP addresses in aggregate on a single switch.	8.41.01
VSB Problems Corrected in 8.42.01.0007	Introduced in Version:
Heartbeat transmit errors to remote bonded slot may occur during initialization. When this happens, a message similar to the following appears: "<3>FtmLj[5.tHBChk]heartBeatCheck: Transmit errors(10) to slot 12 are preventing heartbeat checks."	7.70.00
If software VSB is enabled on either (1) pre-S140 S-Series chassis module assemblies that include SFP+ ports, or (2) S140/S180 assemblies that include option module(s) with SFP+ ports, an SFP module present in one or more of the SFP+ ports prior to system initialization may fail to properly link up. To resolve this problem, remove, and then reinsert, each such SFP module after the system has booted.	8.41.01

Features Enhancements 8.41.01.0004

MACsec Enhancements in 8.41.01.0004
MACsec is defined by IEEE802.1AE-2006 and 802.1X-2010 and can be used to provide hardware-based point-to-point link layer security using authentication and encryption with pre-shared key exchange between two MACsec capable devices. S-Series MACsec capability is hardware dependent and requires a license; see the hardware support table for details.
New Licenses Enhancements in 8.41.01.0004
MACsec 1Gb module/SSA/uplink License: In support of the MACsec feature set, the S-EOS-MACSEC license is required per module to enable MACsec for 1Gb modules, all capable uplink modules and capable SSA ports. See the MACsec license table for details. MACsec 10Gb module License: In support of the MACsec feature set the S-EOS-KMACSEC license is required per module to enable MACsec for 10Gb S140 and S180 modules. See the MACsec license table for details.
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
The switch now allows AES CTR Ciphers. -----

Cryptography Enhancements in 8.41.01.0004

The allowed ciphers and allowed MACs lists used by the switch's SSH Client and SSH Server are hardcoded as follows:

Ciphers:

aes128-cbc,aes192-cbc,aes256-cbc,3des-cbc,none,blowfish-cbc,
cast128-cbc,rijndael-cbc@lysator.liu.se

MACs:

hmac-sha1-96,hmac-sha1,hmac-md5,hmac-md5-96,hmac-ripemd160, hmac-ripemd160@openssh.com

One (1) cipher has been removed from SSH:

none ("none" cipher is used to bypass encryption)

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

Five (5) new Encrypt-then-MAC (ETM) MACs have been added to SSH:

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 via 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 (revert to default ciphers list)

clear ssh macs (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*,

Cryptography Enhancements in 8.41.01.0004

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*

Priority Based Flow Control (PFC) Enhancements in 8.41.01.0004

IEEE 802.1Qbb Priority Flow Control (PFC) enhances existing Data Center Bridging (DCB) functionality to allow for End-End congestion management for protocols that are capable of transmission rate limiting to avoid frame loss. Priority Based Flow Control (PFC) is hardware dependent and only supported on S140 and S180 hardware platforms.

QSFP+ Parallel Single Mode (PSM) Transceiver Enhancements in 8.41.01.0004

QSFP+ 40Gb Parallel Single Mode (PSM) transceiver (10326) support—Provides support for 4x10GbE links on a QSFP+ port using parallel single mode fiber interface for distances up to 10km. Use 10327 MPO to 4xLC SMF patch cord to break out 4 fiber pair in parallel fiber to separate 4xLC fiber pair.

VRF Capacity Increase Enhancements in 8.41.01.0004

The number of static VRFs supported has increased to 256 static VRFs from 128. The number of dynamic VRFs supported has also increased to 128 dynamic VRFs from 64. This increase is only applicable to modules with 2GB of RAM. (S180/S140/S155, SSA180/SSA150A)

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.

Route Table Capacity Increases Enhancements in 8.41.01.0004

IPv4 Route table Capacity:

The S-Series IPv4 FIB has been increased to 800k from 525k. The IPv4 RIB remains unchanged at 1.6M

IPv6 Route table Capacity:

The S-Series IPv6 RIB has been increased to 128k from 50k and the IPv6 FIB has also been increased to 100k from 50k

SPB CLI Enhancements in 8.41.01.0004

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

```
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

40GbE Problems Corrected in 8.41.01.0004	Introduced in Version:
Packets containing tag-based priority information that are received on an S180 40G front-panel port can subvert the "flex-edge" reprioritization scheme within the MAC chip. As a result, such packets could be given significantly higher internal priority, potentially contending with true "high value" (for example, L2/L3 protocol) packets; or they could be given lower priority, and potentially dropped in an oversubscription scenario.	8.11.01
When a 40G port is configured for 4x10 mode and one of the 10G ports is administratively disabled, the link LED turns off even though the remaining 10G ports are up and passing traffic.	8.22.01
CLI command "show port status" may display improper type for "4x10g-f100" hydra cable.	8.32.01

802.1D Priority Problems Corrected in 8.41.01.0004	Introduced in Version:
An unexpected reset may occur while configuring interfaces.	8.11.01

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 then re-create it.	8.32.01
When ACL logging is enabled on a Policy ACL, the Policy ACL specific field "set-dscp <value>" does not appear in the log message.	8.32.01

Auto Negotiation Problems Corrected in 8.41.01.0004	Introduced in Version:
A 10GBASE-T port sometimes does not establish a link with a Intel Quad i340-T4 system. When connecting/disconnecting the cable repeatedly after a number of interactions (not always the same from 2 to 15), the port no longer links. The link can be recovered by disabling/enabling negotiation on the system. This problem is addressed by implementing: set port low-power-mode <port-string> disable Execute this command to prevent the problem from happening. Using this command disables auto power-down mode on port that is linked to the Intel Quad i340-T4 system.	7.91.03

BGP Problems Corrected in 8.41.01.0004	Introduced in Version:
BGP peer is not immediately deleted causing update-source to fail.	8.11.01
Reset may occur after deleting an active BGP router with message log: "Assert in ntlavll.c line 644".	8.01.01
Graceful restart may fail when switch has a second VRF provisioned, but not configured.	8.31.01
DSI Reset may occur after adding or changing loopback address of active BGP router.	8.01.01
BGP update message containing duplicate MEDs of zero are accepted by the router.	8.21.01

BGP Problems Corrected in 8.41.01.0004	Introduced in Version:
BGP may reject routes if aggregate-address is used.	8.22.01
With BGP graceful-restart configured, MPLS label table may not be synchronized on all modules in a chassis after failover.	8.31.01

Data Center Bridging Problems Corrected in 8.41.01.0004	Introduced in Version:
On a bonded S6 or S8 chassis, congestion on a software bonding port may fail to generate CNMs if the ingress port is on a different module than the bonding port. If the host is on a non-capable slot, the "show dcb cn congestion-point" and "show dcb cn cp-mapping" commands do not display all the CPs for a software bonding port.	8.31.01
The "show config" command stops responding for a very long time at the "dcb cn" application on an S8 with no modules present in slots 1 through 4. Modules must be present in slots 5, 6, 7, and optionally in slot 8.	8.31.01

DHCP Problems Corrected in 8.41.01.0004	Introduced in Version:
Whenever an IPv6 DHCP lease is released the router crashes.	8.32.01
DHCP relay agent does not work over L3VPN.	8.01.01

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 "this server has been invalidated" warning message printed to the console.	1.07.19

EthernetOAM Problems Corrected in 8.41.01.0004	Introduced in Version:
OAM LoopBack sessions are not maintained for longer than three seconds.	8.31.01
OAM enable causes increase in CPU usage, but this usage does not decrease when disabled.	8.31.01

File Management Problems Corrected in 8.41.01.0004	Introduced in Version:
Programmed boot image fails to load with commands "clear config 1" or "reset 1" with a stand-alone chassis.	4.05.08

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:
<p>Messages similar to the following might appear causing packets to be dropped:</p> <ul style="list-style-type: none"> - <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
SSA180 unit may unexpectedly reset upon the occurrence of a "soc_counter_thread: DMA did not finish" error.	8.01.01
HAU Problems Corrected in 8.41.01.0004	Introduced in Version:
The following Syslog message may appear when completing an HAU upgrade: RtrVRRP[13.tVrrpRX]dispatchLib error -25	8.01.01
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
IPv6 Forwarding Problems Corrected in 8.41.01.0004	Introduced in Version:
<p>When using the router as an IPv6 DHCP (DHCP6) relay agent the router may generate a response to a discovery packet even if DHCP6 has not been configured on the router. The response tells the host that the address range is not valid (since it was not configured on the router) and the host never assigns an IP address even if the valid packet from the DHCP server pointed to by the relay agent returns a valid packet.</p> <p>For Example: H1 ----- Rtr1 ----- Rtr2 ---- DHCP6S</p> <p>H1 = Host Number 1</p>	7.62.10

IPv6 Forwarding Problems Corrected in 8.41.01.0004	Introduced in Version:
<p>Rtr1 = Router Number 1 Rtr2 = Router Number 2 DHCP6S = DHCP6 Server</p> <p>In the above example, Rtr1 is a relay agent for the DHCP6 Server.</p> <p>Normally, when H1 send a discover packet to DHCP6 DHCP6 should respond with a valid IP address and the normal operation DHCP operation should complete.</p> <p>Currently, Rtr1 also processes the DHCP Discover packet (even though it is not configured to do so) and responds to the DHCP packet with an error. Because of this the host does not complete the DHCP process despite getting a valid packet from DHCP6S.</p>	

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.	Unknown

Management Problems Corrected in 8.41.01.0004	Introduced in Version:
EDR memory in free list error appears 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:
<p>The following message may appear in S180 class products when multicast traffic is oversubscribed:</p> <pre><163>Jul 8 07:51:05 100.10.10.22 Dune[13.dTcmTask]Petra[0] Received Interrupt PB_EGQ_INT_DIF_AF instance 0, count 6731, value= 0x1</pre>	8.11.01
For single chassis which has both S180 fabric modules and S140/S150 IO modules, enabling IGMP(v4) or MLD (v6) potentially degrades the fabric performance.	8.11.01
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 Problems Corrected in 8.41.01.0004	Introduced in Version:
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

NAT Problems Corrected in 8.41.01.0004	Introduced in Version:
The reply of IPv6 ICMP NATTED packets may be dropped if the NAT outside interface is also configured as a NAT inside interface.	7.91.01
An existing IPv6 NAT binding may continue to be used after the NAT outside interface has been de-configured.	7.91.01

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 are slightly greater than MTU requiring unnecessary IP fragmentation.	7.00.01
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:
Installing a new module might cause network disruption.	7.00.01
Module might reset with messages similar to: "application Chassis Coherency (2) failed to run".	6.11.01
"watchDogTask() application Chassis Coherency failed to run in 300 seconds" message might appear.	Unknown
When tunneled bridge ports are active, infrequently, messages similar to the following may appear: =====	8.20.02
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	

Platform Problems Corrected in 8.41.01.0004	Introduced in Version:
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.	
Module might fail to initialize with message similar to: <163>Mar 13 20:14:57 0.0.0.0 MII[1.tusrApplnit]MdioBcastVerifyOnBus: invalid checksum on phyAddr 89!	8.22.01

Port Status/Control Problems Corrected in 8.41.01.0004	Introduced in Version:
Setting port speed to 1000 on a SOK2209-0204 module may reset the device.	8.32.01
Setting port speed on a SOK2209-0204 module while auto negotiation is enabled may cause the link speed to change.	8.32.01
Transitioning from a port speed of 10G to 1G on a SOK2209-0204 module may cause the following message to appear: "IntrHand[5.tNimIntr]PCIe Core Status Interrupt on NIM 0 inst 0 (sts=0x00300000)"	8.32.01
Traffic indication LEDs on SOK2209-0204 modules may not appear correctly when transitioning from 10G to 1G operation.	8.32.02

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
CoS Flood control is applied to protocol packets.	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 cause 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:
Running Multicast (IGMP Snooping) over SPB-V may lead to mis-programmed hardware leading to packet processing errors and Syslog messages similar to the following: <163>Apr 16 14:14:32 100.10.10.4 System[1]Switch Chip 1 (Slot 1 Mainboard) detected excessive number of reframer interrupts pointing to misprogrammed hardware (code 0x0C008142)	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
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 be seen when issuing SPB CLI 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 that port's spantree adminPathCost to a value greater than 16777213.	8.32.02
After clearing and recreating a static multicast MAC address, traffic destined through a Shortest Path Bridging network are 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:
Module might reset with messages similar to: "serverWatchDog.3 client 18(MstpClnt) not ready in 3365".	6.11.01
When a switch which has backup root activated relinquishes root status, the bridge priority does not automatically revert to the configured bridge priority. This has no effective impact, but it is an unexpected value. Once the switch is no longer root bridge for the given SID, there is no reason to maintain the temporary bridge priority value. It should return to the configured value.	7.00.01
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 do not receive it. 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:
When loading a configuration from a file, the tunnel configuration may fail with the error message: Error: Incompatible ip version with source endpoint or tunnel mode.	7.41.03
Host sourced IPv6 packet entering an L2 tunnel may not forward if the resulting packet is too big. This causes router protocols to fail over tunnels.	8.31.01
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 amount (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-module system.	7.91.01
DHCP relay agent does not work over L3VPN.	8.01.01

VRRP Problems Corrected in 8.41.01.0004	Introduced in Version:
Some routing protocol traffic received through a Shortest Path Bridging network is not correctly reflected in the protocol's operations.	8.31.01
When Host Mobility is enabled, OSPF routes redistributed within an NSSA area are not translated into the backbone as the propagate bit is not set.	8.31.01
The router may stop responding 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
When a host route created by host mobility is removed, the corresponding ARP is not removed.	8.32.01

VSF Problems Corrected in 8.41.01.0004	Introduced in Version:
A module added to a bonded system infrequently gets stuck in a reboot loop enabling, and then disabling, bonding.	8.11.01
VSF can cause watchdog timeout when slots are added or removed from extremely busy systems.	7.62.00
In a bonded chassis, the reframer DB may remove entries used by a connection.	7.11.01

VSB Problems Corrected in 8.41.01.0004	Introduced in Version:
<p>In software VSB chassis that have Shortest Path Bridging or Tunnels active, very infrequently a module may reset with a message similar to:</p> <pre><0>Fuji[12.tNimIntr]Switch Chip 3 (Slot 12 Mainboard) detected fatal condition. (0x00ebfdb0 0x01827aa8 0x01821394 0x00443130 0x00470ea4 0x00472498 0x01b2de24 0xeeeeeeee),</pre> <p>or</p> <pre><1>MCNXMGR[7.tFujiAge]ASIC failed to write callback, transaction = 0, chip = 1 (0x00ec0078 0x019e8acc 0x0184b054 0x0184b3f0 0x01b2f3e4 0xeeeeeeee),</pre> <p>or</p> <pre><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 bmgrIntf:0x2136b9 getBuf:0x84b4a9ca link0Buf:0xa9c6a9c6 link1Buf:0x84b084b0 walk0Buf:0 walk1Buf:0 freeBuf:0 freeChainPkt 0:0xd743d743 1:0xdee1dee1 2:0x6e406e40 (0x00ec0078 0x01858574 0x01856bbc 0x00470058 0x000000518 0x00000004 0x01b1a9f0 0x007116d0 0x0070096c 0x007009cc 0x03369dd4 0x006f1bb8 0x0070c94c 0x0070c8c4 0x0070d018 0x018611dc 0x018557d4 0x0185eb24 0x01b2f424 0xeeeeeeee)</pre> <p>The following non-reset level messages may also be logged:</p> <pre><3>mazama[8.tDispatch]Host buffer already free. freeBuffer:10897</pre> <p>If that message is logged, there is vulnerability to the resets mentioned above."</p>	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

Host Problems Corrected in 8.32.02.0008	Introduced in Version:
When a chassis is booting and one or more blades in the chassis continually reset or halt at initialization (usually, but not limited to, hardware failure reasons), this may cause other blades to reset before fully booting. A message similar to: <0>Rdy2swch[5.tRdy2Swch]Initialization failed to complete(388 seconds allotted, 413 seconds elapsed). DistServ:[13], incomplete. (0x00e8a334 0x00990d8c 0x01af23e4 0xeeeeee) Dwill be logged. It is possible that the chassis will never fully boot until the originally resetting blade(s) are removed.	7.00.01
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
MPLS Problems Corrected in 8.32.02.0008	Introduced in Version:
When running BGP in multiple VR's on a multiblade system, module reset may occur with the following message logged: "SMS assert in asejoifm.c at line 2370 : (null) INVALID BRANCH 0 (null) 0"	8.31.01
On failover, MPLS labels remain in forwarding plane leading to connectivity issues.	8.31.01
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
Port Status/Control Problems Corrected in 8.32.02.0008	Introduced in Version:
Setting port speed to 1000 on a SOK2209-0204 module may reset the device.	8.32.01
Setting port speed on a SOK2209-0204 module while auto negotiation is enabled may cause the link speed to change.	8.32.01
Transitioning from a port speed of 10G to 1G on a SOK2209-0204 module may display the following: "IntrHand[5.tNimIntr]PCle Core Status Interrupt on NIM 0 inst 0 (sts=0x00300000)"	8.32.01
Setting port speed on 40G ports to 40000 and performing module reset might cause the board to enter into reset loop.	8.32.01
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
VRF Tree Problems Corrected in 8.32.02.0008	Introduced in Version:
DHCP relay agent does not work over L3VPN.	8.01.01

Features Enhancements 8.32.01.0021

Additional Modules Supported in 8.32.01.0021

SKL8008-0810-F8

S-Series S180 Class I/O-Fabric Module, Load Sharing - 8 Ports 10GBASE-X via SFP+ and 2 ports 40GBASE-X Ethernet via QSFP+

STL2006-0850

S-Series S140 I/O Module - 48 Ports 10/100/1000BASE-TX via RJ45 with PoE (802.3at) and 2 ports 40GBASE-X Ethernet via QSFP+, EEE

SGL2001-0850

S-Series S140 I/O Module - 48 Ports 1000BASE-X ports via SFP and 2 ports 40GBASE-X Ethernet via QSFP+

SOK2209-0204

S-Series Option Module (Type2) - 4 Ports 10GBASE-T with PoE (802.3at)

Additional Transceivers Supported in 8.32.01.0021

Part Number: 10318 - 40Gb, Active Optical DAC with integrated QSFP+ transceivers, 100 m

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.

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.

Routing Capacity Change Support in 8.32.01.0021

FIB Capacity Increase: Capacity has been increased from 525k to 600k

Problems Corrected in 8.32.01.0021

BGP Problems Corrected in 8.32.01.0021	Introduced in Version:
BGP may reject routes if aggregate-address is used.	8.22.01
When entire internet route table is loaded the following message might be seen: "FwdMgr[1.tRtrPtcls]FIB entry pool exhausted". The IPv4 FIB can now support approximately 600,000 routes.	8.32.01
Graceful restart may fail when DUT has a second VRF provisioned, but not configured.	8.31.01
With BGP graceful-restart configured, MPLS label table may not be synchronized on all blades in a chassis after failover.	8.31.01

Data Center Bridging Corrected in 8.32.01.0021	Introduced in Version:
The "show dcb cn cp-mapping" command may display an invalid CP index of 33 following a system reset with multiple CNPVs configured.	8.31.01
A message similar to "DCB[2] setCnCpEnableDD:Unable to convert port 412 portCpHwIdx 7 to a txQueue" or "DCB[2] setCnCpEntryDD:Unable to convert port 412 portCpHwIdx 7 to a txQueue" may be seen at bootup with multiple CNPVs configured. As a result, the CP configuration may be lost.	8.31.01
The user may see a message similar to "DCB[10]HW: Failed to enable CP; sport=1336, hwPort=12259, pcpHwIdx=3, qpldx=0, dunePri=4" with CN configured on a hardware bonded S-Series system.	8.31.01

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
Installing a new blade might cause network disruption.	7.00.01
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>may be displayed. Whenever one of these messages is displayed, a packet that should have been forwarded will be dropped.</p>	8.11.01
<p>Under unknown circumstances, an S-Series module may begin to transmit malformed packets associated with certain traffic flows, resulting in degradation of these flows. Workaround is to reboot the module. The susceptible modules include:</p> <p>SK2008-0832 SK2009-0824 SL8013-1206-F8 SL8013-1206-F8A SK8008-1224-F8 SK8009-1224-F8 SL8013-1206 SL8013-1206A SK8008-1224 SK8009-1224</p>	8.01.01
<p>It is possible that during initialization a module might halt with the following message: "This device has encountered a hardware failure. Fabric chips SPI4.2 initialization failed. Please contact Support for a possible repair/replacement. Press <r> to reset board."</p> <p>If you do not see this message on subsequent board initializations you have hit this case and the error can be ignored.</p>	8.01.01
Host Mobility Problems Corrected in 8.32.01.0021	Introduced in Version:
<p>The host-mobility aging feature has been removed in order to support host-mobility in a segmented network design whereby the routers' VRRP interfaces are not connected.</p>	8.11.01
<p>Host-mobility is now supported in a segmented VRRP network design where the router's VRRP interfaces are not connected.</p>	8.31.01

Host Services Problems Corrected in 8.32.01.0021	Introduced in Version:
<p>When tunneled bridge ports are active, infrequently messages similar to:</p> <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, you could see an exhaustion of resources that prevent IP host communication and loss of device management.	7.91.01
IPv4 Forwarding Problems Corrected in 8.32.01.0021	Introduced in Version:
After router failover, any nexthops no longer reachable remain in the route table.	8.31.01
IPv6 Forwarding Problems Corrected in 8.32.01.0021	Introduced in Version:
Traffic from server to client might not be forwarded by the hardware over SPB due to incorrect VID being used.	8.31.01
LSNAT Problems Corrected in 8.32.01.0021	Introduced in Version:
Using LSNAT or NAT an FTP binding will not be deleted immediately when the FTP connection is removed with a "FIN".	7.31.02
Management Problems Corrected in 8.32.01.0021	Introduced in Version:
SFP entity MIB sensor does not support VSB ports. "Show port transceiver" CLI output indicates that sensor data is not available for these ports.	8.31.01

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

MPLS Problems Corrected in 8.32.01.0021	Introduced in Version:
Cutting/pasting a large configuration from a text file into the router under test may result in message similar to: "<0>sms[2.tRtrPtcls]SMS assert in ahlij.c at line 737 : == mj_cb->desired_state 1 AHL_MJI_DEL_JOIN 2".	8.31.01
The output of 'show mpls forwarding-table' cannot be stopped with ctrl-C or by entering 'q' during pagination.	8.31.01
The results from the 'show ip-binding' command may appear in random order when presented to the user.	8.31.01
When setting the ldp-label-retention mode this message will be displayed: "Err setting LDP entity admin up."	8.31.02

MSDP Problems Corrected in 8.32.01.0021	Introduced in Version:
MSDP can't be removed when only msdp originator-id exists.	8.20.02
Exceeding MSDP multicast flow limit may result in a SMS assert similar to "sms[3.tRtrPtcls]SMS assert in nbbpd.c at line 691 : != proc_type 0x0x 0 NULL 0x0x0 (0x00e86580 0x01e9d714 0x004c54ac 0x004c556c 0x004ccfec 0x02629250 0x02629938 0x0263e9fc 0x0263f380 0x0260dfb0 0x004b4cc8 0x004b5184 0x004b54b8 0x004bad08 0x004e64b4 0x01ea7674 0x01ea5ce4 0x01ad56e4 0xecececece)"	8.11.01

Multicast Problems Corrected in 8.32.01.0021	Introduced in Version:
Static layer 2 multicast traffic is not forwarded through a hardware VSB device that is using Shortest Path Bridging.	8.31.01
Multicast cache entries show up in the router even without a multicast routing protocol enabled on an interface.	8.31.01

Multicast Problems Corrected in 8.32.01.0021	Introduced in Version:
<p> Toggling LACP interfaces may result in module reset with message similar to: "<0>sms[1.tRtrPtcls]SMS assert in msnmcfsm.c at line 1023 : (null) INVALID BRANCH 0 (null) 0 (0x00ea7d44 0x01ed54a0 0x004c8034 0x004c835c 0x027a9c80 0x0252e2b4 0x02793444 0x0278ae04 0x027889a8 0x02788dd8 0x004b6648 0x004b6c0c 0x004b6f40 0x004bd890 0x004e903c 0x01edf31c 0x01edd974 0x01b0c064 0xeeeeeeeee)". </p>	8.11.01
Multicast frames that are buffered and forwarded do not have TTL decremented.	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

MVRP Problems Corrected in 8.32.01.0021	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
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

NAT Problems Corrected in 8.32.01.0021	Introduced in Version:
An existing IPv6 NAT binding may continue to be used after the NAT outside interface has been de-configured.	7.91.01
The reply of an IPv6 ICMP NATTED packet may be dropped if the NAT outside interface is also configured as a NAT inside interface.	7.91.01
An FTP transfer of data may fail while on an inside interface and using NAT.	8.31.01
If Tunneled Bridge Ports or Network Address Translation are active, traffic may not be forwarded correctly.	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:
<p> Negating PIM DM incoming interface may result in an assert similar to "SMS assert in qptukrcv.c at line 819 : == group_mode 5 QPTM_RPM_PIM_MODE_UNROUTABLE 1 (0x00e84f50 0x01e9c158 0x004c54ac 0x004c57d4 0x0216f6c8 0x02170958 0x02168a04 0x0211c948 0x004b4cc8 0x004b5184 0x004b54b8 0x004bad08 0x004e64b4 0x01ea60b8 0x01ea4728 0x01ad4124 0xeeeeeeeee)" </p>	8.21.01
"show ip mcache" shows a corrupted/incorrect Source/Destination IP in the display output.	8.31.01

PIM-SM Problems Corrected in 8.32.01.0021	Introduced in Version:
Non-DR may forward first multicast packets.	8.31.02
PIM-SM non-Designated Router (DR) may forward Register packets to the Rendezvous Point (RP).	8.31.01

Platform Problems Corrected in 8.32.01.0021	Introduced in Version:
Very rarely, an SSA180 unit may unexpectedly reset upon the occurrence of a "soc_counter_thread: DMA did not finish" error.	8.01.01
System logs the message "bcmStrat[1.tNimIntr]MEM_FAIL_INT_STAT=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000, IP3_INTR_STATUS_1=0x00000000, IP3_INTR_STATUS_2=0x00000000, EP_INTR0_STATUS=0x20000000, EP_INTR1_STATUS=0x00000000, GPORT_INTR_STATUS0=0x00000000, GPORT_INTR_STATUS1=0x00000000, XPORT_INTR_STATUS0=0x00000000, XPORT_INTR_STATUS1=0x00000000, XPORT_INTR_STATUS2=0x00000000, XPORT_INTR_STATUS3=0x00000000, XQPORT_INTR_STATUS0=0x00000000, XQPORT_INTR_STATUS1=0x00000000, XQPORT_INTR_STATUS2=0x00000000, XQPORT_INTR_STATUS3=0x00000000, XQPORT_INTR_STATUS4=0x00000000, XQPORT_INTR_STATUS5=0x00000000, SPORT_INTR_STATUS=0x00000000" and resets.	8.31.02
A S180 fabric blade, when inserted into a chassis where multicast traffic (which is other than VLAN flooded traffic) is present, may get stuck on constant reboot with error message: <3>NonVol[13.tNvFinIn]nonvol_copyStream:failed to read (75),(S_rpcLib_RPC_TIMEDOUT)	8.12.04
When uploading new image to chassis, errors noticed when distributing image to compatible slots in the chassis.	4.21.19
"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:D<163>Feb 12 10:42:00 100.10.10.22 dot3Mgt[4.tEmanate10]dot3MgtDist::ifJackEntryGet():sendMessage(ackReq)!=kDs_good;sendMask=0x10000	7.70.00
System logs the message "Default[1.tNimIntr]Assertion failed: SOC_REG_IS_VALID(unit, reg), file /firmware/common/bcmStrataDrv/04_12_50/sdk/src/soc/common/.reg.c, line 2897 (0x00dcd2dc 0x004108f8 0x010a9d80 0x010ad244 0x00ea94b4 0x00e9e82c 0x00444bc4 0x0046e4ec 0x0046ee40 0x019eec64 0xeeeeeeee)" and resets.	8.11.01
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

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:
Changes to the Shortest Path topology may take longer than expected time to converge on a new topology.	8.31.01
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
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 0xeeeeeeee)".	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
Routed traffic will fail on a helper router when Routing-As-A-Service is configured and the SPB domain fails over to Cist.	8.31.01

Shortest Path Bridging Problems Corrected in 8.32.01.0021	Introduced in Version:
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 BPDUs 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
In a shortest-path-bridging network, hardware connections for routed traffic may not be removed when the routes affecting the traffic change.	8.31.01
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 BPDUs with an agreement BPDUs. 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 Ip (Loop Protect), it will detect a loop protect event and remain in the listening state.	7.60.01
The Multisource function detects multiple BPDUs 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 BPDUs 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

Spanning Tree Problems Corrected in 8.32.01.0021	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

Transceiver Problems Corrected in 8.32.01.0021	Introduced in Version:
Some QSFP transceivers may not achieve link on SL8013-1206A or SL8013-1206-F8A module.	8.31.01
Some SFP transceivers may cause port link flap on SL8013-1206A or SL8013-1206-F8A module.	8.31.01
Some QSFP+ transceivers will come up as "Unauthenticated".	8.31.02
QSFP+ 40GB to 4x10GB cables show up as unauthenticated with no link.	8.31.02

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 requires only a source tunnel endpoint to be configured on a tunnel, causes the tunnel to be deleted.	7.41.02
The switch may crash when changing the VLAN membership of a tunnel bridge port.	8.21.01

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

Hardware VSB Problems Corrected in 8.32.01.0021	Introduced in Version:
Very infrequently, after a module resets in a Hardware Bonded Chassis, other modules may reset on a system watchdog timeout.	8.11.01
On a hardware or software bonded system with LFR ports, the following assertions (or similar) are seen during startup or when an individual slot leaves or joins the stack: <3>Dune[5.tDSrecv3]assert: ../fapPort_dist.hxx:1704 32768 !< 6144 <0>Default[5.tDSrecv3]Assertion failed: 0, file ../dune_utils.cxx, line 4754 (0x00ed9ae0 0x004108f8 0x013901d0 0x035484a0 0x0135ccac 0x013545e4 0x006c74c4 0x006aa7c0 0x006aafc0 0x006ace3c 0x006b5c14 0x00727dc8 0x00738668 0x007399c4 0x01b5bfe4 0xe0000000) The devices subsequently start rebooting, which may never stop. The underlying issue involves a mishandling of remote ports. Abandoning the use of LFR ports avoids the issue.	8.31.01

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

Feature Enhancements 8.31.02.0014

Enhancements in 8.31.02.0014
FW support for the Purview Application Sensor model: PV-FC-180.

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

802.1q Relay Problems Corrected in 8.31.02.0014	Introduced in Version:
The switch may crash when changing the VLAN membership of a tunnel bridge port.	8.21.01

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

BGP Problems Corrected in 8.31.02.0014	Introduced in Version:
BGP update message containing duplicate MED's of zero are accepted by the router.	8.21.01

Chassis Problems Corrected in 8.31.02.0014	Introduced in Version:
Module might reset with messages similar to: "application Chassis Coherency (2) failed to run".	6.11.01

Chassis Problems Corrected in 8.31.02.0014	Introduced in Version:
Configuration associated with the module will be lost when: - replacing ST2206-0848 with ST2206-0848A - replacing ST8206-0848-F8 with ST8206-0848-F8A - replacing SL8013-1206 with SL8013-1206A - replacing SL8013-1206-F8 with SL8013-1206-F8A	8.22.01
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
COS TX queue rate shaping values are not being divided among the 40 Gig MUX ports, instead the whole value is applied to each 40 Gig MUX port.	8.11.01
Data Center Bridging Problems Corrected in 8.31.02.0014	Introduced in Version:
DCB application priority will restore only one port's configuration.	7.91.01
The user may see a message similar to: "DCB[10]HW: Failed to enable CP; sport=1336, hwPort=12259, pcpHwIdx=3, qpIdx=0, dunePri=4" with CN configured on a hardware bonded S-Series system.	8.31.01
Distributed Services Problems Corrected in 8.31.02.0014	Introduced in Version:
Bonded Chassis module may infrequently reset after a message similar to the following is logged: "DistServ[4.tDsBrdOk]serverWatchDog.6(PortInfo) client 106(Bonding) not ready in 18399".	8.12.01
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:
USB drives formatted as FAT may not show up as "inserted" or show up when issuing a "dir" CLI command.	8.21.01
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	8.21.01

Hardware Problems Corrected in 8.31.02.0014	Introduced in Version:
These resets will occur more frequently if the switch has Tunnels or Shortest Path Bridging enabled.	
It is possible that during initialization module might halt with the following message: "This device has encountered a hardware failure. Fabric chips SPI4.2 initialization failed. Please contact Support for a possible repair/replacement. Press <r> to reset board." If you do not see this message on subsequent board initializations you have hit this case and the error can be ignored.	8.01.01
HAU Problems Corrected in 8.31.02.0014	Introduced in Version:
High-availability upgrade groups may not be displayed correctly in the show config output.	7.60.01
Host Services Problems Corrected in 8.31.02.0014	Introduced in Version:
"show neighbor" CLI command may generate the syslog message "DistServ[4.tCLI1]sendMsg.1(Config) msg (2147483660) too long(1405) from client 37(NDS)" and some neighbor info may be missing.	8.21.02
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
IPv6 Forwarding Problems Corrected in 8.31.02.0014	Introduced in Version:
IPv6 traceroute ignores the -s source IPv6 address option.	7.00.01
The SPB Vid value was incorrect due to retrieving the Vid from the packet transformation instead of from the Forwarding Attributes for the flow.	8.31.01
LSNAT Problems Corrected in 8.31.02.0014	Introduced in Version:
While running SLB traffic and creating bindings the module might reset with messages similar to: Message 5/302 Exception PPC750 Info 08.31.02.0003 10/02/2014 16:49:03	8.31.01

LSNAT Problems Corrected in 8.31.02.0014	Introduced in Version:
Exc Vector: DSI exception (0x00000300) Thread Name: tRtrASvcMain Message 20/302 Exception PPC750 Info 08.31.02.0003 10/02/2014 15:56:57 Exc Vector: DSI exception (0x00000300) Thread Name: tDSrecv5	
Management Problems Corrected in 8.31.02.0014	Introduced in Version:
SFP entity MIB sensor does not support VSB ports. "Show port transceiver" CLI output indicates that sensor data is not available for these ports.	8.31.01
MPLS Problems Corrected in 8.31.02.0014	Introduced in Version:
The CLI command 'show mpls propagate-ttl' does not display the current system settings.	8.31.01
Toggleing MPLS admin state causes 'bgp no enable' command to be removed from config file.	8.31.01
Multicast Problems Corrected in 8.31.02.0014	Introduced in Version:
mgmdStdMib InverseRouterCacheTable may not SNP walk properly.	7.60.01
IGMP/MLD snooping crashes with message similar to: "Recv base index out of range baseidx:3679 flowIdx:3679".	8.31.01
Static layer 2 multicast traffic is not forwarded through a hardware VSB device that is using Shortest Path Bridging.	8.31.01
When running S130/S150/S155/SSA130/SSA150 only software it is possible to see the following message at the CLI when exceeding the limits of the number of hardware resources for reframing when using IGMP: "Invalid MCI -1, for ASIC ""x"", where ""x"" can be any number"	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
NAT Problems Corrected in 8.31.02.0014	Introduced in Version:
'overloaded' option of a NAT list rule may not be allowed if "inside_vrf" is configured.	8.31.01
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
<p> Toggling LACP interfaces may result in module reset with message similar to: "<0>sms[1.tRtrPtcls]SMS assert in msnmcfsm.c at line 1023 : (null) INVALID BRANCH 0 (null) 0 (0x00ea7d44 0x01ed54a0 0x004c8034 0x004c835c 0x027a9c80 0x0252e2b4 0x02793444 0x0278ae04 0x027889a8 0x02788dd8 0x004b6648 0x004b6c0c 0x004b6f40 0x004bd890 0x004e903c 0x01edf31c 0x01edd974 0x01b0c064 0xeeeeeeee)". </p>	8.11.01
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
<p> The following message may be shown in S180 class products when multicast traffic oversubscribed fabric chips: <163>Jul 8 07:51:05 100.10.10.22 Dune[13.dTcmTask]Petra[0] Received Interrupt PB_EGQ_INT_DIF_AF instance 0, count 6731, value= 0x1 </p>	8.11.01
<p> 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" </p> <p> 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" </p>	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
Inserting a 40Gb 0.5m copper QSFP cable into a 40Gb port will result in the board resetting.	8.31.01
SFP sensors information may not display 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
<p> System logs the message: "Default[1.tNimIntr]Assertion failed: SOC_REG_IS_VALID(unit, reg), file /firmware/common/bcmStrataDrv/04_12_50/sdk/src/soc/common/.reg.c, line 2897 (</p>	8.11.01

Platform Problems Corrected in 8.31.02.0014	Introduced in Version:
0x00dcd2dc 0x004108f8 0x010a9d80 0x010ad244 0x00ea94b4 0x00e9e82c 0x00444bc4 0x0046e4ec 0x0046ee40 0x019eec64 0xeeeeeeee)" and resets.	
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
RaaS Problems Corrected in 8.31.02.0014	Introduced in Version:
The command 'show raas' on a RaaS helper router may display more than 8 main routers even though only 8 are used in forwarding packets.	8.31.01
On a RaaS help router, when a main router leaves the forwarding list, connections for flows forwarded to that main router may not be deleted.	8.31.01
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
Routing Problems Corrected in 8.31.02.0014	Introduced in Version:
TLV with no in-service bit set is generated after router failover.	8.31.01
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

Shortest Path Bridging Problems Corrected in 8.31.02.0014	Introduced in Version:
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, 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.	8.31.01
Traffic may not recover after disable/re-enable SPB.	8.31.01
SPB-ISIS packets are flooded when SPB is globally disabled.	8.31.01
IP directed broadcasts transmitted to their destination subnet are not broadcast across SPB domain.	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:
Module might reset with messages similar to: "serverWatchDog.3 client 18(MstpClnt) not ready in 3365".	6.11.01
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 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

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
In a Software VSB chassis, if a blade with active VSB ports resets, there may be a failure to forward packets received on the chassis the blade reset in, to ports on the remote chassis.	7.40.01
After a module resets in a Hardware Bonded Chassis, other modules may reset on a system watchdog timeout.	8.11.01

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 S-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.
Congestion Notification (CN) - IEEE 802.1Qau Congestion Notification (CN) enhances existing Data Center Bridging (DCB) functionality and provides End-End congestion management for protocols that are capable of

Enhancements in 8.31.01.0006
transmission rate limiting to avoid frame loss. Congestion Notification (CN) is hardware dependent and only supported S140 and S180 hardware classes.
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.
IP Tunneling Enhancements: <i>NAT Services Support</i> - NAT /LSNAT/TWCB functionality is now supported for IP tunnel interfaces. <i>Layer Tunnel Switching</i> - Support has been added to allow traffic to be switched directly from a tunnel interface to another tunnel interface. Previous releases supported switching to/from non-tunnel to tunnel interfaces.
L3VPN Enhancements: <i>L3VPN over SPB Core</i> - Support has been added to support L3VPN over BGP/Shortest Path Bridging Core network. This feature allows the use of a SPB transport for L3VPN functionality. <i>L3VPN over an IPv6 BGP/MPLS Core</i> - Support has been added to support L3VPN over IPv6 BGP/MPLS Core network. Previous releases supported L3VPN over IPv4 BGP/MPLS Cores only. <i>L3VPN Graceful Re-start</i> - Graceful re-start is supported with GRE / L3VPN implementations allowing forwarding of existing L3VPN traffic during protocol restart events.
NAT Services Enhancement - Basic Stateful Firewall like functionality utilizing the embedded NAT services engine to provide connection initiation only from designated “inside” to designated “outside” interfaces. This feature can help secure internal resources from being directly compromised from devices entering the network from “outside” interfaces. This feature uses the NAT services engine but does not require address translation.
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. <i>UDP Timing Probe</i> - 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. <i>DNS Timing Probe</i> - Uses the DNS protocol to transmit a DNS query a destination device to determine the round-trip-delay of the DNS answer. <i>UDP/DNS State Probe</i> - Provides the ability to verify data in a DNS resource record in the answer section of the DNS response packet.
VRF Capacity Increase - The number of dynamic VRF’s supported has increase to 64 dynamic VRF’s from 16. This increase is only applicable to modules with 2GB of RAM. (S180/S140/S155, SSA180/SSA150A)
New Feature Licenses: <i>S1/S1A Policy User Capacity License, S-EOS-USER</i> - Policy User Capacity license to allow support up to 8k users. Used for single fabric systems installed in S1/S1A chassis.

Enhancements in 8.31.01.0006

Flow Capacity License, S-EOS-Flow - Flow capacity license for SSA /Purview appliance 10G ports to allow 1M flows per CoreFlow2 ASIC. Applicable to 10G ports only, see port to ASIC mapping.

Problems Corrected in 8.31.01.0006

802.1d Filter Database Problems Corrected in 8.31.01.0006	Introduced in Version:
In customer bridge mode, packets destined for the provider bridge network group address are forwarded.	8.21.01
MAC addresses that should age out form filter database will fail to do so. The frequency of this will increase with lower mac age times.	1.07.19
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:
Packets sent from a VRF host may be delayed while the ARP is resolved. Additionally the destination host may see 10 or more ARP requests to the destination IP address.	7.60.01
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
BGP Problems Corrected in 8.31.01.0006	Introduced in Version:
The 'show ipv6 bgp dampened-routes' command does not support a 'wide' option in which all output appears on one line.	8.01.01

BGP Problems Corrected in 8.31.01.0006	Introduced in Version:
There is no CLI command to display all of the parameters associated with a BGP route-flap dampening configuration table.	8.20.02
The BGP peering session between two BGP peers may bounce if BGP is running over a tunnel interface, BGP is advertising a large number of routes, and jumbo packets are disabled on ports in which the BGP packets are transmitted.	8.21.01
A system reset may occur if BGP peering sessions bounce while loading the full Internet routing table. The following error message will appear: 'SMS assert in qbnmutil.c at line 478'	7.20.01
BGP does not provide a CLI command to allow the user to specify a per peer local AS number.	7.20.01
BGP prefix-lists referenced from a BGP route-map may not function in cases where the prefix-list has been configured, removed, and reconfigured in a short time period.	7.20.01
If a BGP Update message is received with no NLRI path attribute the peering session is torn down.	7.20.01
The "orf-association" setting for BGP route-maps matching the mpls-bgp-vpn sub address-family will not be internally programmed after the vpnv4 or vpnv6 address-family is disabled and then reenabled.	8.01.01
The BGP next-hop for ipv6 L3VPN routes will not change to the local ipv6 address after negating the vpnv4 address-family. The mapped ipv6 address based on the local ipv4 address is still used in this case.	8.01.01
The 'show ipv6 bgp' command does not support a 'wide' option in which all output appears on one line.	7.30.01
The 'show running-config' and 'show config' outputs may display extra white space between options for the BGP 'aggregate address' command.	8.22.01
Configuring OSPFv2/v3 administrative distance can cause a DSI in tRtrPtcls if BGP is also configured.	8.22.01

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

Data Center Bridging Problems Corrected in 8.31.01.0006	Introduced in Version:
CN does not properly update the automatic alternate priority when a new CNPV is created with a value one less than an existing CNPV. The existing CNPV will continue to remap priorities to the new CNPV on ingress.	7.91.01

DHCP Problems Corrected in 8.31.01.0006	Introduced in Version:
When using an ip-helper and the option 82 dhcp relay information option for VPNs, DHCP packets may not be properly relayed.	7.40.01
DNS Problems Corrected in 8.31.01.0006	Introduced in Version:
When a VLAN interface is removed from the configuration the associated DNS-Names are also removed but the DNS-Names count is not decremented. This makes it appear that more DNS-Names are configured than actually are and may cause an error when new DNS-Names are configured. The count will be correct if the router is rebooted.	7.91.01
Graceful Restart Problems Corrected in 8.31.01.0006	Introduced in Version:
Graceful restart for BGP/OSPF is not working with VPN routes.	8.01.01
Graceful restart in an L3VPN environment may fail if the OSPF pe-ce option is enabled.	8.22.01
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

IPSLA Problems Corrected in 8.31.01.0006	Introduced in Version:
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 Forwarding Problems Corrected in 8.31.01.0006	Introduced in Version:
Given overlapping IPv6 routes with different prefix lengths, the 'show ipv6 route <prefix/prefix length>' CLI command will only successfully display the route with the shortest prefix length.	7.00.01
When overlapping prefixes exist in the IPv6 route table, the CLI command 'show ipv6 route <ipv6-address>' will return the route with the shortest prefix length instead of the longest prefix length.	8.01.01
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
The router will ignore IPv6 Neighbor Advertisements if the Neighbor Advertisement does not contain a Target Link-layer address. This is true even if the neighbor solicitation was a unicast request.	7.41.02
If a router receives a Neighbor Solicitation without a Source Link-Layer address then the response for the solicitation may contain a destination MAC address of 00:00:00:00:00:00.	7.01.02
ISIS Problems Corrected in 8.31.01.0006	Introduced in Version:
CLI commands 'ip router isis' and 'ipv6 router isis' are not provided under loopback interfaces.	7.30.01
ISIS Hostnames exceeding 10 characters cannot be displayed with the 'isis database' command.	8.21.01
Processing very high rates of unknown traffic may cause ISIS routes to be lost.	7.91.01

ISIS Problems Corrected in 8.31.01.0006	Introduced in Version:
When displaying complete ISIS database information for a router that has a hostname defined, the system id may be displayed instead.	8.12.01
Jumbo Problems Corrected in 8.31.01.0006	Introduced in Version:
Port Jumbo MTU settings allowed for values below 1519.	8.01.01
LSNAT Problems Corrected in 8.31.01.0006	Introduced in Version:
Fragmented packets are not allowed to traverse across an LSNAT6/4 or LSNAT46 vserver, the packets will be dropped.	7.91.01
No translation of ICMPv6 Packet Too Big to ICMPv4 across an LSNAT64 router.	7.91.01
It is possible for the command "show ip slb statistics" to display '0' deleted Sticky Entries, regardless of how many sticky entries have been deleted.	7.11.01
The inner UDP/TCP hdr checksum on an "ICMP Error" may be incorrect.	6.00.02
The serverfarm must be removed before modifying the vserver serverfarm configuration.	7.20.01
When issuing an ICMP request to a VIP that is accessed using MPLS the ICMP Request will not reply.	7.99.00
LSNAT66/MPLS reverse translation has malformed IPv6 header (and bad TCP chksum).	8.20.02
The CLI command 'show ipv6 route <address>' does not work for LSNAT virtual IP addresses.	7.00.01
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
Entity MIB is missing entries for fan tray modules present in the SSA180 class.	8.01.01
Entity MIB is missing Physical Description (1.3.6.1.2.1.47.1.1.1.2) for fan tray slots present in the SSA180 class.	8.01.01
WebView Chassis display is limited to only 13 slots.	1.07.19
MLD Problems Corrected in 8.31.01.0006	Introduced in Version:
When MLD includes sources different from the real multicast source, the following error message will appear: 'qptuaapi.c at line 1324'.	8.21.01
MPLS Problems Corrected in 8.31.01.0006	Introduced in Version:
Ingress MPLS Labeled packets on a Service Provider port contains unexpected data after the MPLS header.	8.22.02

MPLS Problems Corrected in 8.31.01.0006	Introduced in Version:
After router failover, some filter connections may not be removed in MPLS networks.	8.11.01
If IPv6 MPLS is deleted from the configuration and then restored, MPLS labels are not assigned to FECs for routes learned in the core.	8.21.01
The CLI command 'show mpls forwarding-table' consumes a large amount of available CPU to generate its output with large route tables in place.	8.20.02
The CLI command 'show mpls forwarding-table' operates very slowly with large routes tables in place.	8.20.02
The L2 MAC address were not being updated in the Soft Forwarding path for MPLS flow.	8.20.02
The 'prefix' suboption to the 'show mpls forwarding-table' command is not functional.	8.20.02
Transformation was using MPLS label TTL instead of Inner IP' TTL.	8.20.02
When MPLS is enabled after system boot, MPLS FECs are unnecessarily created for each connected subnet.	8.20.02
When there is more than one nexthop for a given route via ECMP, only one nexthop appears in output of 'show mpls forwarding-table'.	8.20.02
Wrong MTU value set in ICMP Fragmentation Needed/ICMPv6 Packet Too Big packet sent when MPLS/L3VPN labels added to packet making it too big for egress.	8.20.02

Multicast Problems Corrected in 8.31.01.0006	Introduced in Version:
If a VLAN has an egressing tunnel port, and IGMP/MLD is enabled on that VLAN, multicast may not follow proper forwarding rules.	8.20.02
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:
A dynamic MVRP VLAN will not be deregistered on a port when the "set mvrp vlan restricted enable" command is entered.	7.91.01
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

NAT Problems Corrected in 8.31.01.0006	Introduced in Version:
Configuring a NAT static rule with the "allow_frag" option would not result in the "allow_frag" set on the configuration line.	8.11.01
For MPLS Soft Forwarding Path flows using NAT the ether offset was incorrect so the TCP/UDP checksum was not being updated.	8.20.02
If ACLs configured in multiple list rules configured on different VRF's yield the same permissions, a NAT binding may be created on the wrong VRF's.	7.20.01

NAT Problems Corrected in 8.31.01.0006	Introduced in Version:
It is possible to create a hairpin connection when an inside address is destined to a NAT global address used in another NAT binding.	7.91.01
MTU discovery packets may not be transmitted if the ingress is a jumbo and the packet is NAT'd.	8.21.01
Router does not respond to neighbor solicitations for NAT addresses.	8.11.01
When a "nat inside" and a "nat outside" is configured on the same interface it is possible that the return NAT'd packet will be dropped and the nat translation will not work.	7.91.01
When accessing NAT over MPLS the packet will not be NAT'd.	7.99.00
When using NAT route leak with an inside_vrf over mpls the NAT translation will be dropped.	7.99.00
While Processing many LSNAT TFTP packets it is possible that some TFTP data packets may be dropped.	7.91.01
Ingress MPLS labeled packets Soft path Forwarding offset was incorrect by 4 bytes.	8.20.02
When running NAT over MPLS the NAT packets will not be processed.	7.99.00

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
If a saved config file contains an invalid OSPF area range command, it will log an error on configure. The invalid command had no effect.	8.22.02
The "debug ip ospf packet" display for virtual interfaces reads "Interface not found for ifIndex 0".	8.11.01

OSPF Problems Corrected in 8.31.01.0006	Introduced in Version:
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
OSPFv2 default-information originate metric command incorrectly shows a range of 0-65535 instead of 1.	8.22.01
An assert in thread tRtrPtcls may occur with the following message "SMS assert in qodmbld3.c at line 377 : (ls_id == NULL) 0 (!NBB_MEMCMP(&return_cb->lsa_header.ls_id, ls_id, QOPM_ID_LEN" if OSPF AS external lsa's exist that fall under rfc2328 appendix E.	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
When running OSPFv3 with sham-links configured, an assert in tRtrPtcls can occur with the following log "SMS assert in qoamack.c at line 753 : != (if_cb->delayed_ack_list).prev 0x0x0 NULL 0x0x0".	8.01.01
The OSPFv3 default-information originate metric command incorrectly shows a range of 0-65535 instead of 1.	8.22.01

Persistence/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.tNVolCU]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:
For S-chassis (S1/S4/S6/S8) which contains S180 class modules, when there are more than 5 backplane fabric links are down, a customer warning message such as below should be displayed: <3>Dune[2.dTcmTask]WARNING: backplane performance degraded on fe_idx=0. (5-0x0bd65cf6-0x0bd65cf6-0x02d614b6-0x10010101-0x00000000)	8.11.01
"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 rcv for 6007 tics "(0x00d0f9e4 0x0067b420 0x006707ac 0x01683264 0x00000000)" while PoE Controller is being updated.	4.21.09
Replacement of a chassis line card in a Bonded chassis with the same hardware and firmware failed, resulting in a rolling DSI reset.	7.00.01
After slot removal and re-install, board present status is sometimes incorrect.	8.21.02
Setting the MAC age time to 10 seconds may cause the tNtpTmr task to use high amounts of CPU processing time.	8.21.01
MAC addresses were not learned before setting up connection. Once addresses are resolved the reframe RIB programs and CNX is established, message is removed.	8.20.02
An incomplete core file will be left after a reset for the following error message: <0>System[4]watchDogTask() application Chassis Coherency (2) failed to run in 300 seconds (0x00c77d1c 0x0045fcb0 0x01830ea0 0xeEEEEEEEE).	1.07.19

Platform Problems Corrected in 8.31.01.0006	Introduced in Version:
Chassis may reset without any indication if too many boards are inserted causing the consumed power to exceed the chassis power supply available power configuration limit.	7.00.01
The CLI command 'show system' continues to display stale view of S-AC-PS after blade swap.	7.00.01
S140/S180 class modules might reset with message similar to: "<0>Fuji[4.tNimIntr] Switch Chip 2 (Mainboard) detected fatal condition. (0x00c6c454 0x015155f8 0x0150f404 0x0215fbfc 0x0044489c 0x00465f04 0x004678d0 0x017f7ee0 0xeeeeeeee)".	8.01.01
For S-chassis (S1/S4/S6/S8) which contains S150/S130 class modules, a backplane fabric link may rarely get into a bad state and cause packets dropped in fabric.	7.00.01
For S-chassis which contains S150/S130 class modules, a bad backplane fabric link may prevent the blade from booting up.	7.42.02
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 ifInDiscards counter for a front-panel port inadvertently reflects the number of MAC Pause frames received on that port (SK2008, SK2009, SK8008, SK8009, and SL8013 cards only).	8.01.01
When an S-Series system is initialized with SFP modules (operating at 1Gbps) inserted into multiple ports of a SOK2208-0104 or SOK2208-0204 option module or 10G ports of S130, S150 or S155 class module, one or more of these ports may fail to successfully link with its respective peer.	8.21.01

PoE Problems Corrected in 8.31.01.0006	Introduced in Version:
PoE configuration might be lost when performing 'configure append' operation while PoE Controller is being updated.	7.70.00

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
On an S-Series blade assembly with 40G-capable ports (i.e., SL8013-1206, SL8013-1206-F8, SL8013-1206A, or SL8013-1206-F8A), if a TOS policy rule is configured which is mapped to a CoS level with drop precedence specified, the operation will fail and the blade will be rendered unbootable.	8.11.01

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]generateIfIndex():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:
BPDUs 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
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 the S-Series main board modules will not display any messages indicating it was pressed.	8.01.01
Tracked Objects Problems Corrected in 8.31.01.0006	Introduced in Version:
"Failed to set -101" error is seen during logging configuration.	3.11.04
Tunneling Problems Corrected in 8.31.01.0006	Introduced in Version:
After router failover, some filter connections may not be removed in L3VPN networks.	7.91.01
ICMP need fragmentation messages sent to a L2 tunnel source were not properly decoded and forwarded.	8.21.02
LSNAT is incompatible with tunnels.	7.22.01
NAT commands are not available on tunnel interfaces.	7.31.05

Tunneling Problems Corrected in 8.31.01.0006	Introduced in Version:
The TOS or Traffic Class 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
Layer 3 VPN filter connections created on router failover are not removed when new labels are sent to forwarding plane.	7.91.01

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
Ingress Service Provider tagged packets with Customers Q-tag and MPLS Labels are now supported.	7.91.01

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
A Neighbor Advertisement or Router Advertisement may be sent from an IPv6 interface using the hardware MAC address instead of the Virtual MAC address when VRRP is configured.	7.20.02
A VRRP critical IP address that is configured on a local interface may be reported as "down" when it is really up.	7.73.01
Adding or removing IP or IPv6 addresses to VRRP VRIDs may cause the Master IP address to be blank until the VRRP interface is disabled and re-enabled.	7.41.02
After a high-availability upgrade the router will not respond to pings to the VRRP Virtual IP address.	7.00.01
Cannot use an ACL Name that is 64 characters long for the "vrrp host-mobility-acl" configuration command.	7.00.01
IPv4 VRRP advertisements may be transmitted with a TTL of 64.	8.20.02
The Master advertisement interval is reported for VRRPv2 when using the "show ip vrrp verbose" command.	8.21.01
The master down timer for VRRPv3 incorrectly uses the VRRPv2 protocol values.	7.00.01
The router may generate "Unable to bind to address" syslog messages when an IPv6 VRRP VRID becomes master.	8.02.01
The router may not respond to echo requests on the backup router after enabling VRRP "accept-mode".	7.00.01
The task tVrrpEvt may cause a core dump and result in a module reset.	7.00.01
When a VRRP VRID in backup mode receives an advertisement with a priority of zero it should become master after skew_time but waits 3 * ADV_INTERVAL + skew_time instead.	8.21.01

VRRP Problems Corrected in 8.31.01.0006	Introduced in Version:
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 using the "probe-name" option for the "vrrp critical-ip" command the "remote" option is required but the CLI allows users to enter the command without the option causing the command to fail.	7.21.03
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
VRRP drops packets from routers that use an IP pseudo header for IPv4 packets when running VRRP v2-ipv4.	7.00.01

VSB Problems Corrected in 8.31.01.0006	Introduced in Version:
If a card that contains a VSB port in a Hardware bonded system is replaced with a card that doesn't have a VSB port, or the VSB port has been moved (from one option slot to another), it is possible that all the slots in the remote chassis will reset. This reset will only occur once.	8.11.01
When bonding is disabled on a VSB port, a message similar to: 'mazama[4.tDispatch]txPacket Invalid txBlock:255 txFujiBlock:15' maybe logged. There are no negative effects, other than the logging of the message.	8.11.01
When the set bonding port enable command is given a list of ports, ports may fail to be enabled. This error will be accompanied by the syslog message ""Bonding[7]The system is currently too busy to validate bonding port xx.yy.zz enable.	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

Platform Problems Corrected in 8.22.03.0006	Introduced in Version:
S140/S180 class modules might reset with a message similar to: "<0>Fuji[4.tNimIntr]Switch Chip 2 (Mainboard) detected fatal condition. (0x00c6c454 0x015155f8 0x0150f404 0x0215fbfc 0x0044489c 0x00465f04 0x004678d0 0x017f7ee0 0xeeeeeeee)".	8.01.01
When accessing the SSA180 class or the SSA150A, it is possible for a failure to occur causing the device to reset.	8.02.01

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

Feature Enhancements in 8.22.02.0012

Hardware Support Enhancements in 8.22.02.0012
ST8206-0848-F8A S-Series S180 Class I/O-Fabric Module, Load Sharing, - 48 Ports 10/100/1000BASE-T via RJ45 with PoE (802.3at) and two Type2 option slots (Used in S1A/S4/S6/S8) EEE Capable
ST2206-0848A S-Series S140 I/O Module, - 48 Ports 10/100/1000BASE-TX via RJ45 with PoE (802.3at) and two Type2 option slot (Used in S3/S4/S6/S8) EEE Capable

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

BGP Problems Corrected in 8.22.02.0012	Introduced in Version:
MD5 Authentication over IPV6 BGP peering sessions will not allow peering sessions to establish with third party devices.	7.30.01

Hardware Problems Corrected in 8.22.02.0012	Introduced in Version:
System may log the message: "Default[4.tNimIntr]Assertion failed: SOC_REG_IS_VALID(unit, reg), file /firmware/common/bcmStrataDrv/04_09_22/sdk/src/soc/common/.reg.c, line 2884 (0x00c77d1c 0x00414b2c 0x00f421f4 0x00f45d90 0x00d4be0c 0x00d43b9c 0x00447858 0x0046f450 0x01830ea0 0xeeeeeeee)" and reset.	8.02.01

IPv6 Problems Corrected in 8.22.02.0012	Introduced in Version:
A deferred IPv6 packet destined to subnet via a link-local nexthop address and deferred to neighbor discovery for MAC resolution may not be successfully transmitted when neighbor discovery completes.	8.22.01

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

LSNAT Problems Corrected in 8.22.02.0012	Introduced in Version:
When packets are Fragmented, only the first packet should change the L4 Header, filters have been added to accomplish this.	7.00.01
Management Problems Corrected in 8.22.02.0012	Introduced in Version:
Console session or telnet session gets stuck after login.	4.11.18
A message similar to "serverWatchDog.1(Config), client 75(ifMIB) in recv for 6058 tics" may be logged (followed by a system reset) when a card with SFP+ ports is booted and one or more of these ports contain SFP modules.	8.21.01
MPLS Problems Corrected in 8.22.02.0012	Introduced in Version:
A reset occurs when the MPLS/L3VPN data structures are exhausted.	8.21.01
In MPLS with minimum of two P routers in network, ICMP errors generated by all but the last P router are not successfully transmitted to the source of the packet causing the error. For example, traceroute will report loss of packets.	8.20.02
'mpls ip propagate-ttl' and 'mpls ldp-label-allocate' settings are not cleared when the command 'clear router vrf global' is executed.	8.22.01
Multicast Problems Corrected in 8.22.02.0012	Introduced in Version:
For S140/S180 class of products, one may see these error messages when fabric back plane is oversubscribed: <165>Jun 18 08:52:43 192.168.100.11 Dune[5.tDuneErrM]Petra[1] Interrupt PB_IPS_CREDIT_OVERFLOW instance 0 is off <165>Jun 18 08:52:45 192.168.100.11 Dune[5.dTcmTask]Petra[1] Received Interrupt PB_IPS_CREDIT_OVERFLOW instance 0, count 47377, value= 0x1	8.11.01
Changing remote route to connected route may result in an assert similar to "sms[1.tRtrPtcls]SMS assert in qptufsms.c at line 1259".	8.20.02
OSPF Problems Corrected in 8.22.02.0012	Introduced in Version:
When OSPF routes are filtered from the RTM using a distribute-list, if a route that matches the filtered route, but does not match the distribute-list route-map, is introduced into the RTM, it becomes stale and can never be removed.	7.00.01
OSPFv3 Problems Corrected in 8.22.02.0012	Introduced in Version:
If running OSPFv3 and attempting to aggregate a range with more than the first six bytes containing a non-zero value will cut off bytes and neither aggregate correctly nor display the range correctly.	7.30.01

Persistence/Nonvol Problems Corrected in 8.22.02.0012	Introduced in Version:
When seeing the following error: <3>Default[1.tusrAppInit]moduleIsInSameLocation():Unable to open "/flash 2/moduleRecords/chassisSlotInfo.rec" for reading is usually do 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

Spanning Tree Problems Corrected in 8.22.02.0012	Introduced in Version:
When the spanning tree version is set to stpCompatible, it is possible for a multi-blade or stacked device to reset due to a watchdog timeout. The workaround is to set the version to any other value, which are all backwards compatible. The only configuration that requires the setting of stpCompatible is when an attached device will malfunction due to trying to process a type 2 BPDU (RSTP or higher) or due to not receiving an STP BPDU in a timely manner. It takes approximately 3 seconds for the Port Protocol State Machine to recognize a legacy device and switch to transmitting type 0 BPDUs.	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

VSB Problems Corrected in 8.22.02.0012	Introduced in Version:
A blade introduced into a bonded system may not be able to get a copy of the running image and get stuck in a reboot loop.	7.00.01
VSB enabled module stuck in a reboot loop, disabling then re-enabling VSB when added to VSB disabled chassis. This failure will occur when the module is installed in the same location it last occupied. A workaround is to first move the module to a different location. By doing this the module's previous VSB configuration will be deleted.	8.21.01
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.

BGP "Pass Through" Route Feature Target Support Enhancements in 8.22.01.0022

This enhancement provides the ability to adjust the route targets applied to routes exported from a VRF to the BGP backbone in an L3VPN network. Functionality includes the ability to merge existing route-targets with export route targets configured on a VRF or to replace export route targets configured on a VRF with the existing (pass through) route-targets.

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 system sends packet to a remote IP address, an ARP request for the remote IP address may be transmitted on a configured interface.	8.21.01
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

Auto-Tracking Problems Corrected in 8.22.01.0022	Introduced in Version:
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

BGP Problems Corrected in 8.22.01.0022	Introduced in Version:
When MPLS is disabled, established state with BGP peers are lost.	8.02.01
The S-Series router currently does not have a mechanism to replace the export route targets defined on a VRF with the existing route targets on an L3VPN route.	8.01.01
The 'show ip protocols' command output does not display the BGP max-as limit.	8.21.01
The "show ip bgp" output of the AS-Path will display incorrect AS numbers if the AS-Path is longer than 30 AS numbers.	7.20.01
The BGP Autonomous system number of 0 is accepted at the CLI even though the help indicates the minimum value is 1. In this case, "show configuration" output will not display the "router bgp <AS>" command.	7.20.01
The "show ip bgp <prefix> detail" command will display repeated instances of the same community and extended-community values in some cases.	7.91.01
The 'show ip bgp peer <ip> advertised-routes <prefix/length> detail' command does not always display the correct communities and extended-communities associated with the route.	7.20.01
Negating the BGP peer-group soft-reconfiguration command does not take effect. The show running-config output will indicate the command is negated, however the setting is not negated internally.	7.20.01
Redistribution of ISIS into BGP under non-vrf address-family mode will result in show running-config output that is inconsistent with the required command syntax for the "match" path type options.	7.30.01
The following error message may occur if deleting an instance of a routing protocol which contains redistribution entries with multiple references to the same route-map: "Error decrementing route map <name>."	7.20.01
The output of the "show ip bgp" command does not display any information under the AS-Path heading if the actual patch contains approximately 70 or more AS numbers.	7.22.01

Chassis Bonding Problems Corrected in 8.22.01.0022	Introduced in Version:
During a time of chassis instability, a module in a Bonded chassis may reset after logging an error with format similar to: "<0>Default[12.tBondProto]Assertion failed: hdr>reqGeneration == generation, file /firmware/common/chassisBond/01_06_16/src/chassis_bond_protocol.cxx, line 599".	7.72.01
A module in a Bonded chassis may reset soon after power up. When this failure occurs, a message with format similar to: "Default[14.tBondProto]Assertion failed: hdr->reqGeneration == generation, file /firmware/common/chassisBond/01_05_15/src/chassis_bond_protocol.cxx, line 599" is logged.	7.72.01
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

Chassis Bonding Problems Corrected in 8.22.01.0022	Introduced in Version:
Bonding mode may be changed from software-assist to hardware at boot with mix of 8.12 and older firmware images in the chassis. When this occurs modules will reset and a message with following format will be logged: "Received Bonding mode = hardware from master. Rebooting ..."	8.11.01
A module in a software bonded system may reset while bonding ports are being enabled. A message similar to "1>DistServ[2.tDsBrdOk]serverWatchDog.6(PortInfo), client 106(Bonding) in recv for 6300 tics" is logged on this failure. A workaround is to wait 1 minute between bonding port enables.	8.21.01
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
ECMP Problems Corrected in 8.22.01.0022	Introduced in Version:
If an interface that is part of an equal cost multipath route goes down, host originated traffic to destinations in the route's subnet may temporarily fail.	8.21.01
Host Services Problems Corrected in 8.22.01.0022	Introduced in Version:
Some devices may reset after logging a message similar to the one listed below. This may occur intermittently on S140 modules during initialization of the onboard power controllers. Message 6/213 Exception PPC750 Info 08.11.04.0005 01/09/2014 11:14:31 Exc Vector: DSI exception (0x00000300) Thread Name: tRootTask Exc Addr: 0x0168ba70 Thread Stack: 0x7dffd10..0x7dfec7c0 Stack Pointer: 0x7dff4f0 Traceback Stack: GENERAL EXCEPTION INFO.....	8.11.03
Messages like the following can be seen during slot resets in busy systems. Transmit errors(8) to slot # are preventing heartbeat checks.	7.72.01
No eligible master messages are misleading because slot/peer ID is mistaken for server ID.	7.03.05
"show system utilization storage" will report inaccurate size and available size for USB drives greater than 2G.	7.60.01

Host Services Problems Corrected in 8.22.01.0022	Introduced in Version:
<p>When updating to a new image that has microcode updates there are error messages displayed about the DOSFS/DOS volume. Example:</p> <pre>***** This board is updating its hardware microcode. This process may take up to 3 minutes to complete. Please do not reset or disconnect power during this update. The board will be reset when this process is complete. ***** Writing primary host controller : bank 1 block 0 Programming flash. Complete: 100% Verify succeeded. Writing alternate host controller : bank 0 block 0 Programming flash. Complete: 100% Verify succeeded. /flash0 Not a DOSFS Device /flash1 Not a DOSFS Device /flash2 Not a DOSFS Device <163>Dec 18 08:27:57 0.0.0.0 System[1]Resetting after host controller microcode update. <163>Dec 18 08:27:57 0.0.0.0 System[1]DOS Volume /flash0 cannot be set to read-only on unmount. <163>Dec 18 08:27:57 0.0.0.0 System[1]DOS Volume /flash1 cannot be set to read-only on unmount. <163>Dec 18 08:27:57 0.0.0.0 System[1]DOS Volume /flash2 cannot be set to read-only on unmount.</pre>	8.21.01
<p>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:</p> <pre>/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.</pre>	7.30.01
<p>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.</p>	7.40.00
<p>The "show running-config" command may not display all static ARP/ND entries that are configured.</p>	7.00.01
<p>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.</p>	8.21.01
<p>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:</p> <pre>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.....</pre>	7.00.01

Host Services Problems Corrected in 8.22.01.0022	Introduced in Version:
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 0x00a81504 0x01686324 0x00000000) A core file will be generated.	8.21.01
<p>Infrequently, when the switch is adding (encapsing) 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</p> <p><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</p> <p><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.</p>	7.40.00
IGMP Snooping Problems Corrected in 8.22.01.0022	Introduced in Version:
IGMP/MLD database entries (primarily, but not limited to IGMPv3/MLDv2 reporter state) do not age out correctly.	7.30.01
Legacy S-Series modules (S130/S150) with IGMP/MLD snooping enabled log messages similar to the following: <188>Jan 6 07:26:20 172.20.1.20 RfrmrHw[3.tDSrecv2]Invalid MCI - 1, for ASIC 0 <188>Jan 6 07:26:20 172.20.1.20 RfrmrApp[3.tDSrecv2]addPortReframing, Error: Could not convert mciIndex 8113 to UNTAGGED error status -2, then reset with a DSI exception in thread tDispatch.	8.21.01
When using IGMP over 40G ports, data may not reach the proper egress.	8.11.01
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
Some IPv6 addresses may remain in the tentative state when the master blade changes from one slot to another.	7.30.01
The IPv6 /128 host address of tunnel interfaces appears in output of 'show ipv6 route'.	8.21.01
IPv6 Neighbor Discovery Problems Corrected in 8.22.01.0022	Introduced in Version:
The router may not accept router advertisements to generate IPv6 addresses when the "ipv6 address autoconfig" command is applied to an interface.	8.21.01

LAG Problems Corrected in 8.22.01.0022	Introduced in Version:
Prior to 8.11.1 S140/S180 cards 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 on S140/S180, and from initial S release on Legacy S cards, if you have LAGs with more that 16 ports, Source Port Exclusion may not work.	8.01.01
LSNAT Forwarding Problems Corrected in 8.22.01.0022	Introduced in Version:
It is possible that while processing using sticky entries on a multiple blade system, that the sticky entry may not be deleted from all blades and subsequent sticky creations will fail causing a failure of processing LSNAT packets.	6.12.01
LSNAT FIN/RST timeout may not be properly set while running with multiple blades.	8.11.01
During debug session the command <code>"/* rtr lsnat show data-plane bindings detail <ID>"</code> caused a reset.	6.00.02
Mirroring Problems Corrected in 8.22.01.0022	Introduced in Version:
Remote mirroring may not operate correctly on a 10Gbase-T port.	8.02.01
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
MPLS Problems Corrected in 8.22.01.0022	Introduced in Version:
The command <code>'no mpls ip propagate-ttl [local]'</code> did not affect packets originated by the local host. The packet's TTL was propagated to the MPLS label.	8.21.01
Learning the internet route table from BGP with LDP configured will exhaust system memory and cause a reset.	8.21.01
First hop may not respond when issuing traceroute from PE router across MPLS network.	8.21.01
The wrong MTU is specified in an ICMP Fragmentation Needed packet sent by an MPLS provider edge router for packets egressing an LSP and exceeding MTU of egress port.	8.12.01
Configuring LDP with the internet route table present in the system will exhaust memory and cause a system reset.	8.21.01
<code>'mpls ip propagate-ttl'</code> settings not stored in persistent storage.	8.21.01
Given MPLS/LDP enabled in the system, if it is disabled and enabled again, additional FECs are unnecessarily created for connected subnets.	8.21.01
When both IPv4 and IPv6 prefixes were in use with label switched paths, display commands showing the MPLS forwarding table contained invalid characters when trying to interpret the next hop addresses.	8.21.01

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

NAT Problems Corrected in 8.22.01.0022	Introduced in Version:
It is possible to for NAT to stop working due to running out of local buffers.	8.11.01

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
Using OSPF with a route-map for redistribution that sets the metric to a number greater than 65535 will result in an assert in thread tRtrPtcls with the following log in OSPFv2: "SMS assert in qodmbld3.c at line 471 : == (NBB_INT ((route_entry_ptr->path_cost) >> 24) 255 nbb_zero 0" and log "SMS assert in qod3bld2.c at line 214 : == (NBB_INT)((route_entry_ptr->path_cost) >> 24) 255 nbb_zero 0" for OSPFv3.	7.00.01
If the display of OSPF passive-interfaces in show running exceeded 80 characters, no interfaces are displayed.	8.21.01

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

OSPFv3 Problems Corrected in 8.22.01.0022	Introduced in Version:
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:
Enabling a PIM-DM upstream interface may result in an assert similar to "sms[1.tRtrPtcls]SMS assert in qptuftmr.c at line 1134 : (null) NTL_CLTIM_IN_LIST(&s_g->sg_join_timer) 0 (null) 0".	8.21.01
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
Rebooting PIM DM source router may result in an assert similar to sms[1.tRtrPtcls]SMS assert in qptuwapi.c at line 602 : == 0 0 (s_g_i->sgi_flags & QPTM_DSM_SGI_WP_ALL) 4096.	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:
Infrequently a board will not boot up and will end up in a halted state after a failure to read chassis type. The following message is output to the console when this error occurs: "A device within this chassis has encountered a hardware failure. Could not read chassis type. Please contact Support for the troubleshooting procedure to determine which device will possibly need to be repaired/replaced. Press <r> to reset board." This may not be a real hardware failure and a module reset will result in successful module initialization.	8.01.01
A watchdog timeout exception message may be logged (followed by a system reset) when a card with SFP+ ports is booted and one or more of these ports contain SFP modules.	8.21.01
Setting a 40Gb port to 4x10Gb using either "set port speed fg.x.x 10000" automatically by inserting a 40Gb fan-out cable into a fg port, the configuration will not be saved and will not display in "show config port".	8.01.01

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:</p> <ul style="list-style-type: none"> Smon stats Rmon Host/Matrix Router ACL Policy Routing Tunneling Policy <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>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/usrroot/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
10G port with 1G SFP doesn't propagate its advertised speed to link partner.	8.11.04
<p>A S140/S180 blade may display messages similar to the following when backplane fabric is oversubscribed.</p> <pre><165>Apr 29 11:36:11 0.0.0.0 Dune[4.dTcmTask]Petra[0] Received Interrupt PB_SCH_FCTFIFOOVF instance 0, count 12, value= 0x1 <165>Apr 29 11:36:17 0.0.0.0 Dune[4.tDuneErrM]Petra[0] Interrupt PB_SCH_FCTFIFOOVF instance 0 still active <165>Apr 29 11:38:07 0.0.0.0 Dune[4.tDuneErrM]Petra[0] Interrupt PB_SCH_FCTFIFOOVF instance 0 is off.</pre>	8.11.01
If a 1G SFP is inserted into one of the 10G ports on a SOTK2268-0212 or SOGK2218-0212 option module, the system will reset.	7.91.01
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.	7.91.01
When doing a "dir" from CLI, if the directory is currently being modified (file being added/deleted) an incomplete listing can be returned.	7.91.01
During a system reset or a module reset, removal or insertion, it is possible to receive a DSI exception containing the text "DuneCB::RemoteSlot". This exception can be ignored once the system completes normal initialization.	8.11.01

PoE Problems Corrected in 8.22.01.0022	Introduced in Version:
POE redundancy shown as Not Supported after POE blade is reset and boots up.	7.60.01

Policy Problems Corrected in 8.22.01.0022	Introduced in Version:
Unable to clear all policy profiles with a single CLI command.	1.07.19
VLAN authorization commands allow for configuration and display of tunnel bridge ports although they are not supported port types for VLAN authorization.	8.21.01
Policy "macsource" rules configured with a mask less than 48 are not applied to traffic immediately upon configuration.	8.21.01
PWA Problems Corrected in 8.22.01.0022	Introduced in Version:
The "set pwa ipaddress <ip-address>" CLI command allows invalid values for the <ip-address> field.	4.00.50
PWA occasionally becomes unresponsive under heavy load. Device resets with this message in the log: <0>PWA[1.tPwaHtWD]pwaHttpReadWatchDog expired!	4.00.50
QOS Problems Corrected in 8.22.01.0022	Introduced in Version:
Default port group COS txq settings are applied to hardware VSB ports. No COS settings, default group or not, should ever be applied to hardware VSB ports.	8.11.01
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:
In rare instances, upon a blade reset in a multi-blade system with a large number of RMON alarms configured DSI exception resets may occur.	5.01.58
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

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 <code>_none_</code> , the non-master blades will still operate as if the mode were <code>_ieee8021_</code> 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 reenables 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
A static host route whose address matches an LSNAT virtual server address will not be displayed in the router configuration if the LSNAT virtual server is up.	7.00.01

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 were not updating the L2 IP's total length field, Dwhen adding the GRE header and Chassis Bondheader to the egress packet.	8.21.01
If tunnels are configured and at least one is up, then the connection database is no longer flushed when a route changes.	8.21.01
The ifMib returns a valid ifIndex row with no other valid leaves for internal ports that should be hidden.	8.21.01
Traceroute does not work from layer 3 VPN when configured over IP tunnels.	8.21.01
Infrequently, when switch is adding (encapng) tunnel headers, a message similar to: D <0>chassis[9.tBcastStRx]powerSupplyComputeModuleConsumedPower :Invalid uplink number 0x00 detected in remote info table, may be logged	7.40.00

Tunneling Problems Corrected in 8.22.01.0022	Introduced in Version:
Infrequently, when switch is adding (encapsing) 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:
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
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
It is possible for show running to erroneously display "set router vrf vrf-management <VRFNAME> disable".	8.21.01

Feature Enhancements in 8.21.02.0001

Hardware Support Enhancements in 8.21.02.0001

SK8208-0808-F8 I/O fabric module – Support for the new S180, 8 Ports 10GBASE-X via SFP+ and two Type2 option slot module.

IEEE 802.3az Enhancements in 8.21.02.0001

Energy Efficient Ethernet (EEE) – Supports the ability to reduce power consumption in Ethernet networks during link idle periods by providing a mechanism to change/lower a ports always-on transmitter during idle periods. EEE is dependent on the physical layer (PHY) of networking equipment and is supported the following S-Series models. SK8009-1224-F8, SK8009-1224, SK2009-0824.

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 (VLANs) 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.

BGP Route-Flap Dampening Enhancements in 8.21.02.0001

Support for BGP Route-flap dampening to suppress routes that are being repeatedly advertised and withdrawn (flapping) due to mis-configuration or a badly behaving (i.e. rebooting or a link flapping) router.

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 190 LAGs in multislot S-Series chassis.

Tunnel Enhancements in 8.21.02.0001

The S-Series IP tunnel capacity has been increased to 62 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.

Remote port mirroring is now supported when user ports are used to create a “software” bonded VSB. Previously hardware VSB ports were required to bond chassis together AND use the remote port mirror feature.

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' lists all multicast groups the VLAN has joined.

A command to disable DHCP server logging has been added.

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

'show running bgp' has add a keyword to only display modal configuration.

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 a 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
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 affect as only the IPv4 Threshold Type is currently supported.	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

ARP Problems Corrected in 8.21.02.0001	Introduced in Version:
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
For LAG ports that are composed of 40G Ethernet ports, if one or more of the 40G ports detach, the LAG port may not reliably switch traffic afterwards.	8.11.01

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:
Help string for auto-tracking port radius-reject-profile command is incorrect.	8.01.01

BGP Problems Corrected in 8.21.02.0001	Introduced in Version:
The BGP network command for the default route (0.0.0.0/0) will not inject the route if a redistribution command exists which uses a route-map/access-list combination and this access-list does not permit the default route. To ensure that the route is injected the access-list must permit the default route.	7.20.01
BGP peering sessions may time out due to the deletion of internal connections when ACL's are deleted.	7.20.01
When redistributing loopbacks from one VRF to another on the same PE the loopbacks will not be redistributed.	7.91.01
A BGP ORF route-refresh message is not sent in a Layer-3 VPN network if the neighboring router reboots and the peering session supports graceful restart.	8.01.01
Filtering of BGP routes based on the length of the AS-PATH is not supported.	7.20.01
BGP ORF capability for multicast routes is not supported. However the CLI allows the user to enable this capability. Attempts to disable the capability will result in the following error message: "Error:Command Failed - :No such object: Setting orf admin_status"	7.20.01
The 'show ip bgp groups' command always displays the BGP peer-group address-family as IPV4/Unicast.	7.30.01
The "show ip protocols" command output is missing the following BGP related information: 1. The default values of peer based timer related variables are missing. 2. Redistribution of ISIS into BGP is not shown.	7.20.01
The 'debug ip bgp notification' log messages do not display text descriptions of the numerical error code and subcode for sent and received BGP notification messages.	7.20.01

BGP Problems Corrected in 8.21.02.0001	Introduced in Version:
The BGP "neighbor <IP address> clear-counters" command resets the counters the first time the command is issued for a given peer. Subsequent attempts do not clear the counters.	7.20.01
The output of the 'show ip bgp neighbors' command displays the established time in seconds instead of the "day, hours:minutes:seconds" format.	7.20.01
The BGP route-map match and set parameters will appear in the show running-config output with their default values after they are negated.	7.20.01
The BGP route-map "set extended-community ospf-route-type" command error message indicates the range of valid values is 1-7. The valid route-type values are actually 1,2,3,5, and 7.	7.20.01
Negating the "set community" or "set extended-community" clause from a BGP route-map yields two copies of the error message if the "action" keyword is incomplete. For example if the abbreviation for "remove" is entered for the action.	7.20.01
BGP prefix lists configured with sequence number 65535 are stored with sequence number 0 and can't be deleted.	7.20.01
The 'show ip bgp' command output does not display the value of the weight attribute.	7.20.01
The output of the 'show ip bgp neighbors' command displays the established time in seconds instead of the "day, hours:minutes:seconds" format.	7.20.01
A system reset may occur if an IPV6 BGP peer is disabled and BGP route aggregation is configured. The following error message will appear when the system resets: SMS assert in qbpmreca.c at line 233 : != *old_route 0x0x0 NULL 0x0x0	7.30.01
The BGP neighbor activate command will not appear under the vpnv4 or vpnv6 address-family configuration if the address-family has not been enabled.	8.01.01
The 'show ip ipv6 bgp summary' output display of the time in established state and time since last message received is in seconds instead of days, hours:minutes:seconds format.	7.20.01
The 'show running-config' and 'show config' output displays extra exclamation points in the BGP section.	7.20.01

Bonding Problems Corrected in 8.21.02.0001	Introduced in Version:
Under heavy traffic conditions bonded system may see "failed to send message - Buffer Full" in the message log.	7.61.02
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
Cabling a VSB port to an ethernet port may cause modules to reset, and a message similar to " <0>Bond[14.tDispatch]getVsbInPort: learn inport:000037e3 outport:00000be7 binding failed" logged.	8.11.01
Configuration of RMON stats and history options on existing default entries will be lost on reboot.	8.11.01
Sometimes after a configure, bonding ports that are attached to a partner port are not activated for bonding. When this happens the following message is logged on the partner chassis for the partner port: "No Bond Partner found on port x.y.z, possible misconfiguration." To fix this issue disable the port for bonding, then re-enable it.	7.72.01

Bonding Problems Corrected in 8.21.02.0001	Introduced in Version:
In a bonded system, with logging for bonding set to debugging(8), messages with following format are infrequently logged: - Bonding[1]Starting inter-module communication to bonded slot <slot> - Bonding[2]Received first inter-module communication from bonded slot <slot>	7.72.01

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

CLI Problems Corrected in 8.21.02.0001	Introduced in Version:
On bonded systems while copying files from one blade to another or off the system and the bond link goes down, the master blade could reset/DSI.	7.60.01
Syslog message status is OK when setting port duplex and speed fails.	7.70.00
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
Auto-tracking and quarantine-agent CLI is unable to set 40 Gigabit and 100 Gigabit port settings.	8.01.01
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
COS ORL rates less than or equal to 30pps may not work if the traffic is received on a remote blade.	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
With DVMRP configured, the Management Module (Router) resets with a message similar to the following: "SMS assert in ntlcltim.c at line 547 : < duration -296."	7.00.01

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 the maximum amount of MAC entries is attempted to be set to 128K, but all blades in chassis do not have required 2G of memory, the cli command correctly fails. However, the status returned is OK rather than ERROR.	7.91.01
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
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/MLD IP Multicast traffic only utilizes a single underlying physical port of a LAG in a Bonded system.	7.60.01
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: "Cigmp::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

IGMP Problems Corrected in 8.21.02.0001	Introduced in Version:
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
IGMP and MLD frames ingressing at excessively high rates on VLANs that do not have IGMP/MLD enabled, but have a Layer 3 interface can cause system instability including module/system resets.	7.00.01

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
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
Route table updates may result in layer 3 VPN packets using VPN labels present before route updates occurred.	8.11.01

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

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

IS-IS Problems Corrected in 8.21.02.0001	Introduced in Version:
Cisco LSPs are sometimes displayed incorrectly.	8.01.01
'show isis hostname' for a level-1 router displays hostnames for level-2 router instances.	8.01.01
ISIS hostnames do not appear in LSP Summary database.	8.01.01
'show isis topology' does not display configured hostnames.	8.01.01

IS-IS Problems Corrected in 8.21.02.0001	Introduced in Version:
Within a VRF a 32 bit summary address for ISIS is incorrectly displayed in the running config and cannot be negated.	7.73.01
When displaying ISIS LSP database information having Cisco routers, the multiple metric and IP information may be strung together.	8.01.01
When redistributing RIP into ISIS, the wrong metric is displayed.	8.01.01
When displaying the ISIS LSP database, the wide metric values shown are incorrect.	8.01.01
ISIS LSP database wide metrics are displayed incorrectly.	8.01.01
Deleting 'isis lsp-gen-interval' will result in the default value (in seconds) showing up in show running config.	8.11.01
'show running isis' may display additional blank line between authentication statements.	8.11.01
ISIS LSP databases with metrics associated with multiple IPs get displayed incorrectly.	8.01.01
When the connection between ISIS neighbors is tunneled and over-subscribed, the neighbors periodically lose connectivity.	8.01.01
ISIS database display shows an incorrect format after metric IS.	8.01.01
After 'no ip router isis' is done on an interface, hellos may continue to be sent.	8.01.01
When redistributing into ISIS the show running command displays command twice.	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

L3VPN Problems Corrected in 8.21.02.0001	Introduced in Version:
Using L3VPNs when BGP is attached to a route-reflector client will result in routes not propagated to the VRFs.	8.01.01
Creating domain-ID (primary or secondary) with an invalid <6 octet domain id> creates on with FF instead of returning an error message	8.11.01

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
A set of a lag port attribute may fail without a message at the console.	1.07.19

LACP Problems Corrected in 8.21.02.0001	Introduced in Version:
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
For LAG ports that are composed of 40G Ethernet ports, frames flooded out the LAG may be sent to more than one of the 40G ports in the LAG. Frames received on the LAG that need to be flooded may also be sent back out the LAG.	8.11.01

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

LSNAT Problems Corrected in 8.21.02.0001	Introduced in Version:
In a previous release, access to a VIP server from a VRF via a route was not allowed without the 'all_vrfs' configuration command option defined on the VIP server. Access to a VIP from a VRF via a route leak is now allowed.	7.00.01
In a previous release after modifying the global NAT SLB or TWCB binding limit, it may cause runtime issues while processing bindings.	6.12.05
"show config", "show running", and "show ip slb info" will not display the "real server access client" configuration lines.	6.12.08

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:
Infrequently, a chassis module with port mirrors configured resets. On this failure a message similar to "PortMirr[3.tDSrcv1]processMirrorDestination(1,61013): index mismatch detected: smon=2, mirror=5" is logged.	7.41.02
Port mirroring may reset and log a message similar to "PortMirr[12.tDSrcv1]setMirrorIndex(72028,162024): invalid mirror index transition 2->1".	7.00.01

Mirroring Problems Corrected in 8.21.02.0001	Introduced in Version:
The "clear port mirroring orl" 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
In a software VSB chassis, if mirrors are active, very infrequently a blade may reset, leaving a message similar to: "<1>DistServ[12.tDsBrdOk]serverWatchDog.1(Config), client 88(mirrorMgr) in rcv for 6721 tics".	7.60.01

MSDP Problems Corrected in 8.21.02.0001	Introduced in Version:
Multiple removing MSDP by 'no ip msdp originator-id' causes a crash similar to: 'SMS assert in qptujms.c at line 257 : == msdp_mj_cb-> j_mj_cb.mj_row_data.oper_status 8 AMB_NPG_OPER_STATUS_GOING_UP 3 '	8.02.02

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
Quarantine Agent Multiauth Sessions may not idle timeout as expected when port roaming.	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

NAT Problems Corrected in 8.21.02.0001	Introduced in Version:
No counter for NAT packets and NAT drop packets are included in "debug packet show-statistics".	1.07.19
It is possible that once the "ip nat log translations" config has been entered that it will remain persistent even when a "no ip nat log translations" command is entered.	6.00.02

NAT Problems Corrected in 8.21.02.0001	Introduced in Version:
<p>It is possible when upgrading from 08.02.xx or downgrading to 08.02.xx that the following NAT config may be lost:</p> <pre>ip nat ftp-control-port ip nat log translation ip nat inspect dns ip nat translation max-entries ip nat translation 'timeouts'</pre>	8.11.01
NETFLOW Problems Corrected in 8.21.02.0001	Introduced in Version:
<p>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 any where from 1 to 30 netflow records, is dropped and will not be delivered to NetFlow collectors.</p>	8.01.01
<p>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.</p>	8.01.01
<p>NetFlow cannot be enabled on 40G Ethernet ports.</p>	8.01.01
<p>When the 'set default-nexthop[-v6]' option is used in route-maps and NetFlow is in use, a reset may occur when route updates are being processed.</p>	7.20.01
<p>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 generating with invalid source and/or destination interfaces.</p>	8.01.01
OAM Problems Corrected in 8.21.02.0001	Introduced in Version:
<p>Disabling OAM on a port does not clear the OAM or ULD operstatuscause</p>	7.30.01
OSPF Problems Corrected in 8.21.02.0001	Introduced in Version:
<p>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"</p>	7.00.01
<p>A tunnel interface running OSPF will default to network type POINT_TO_POINT. If it is explicitly configured as POINT-TO-POINT and then removed, it defaults to BROADCAST instead.</p>	7.41.02
<p>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.</p>	7.41.02
<p>OSPFv2 will accept the configuration of an invalid nssa-range and display it incorrectly.</p>	7.00.01
<p>Configuring an OSPF cost metric outside the range results in an unclear message error.</p>	8.11.01
<p>When issuing a "clear ip ospf process" and multiple OSPF processes exist, the ambiguous message "Resetting the OSPF process" is seen multiple times.</p>	7.00.01

OSPF Problems Corrected in 8.21.02.0001	Introduced in Version:
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
sham is spelled "shaml" on both OSPFv2 and OSPFv3 debug	8.11.01
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
On bonded systems while copying files from one blade to another or off the system and the bond link goes down, the master blade could reset/DSI.	8.11.04
When using a route map when entering the 'redistribute bgp global' command, the route type will be changed to E2 and not use the correct domain id type.	8.11.01
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

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.	8.11.01
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.	

PKI Problems Corrected in 8.21.02.0001	Introduced in Version:
"If a configuration file which contains PKI data is modified by an external text editor and that editor adds control charactersD (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 the 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
The following syslog messages may be seen on bonded systems when the remote blade is under heavy load or the remote system is resetting/booting: 'NonVol[1.tNvBulk]nonvol_copy: Copying of redundant store will need to retry (FIOSYNC of outFd failed errno(errno = 0x300005))' These are log level 5(NOTICE) messages and usually only an issue if persistent.	7.30.01
Setting port speed on 1G fiber port to a speed the SFP does not support, and disabling auto negotiation causes the port to go down and stay down.	7.91.01
Releases before 8.20.01 may fail at initialization time with large configurations causing continuous resets caused by denoted rdyToSwitch fault.	8.11.01
"set port mdix" CLI commands may display "Error: failed to set mdix configuration of swap on port tg.5.2" on RJ45 ports that don't support mdix commands.	8.01.01
Unsupported Option Module will halt the board and not allow software upgrade until removed.	8.11.01
System may log a message similar to: "<1>DistServ[1.tDsBrdOk]serverWatchDog.5(Host), client 26(Emanate) in rcv for 6446 tics" and then reset.	7.60.01
If a 40g fan-out cable is inserted into a port set for 1x40g mode with a 1x40g transceiver inserted in the paired port, "show port status" on 40g ports will not display as "invalid" for the 40g fan-out cable and port will remain up.	8.11.01
For a S180 dual mode 40g port, if "set port speed tg/fg.x.x 10000" was used to change the port speed to 4x10g mode, using the same port type "set port speed tg/fg.x.x 40000" could not be used to change the mode back to original speed.	8.11.01
When changing from 4x10g mode to 1x40g mode using "set port speed tg.x.x 40000" only the individual port entered was reported in message to be changing when in fact the entire range of quad ports associated with the port tg.x.x-y change.	8.11.01
If a dual mode port has been set to 4x10g mode using ""set port speed fg.x.x 10000"" trying to set it back to original speed 1x40g using the opposing tg port ""set port speed tg.x.x 40000"" says that it will change to original 1x40G speed on the next reset, but in fact will not change and end up in 4x10g mode.	8.11.01
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

Platform Problems Corrected in 8.21.02.0001	Introduced in Version:
Issuing a "set port speed tg.X.1-8 40000" command to change port speed from 10G to 40G which encompasses two paired ports, will result in duplicated speed change messages for each port.	8.11.03
10GB-LRW-SFPP Laser Wire transceivers not acquired through Extreme that do not have a "-EN" part number will display "sfpDataAccess: CI2CBus::Access() failed to write transceiver data for slave 0xa2 on NIM x port y" message.	8.11.01
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
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/<xxx>.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
'show port status' is missing speed and duplex for 10G copper ports without a link.	8.01.01
No message is logged indicating a reason for board shutdown due to over temperature on S-Series.	7.00.01
"At boot a board could get into reset loop with the following syslog output: 'NonVol[1.tusrApplnit]Nonvol reached max fileldx 4080, storeNum 1, major 1'. Sets will be dropped until space is freed."	7.30.01
Unsupported speed of 100M is allowed to be set for 1G SFP.	7.91.01

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:
8.11 RADIUS Enhancement documentation does not clearly indicate that round-robin handling occurs on a per blade basis.	8.11.01
RADIUS authentication server max-sessions configuration is not output as part of "show config" or "show config all" commands.	8.11.01

RADIUS Problems Corrected in 8.21.02.0001	Introduced in Version:
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 authentication, authorization, and accounting server configuration may be lost upon upgrade from any release prior to 7.40 to any release post 8.02.	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
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
Setting radius-snooping port configuration for unsupported ports may not result in proper CLI error messaging.	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 checkspoof setting without a keyword returns an error when checkspoof 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
Polling the SMON Vlan Statistics MIB smonVlanidStatsCreateTime object returns an error. 7.91.01	7.91.01

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
SYSLOG Problems Corrected in 8.21.02.0001	Introduced in Version:
Messages that should be logged to the console as part of the shut down process are not seen.	7.80.01
client 22(Syslog) not ready in <number_of_tics> tics' message seen on console.	5.51.xx
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:
After repeated enable/disable of VRRP accept-mode, the system may not reply to ICMP echo requests to the VRRP Virtual IP address.	8.11.01
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
CWDM support: 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+
Additional DWDM support: 10GB-ER21-SFPP - 10GB-ER, DWDM CH21 SFP+ 10GB-ER24-SFPP - 10GB-ER, DWDM CH24 SFP+ 10GB-ER31-SFPP - 10GB-ER, DWDM CH31 SFP+ 10GB-ER33-SFPP - 10GB-ER, DWDM CH33 SFP+

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

Auto-config Problems Corrected in 8.11.05.0006	Introduced in Version:
On a chassis with 6 or more filled slots running with no/default configuration, if you do a "set configuration" command, during the reset you may see the following messages in the log: "<163>Sep 13 14:12:03 0.0.0.0 autoConfig[4.tDSrecv7]setConfigAtDefaultsBySlot: Unable to send nonvol change to msgQ inslot(6) value(0)" "<163>Sep 13 14:12:03 0.0.0.0 autoConfig[4.tIpAddrCb]autoConfig_IfEventCallback: Unable to send IF_DELETED-event(6), id(1) myid(0) to msgQ"	8.11.01

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
When a device goes through its synchronization process, it is possible for IGMP to cause an ISI exception, if internal structures get corrupted.	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
NAT Problems Corrected in 8.11.05.0006	Introduced in Version:
It is possible on a failover that a NAT Static Binding may be missing causing NAT translations to not function correctly.	8.11.01
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]generatelfIndex() :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 LLMNR 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, LLMNR, 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
NAT 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
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

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
VSB Problems Corrected in 8.11.05.0006	Introduced in Version:
S-180 class product syslog messages indicate that a VSB license is required after successfully bonding.	8.11.01

Feature Enhancements in 8.11.04.0005

Transceiver Enhancements in 8.11.04.0005
Support for the 40GB-ESR4-QSFP transceiver: 40Gb, Extended Reach SR4, MM, 300m OM3, MPO QSFP+
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:
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]ClgmpEts::DistGrpTblRecvDistributedAdd Recv base index out of range baseidx:xxx flowIdx:xxx	8.11.01

L3 VPN Problems Corrected in 8.11.04.0005	Introduced in Version:
After router failover, layer 3 VPN traffic may be transmitted with wrong label.	7.91.01
When configuring L3VPN on an access router the software license does not enable the feature. The user will not see any of the L3VPN commands.	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
Querying the ctAliasInterface table may not return all entries on a given interface in multislot systems.	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]InvFilePtrMgr::verify(3) calcCsum() failed. store=5, fileIdx=10.51, udpSum=0x77e366a, sumCount=65534	3.00.33
At boot time the following errors may be seen in the log: <163>Sep 19 14:46:02 0.0.0.0 NonVol[1.tusrAppInit]validate_files: Unknown record type;store=1,offset=4105,file=0.80, type=0,rawMaj=0,rawMin=0,rawLen=0 <163>Sep 19 14:46:02 0.0.0.0 NonVol[1.tusrAppInit]validate_files: file=1/0.80 rewinding over incomplete record. Truncating to size 4105 <163>Sep 19 14:46:02 0.0.0.0 NonVol[1.tusrAppInit]InvFilePtrMgr::fFlush(5) fflush(0x72b03b0) retval=-1, errno=9 Configuration could have been lost due to file corruption and should be verified.	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]InvFilePtrMgr::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:
Ambient air temperature is inaccurate for S1 chassis, and false warnings about hot ambient temperature are generated.	7.72.01
If chassis eeprom can not be accessed board will reset with no additional cause information displayed to cli or added to message log.	8.01.01
During initialization of a S180 SSA unit, a message similar to the following may be logged and the unit will reboot: bcmStrat[1.]pciMemRead: PcieCoreDeviceAccess::doMemRead() failed!	7.80.01

PLATFORM Problems Corrected in 8.11.04.0005	Introduced in Version:
Some devices may reset after logging a message similar to the one listed below because memory requires an adjustment to the 1.0V power controller. <163>Apr 7 15:05:51 0.0.0.0 Dune[5.tRootTask]PETRA[0] failed to initialize DRAM (0x65535).	8.01.01
Some devices may reset after logging a message similar to the one listed below because memory requires an adjustment to the 1.0V power controller. <163>Mar 27 03:06:57 192.168.100.18 Dune[2.dTcmTask]Petra[0] Received Interrupt PB_IPT_CRC_ERR_PKT instance 0, count 1, value= 0x1	8.01.01
System logs the message "bcmStrat[1.tNimIntr]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" and resets.	7.70.01
System logs the message "bcmStrat[1.tNimIntr]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000001, IP3_INTR_STATUS=0x00000000" and resets.	7.70.01
System logs the message "bcmStrat[2.tNimIntr]MEM_FAIL_INT_STAT=0x00040000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000" and resets.	7.70.01
System logs the message "bcmStrat[1.tNimIntr]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000002" 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

Transceiver Problems Corrected in 8.11.04.0005	Introduced in Version:
When plugging in a QSFP Model number 40GB-C0.5-QSFP copper cable into a 40g port an "fg.x.x unauthenticated pluggable module" message may display.	8.11.01

TWCB Problems Corrected in 8.11.04.0005	Introduced in Version:
When NAT hardware connections are reaped it is possible that subsequent NAT requests will not create a hardware connection.	5.01.58

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

Automated Deployment Feature Enhancements in 8.11.03.0005
Auto Configuration feature requests configuration information from DHCP server when chassis has no configuration. A SNMP trap requesting configuration is now sent to the SNMP server notifying it that the system is ready to be configured.

Problems Corrected in 8.11.03.0005

ACL Problems Corrected in 8.11.03.0005	Introduced in Version:
After updating to 8.11.01, any change made to the ACL configuration will cause any IPv4 and IPv6 ACL's applied inbound to not be applied after a reset.	8.11.01

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

ARP/ND Problems Corrected in 8.11.03.0005	Introduced in Version:
The chassis may crash when performing a distribution sync and when processing several ARP/ND related packets. A syslog produced during the crash will look similar to this: DistServ[1.tDsBrdOk]serverWatchDog.5(Host), client 92(net2Phys)	8.11.01
Autoconfig Problems Corrected in 8.11.03.0005	Introduced in Version:
The Automatic Deployment/Configuration feature will not start in S-chassis with IO modules even when running with default/cleared configuration.	8.11.01
BGP Problems Corrected in 8.11.03.0005	Introduced in Version:
Displaying FIB history via debug CLI may block BGP from maintaining connection to peers.	7.00.01
"Negating a BGP route-map ""match extended-community as-route-target"" command may result in a system reset. The following error message will appear at the CLI: SMS assert in qbmlrex3.c at line 414 : >= string_len 0 (2 * QB_LEN_EXT_COMMUNITY) 16 "	8.01.01
A system reset may occur when running BGP with the full Internet routing table and resetting or changing the export policy of a neighboring router. The following error message will appear: SMS assert in qbdcnhr.c at line 959 : (old_loc_route == ari_route->loc_route) 0 (QBRA_CHECK_FLAG(ari_route->loc_route->flags, QBRA_LOC_FLAG_REMOVAL_DONE)) 0	8.11.01
Multiprotocol BGP peering with third party products may not establish if received update messages contain out of order path attributes such that AS-PATH is the last attribute.	7.30.01
A system reset may occur if peering is attempted with a router supporting multisession BGP. The reset will occur on receipt of a Notification message with the error code of 2 (Open message error) and subcode 8 (grouping conflict). The following error message will appear: SMS assert in qbnmpd.c at line 141 : (null) INVALID BRANCH 0 (null) 0	8.11.01
Bonding Problems Corrected in 8.11.03.0005	Introduced in Version:
When inserting a module running 8.11.01.0001 into a Hardware VSB system, messages similar to the following will be stored the the message logs of the new module. <163>Apr 18 16:45:59 10.227.240.85 PPCtimer[6.tDispatch]PPC TBU has appeared to wrap during get_elapsed_time() <163>Apr 18 16:45:59 10.227.240.85 PPCtimer[6.tDispatch]1728088 17276bc c974ec 5d2314 5cdac8 155ea70	8.11.01
When inserting a module running 8.11.01.0001 into a Hardware VSB system, messages similar to the following will be stored the the message logs of the new module. Message 150/271 Syslog Message 08.11.01.0014 07/02/2013 07:52:56 <3>PiMgr[1.tDispatch]piMgrHwPortRxlcpu(131072,2,62,0,0x7e96e028,1044) RX ICPU message from own slot	8.11.01
Bonded chassis may segment after a slot reset.	7.70.00
Modules in a hardware bonded chassis may reset when a VSB port is connected to a front panel port. A message similar to "<0>Bond[13.tDispatch]getVsbInPort: learn inport:000033eb outport:00002bef binding failed (0x00c77d1c 0x00574058 0x015830e4 0x015756f4 0x0157ebec 0x01830ea0 0xeeeeeeee)" is logged on this error.	8.11.01

Bonding Problems Corrected in 8.11.03.0005	Introduced in Version:
SSA-T8028-0652 and SSA-G8018-0652 erroneously require a SSA-EOS-VSB license to enable chassis bonding.	8.11.01
VSB protocol may reset when enabling/disabling VSB ports.	7.62.02
IGMP flow may pick mismatched VSB ports causing loss of traffic across the Bond links.	7.60.01
DHCP Problems Corrected in 8.11.03.0005	Introduced in Version:
"dhcps6[{{slot#}}.tDSsync5]claimAllData: failed to set option(#) in vxWorks" syslog error message appear at start-up when dhcpv6 server pool is configured.	8.11.01
'ipv6 dhcp relay source-interface' disappears when the master blade is reset in a chassis.	7.30.01
DHCPv6 Problems Corrected in 8.11.03.0005	Introduced in Version:
DHCPv6 server responds to DHCPv6 request on interfaces that do not have 'ipv6 dhcp server' configured.	8.11.01
FDB Problems Corrected in 8.11.03.0005	Introduced in Version:
If the amount of MAC addresses is configured to be 128K, static Unicast and Multicast MAC entries may not function correctly. When attempting to create the entries, messages similar to: FDB: NonVol[2.tDSrecv3]writeData MAJOR_FDB_STATIC_ENTRIES minorTag=66651, may be logged.	7.91.01
When changing the number of MAC addresses supported to between 64K and 128K, a chassis reboot is needed for new value to take effect. If, between the time of the configuration change, and the chassis reboot, a blade resets, it will go into an infinite reboot cycle and display a message similar to: <3>FilterDb[6.tDSrecv3]Resetting for new fdb num entries = 65536, old number entries = 131072	7.91.01
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:
Traceroute using UDP does not work for layer 3 VPNs over tunnels.	8.01.01
After issuing the traceroute command, the string "runTraceroute: ifindex <number>" is displayed before the results.	7.99.00
IPv4 Forwarding Problems Corrected in 8.11.03.0005	Introduced in Version:
It is possible that reframer resources could become disabled while still in use for some tunneled and IPv6Nat flows. The flows associated with these disabled resources would be dropped until it aged out of hardware.	8.11.01

IPv4 Forwarding Problems Corrected in 8.11.03.0005	Introduced in Version:
On router failover, layer 3 VPN filter connections may not be removed if label to VRF mappings change.	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
MAC Authentication Problems Corrected in 8.11.03.0005	Introduced in Version:
MAC-Authentication auth-mode may be set to radius-username when upgrading from older firmware versions.	8.11.01
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
In multiauth sessions-unique-per-port enabled mode, antispoof IP bindings may not be updated for a MAC address with sessions on multiple ports.	8.11.01
NAT Problems Corrected in 8.11.03.0005	Introduced in Version:
It is possible for a NAT Static reserved binding to age out.	8.11.01
If a large number of binding are created with the same global address it is possible for the board to reset when deleting bindings.	7.91.03
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
Node Alias Problems Corrected in 8.11.03.0005	Introduced in Version:
Node Alias is unable to decode packet information for LLMNR and mDNS packets after compression occurs.	8.11.01
In node alias, the protocol setting for LLMNR, SSDP, and mDNS are not displayed in the configuration.	8.11.01
OSPF Problems Corrected in 8.11.03.0005	Introduced in Version:
If OSPFv2 and OSPFv3 are both configured to use the same tracked object on a single interface, and then one of these is removed, a misleading message indicates that the track is in use and will not be deregistered. The track is only removed for the corresponding address-family and continues to be in-use for the other address-family.	8.11.01

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
PIM-SM Problems Corrected in 8.11.03.0005	Introduced in Version:
The "rtr mcast show debug fe" counters within Show Support always display counts of 0.	8.11.01
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
Some devices may reset after logging a message similar to the one listed below: <163>Mar 27 03:06:57 192.168.100.18 Dune[2.dTcmTask]Petra[0] Received Interrupt PB_IPT_CRC_ERR_PKT instance 0, count 1, value= 0x1	8.01.01
Some devices may reset after logging a message similar to the one listed below because memory requires an improved initialization sequence.<163>Apr 7 15:05:51 0.0.0.0 Dune[5.tRootTask]PETRA[0] failed to initialize DRAM (0x65535).	8.01.01
S180 and S140 blades may not automatically restart when the chassis AC power supplies are overloaded. This can occur during an AC power outage when some but not all required AC power supplies lose AC power. Blades must be ejected/reinserted or the chassis must be fully powered down then up to recover from the condition.	8.11.01
S chassis reporting an incorrect ambient temperature of -3C.	7.60.01
40Gb QSFP+ ports that have a QSFP+ 40Gb to 4x10Gb fanout cable inserted do not always come up in the correct 4x10Gb mode which is displayed in "show port status" after a board reset.	8.11.01
Sometimes SFP or SFP+ modules may be missidentified for both type and speed. This can result in the port being non functional when speed is wrong or prone to CRC or Link problems when type is wrong. Miss identification can occur at the time SFP(+) is inserted or during a subsequent boot of the blade. Four port SFP+ option modules, 8 and 16 port SFP+ modules are not affected.	8.11.01
Traffic in both directions may not be established on a 10Gb capable port, with a 10Gb SFP+ installed, on a chassis module or standalone after a 1Gb SFP had been inserted into such port.	8.11.01

Platform Problems Corrected in 8.11.03.0005	Introduced in Version:
A S140/S180 blade may display messages similar to the following when backplane fabric is oversubscribed. <165>Jun 25 10:44:31 10.1.147.12 Dune[3.dTcmTask]Petra[1] Received Interrupt PB_IPS_CREDIT_OVERFLOW instance 0, count 162, value= 0x146b <165>Jun 25 10:44:36 10.1.147.12 Dune[3.tDuneErrM]Petra[1] Interrupt PB_IPS_CREDIT_OVERFLOW instance 0 still active <165>Jun 25 10:44:56 10.1.147.12 Dune[3.tDuneErrM]Petra[1] Interrupt PB_IPS_CREDIT_OVERFLOW instance 0 is off	8.11.01
Transceivers inserted into corresponding ports on each bank of ports (ex. port zero on each bank would be ports 1,9,17) might result in incorrect transceiver detection and functionality.	8.11.01
During module initialization a message may be logged similar to: "i2c[4.tusrApplnit]writeBatchCommand: master 4 empty interrupt timeouts".	8.11.01
Querying the entPhysicalAssetID object for a module that has not yet been programmed might return unexpected string.	8.11.01
A module will sometimes report a message similar to "<163>Jul 15 15:52:54 0.0.0.0 System[1]Module moved from chassis: 20b399559169 to chassis: 20b399559dfd" even when it has not moved.	7.60.01
Routing Problems Corrected in 8.11.03.0005	Introduced in Version:
Layer 3 VPN filter connections created on router failover are not removed when new labels are sent to forwarding plane.	7.91.01
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
SNMP Problems Corrected in 8.11.03.0005	Introduced in Version:
S-Series SK8009-1224-F8 and SK8008-1224-F8 boards have incorrect ENTITY-MIB physical description strings.	8.11.01
SSH Problems Corrected in 8.11.03.0005	Introduced in Version:
If a user's account is configured for local-only authentication, and the account is disabled (administratively or due to excessive login failures), and the user tries to connect (even just once) using SSH with public key authentication, then a port lock out will occur (regardless of the configured number of system lockout attempts).	8.11.01
Tunneling Problems Corrected in 8.11.03.0005	Introduced in Version:
The switch may stop forwarding if an L2 encapped IPv6 in IPv6 GRE packet arrives from a tunnel dedicated to a pseudowire.	8.11.01
Tunnel probes are not restored properly on S-Series modules.	8.11.01

VRRP Problems Corrected in 8.11.03.0005	Introduced in Version:
"RtrVRRP[{MODULE}.tVrrpEvt]Failed: unable to update userData flags for IP {IP ADDRESS} for {INTERFACE}" syslog message is logged from an initializing module.	8.11.01
Checkspoof strict-mode enabled on host-mobility interface would be triggered by host transmitting packets into the router if router had learned about host via OSPF from VRRP host-mobility partner.	8.11.01

Problems Corrected in 8.11.02.0001

Upgrade Problems Corrected in 8.11.02.0001	Introduced in Version:
After updating to 8.11.01, inbound ACLs (IPv4 and IPv6) are no longer functional. This occurs after a reboot when changes have been made to the ACL configuration.	8.11.01

Feature Enhancements in 8.11.01.0014

HW Feature Enhancements in 8.11.01.0014
<p>Support for the S180 Class I/O and I/O Fabric modules including:</p> <p>SL8013-1206-F8 S-Series S180 Class I/O-Fabric Module, Load Sharing - 6 Ports 40GBASE-X Ethernet via QSFP+, 4 ports VSB via SFP+ (Used in S1A/S4/S6/S8)</p> <p>SK8008-1224-F8 S-Series S180 Class I/O-Fabric Module, Load Sharing - 24 Ports 10GBASE-X via SFP+, 4 ports VSB via SFP+ (Used in S1A/S4/S6/S8)</p> <p>SK8009-1224-F8 S-Series S180 Class I/O-Fabric Module, Load Sharing - 24 Ports 10GBASE-T via RJ45, 4 ports VSB via SFP+ (Used in S1A/S4/S6/S8)</p> <p>ST8206-0848-F8 S-Series S180 Class I/O-Fabric Module, Load Sharing - 48 Ports 10/100/1000BASE-T via J45 with PoE (802.3at) and two Type2 option slots (Used in S1A/S4/S6/S8)</p> <p>SG8201-0848-F8 S-Series S180 Class I/O-Fabric Module, Load Sharing - 48 Ports 1000BASE-X via SFP and two Type2 options slots (Used in S1A/S4/S6/S8)</p> <p>SL8013-1206 S-Series S180 Class I/O Module - 6 Ports 40GBASE-X Ethernet via QSFP+, VSB expansion slot (Used in S4/S6/S8)</p> <p>SK8008-1224 S-Series S180 Class I/O Module -24 Ports 10GBASE-X via SFP+, VSB expansion slot (Used in S4/S6/S8)</p> <p>SK8009-1224 S-Series S180 Class I/O Module -24 Ports 10GBASE-T via RJ45, VSB expansion slot (Used in S4/S6/S8)</p>
<p>Support for HW VSB models including:</p> <p>SOV3208-0202 S-Series Option Module (Type2)- 2 port VSB Option Module (Compatible with Type2 option slots on S140/S180 modules only)</p> <p>SOV3008-0404 S-Series VSB Expansion Module - 4 port VSB Module (Compatible with S180 Class 10Gb/40Gb I/O modules only)</p>
S130/S150/S155, SSA130/SSA150 classes must use this image when modules are mixed, physically (in the same chassis) or logically (using VSB) with the S180/S140 or SSA180/SSA150A.

Application Policy Feature Enhancement in 8.11.01.0014
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 8.11.01.0014

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 8.11.01.0014

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 8.11.01.0014

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.

Inter-VRF Access Control List Feature Enhancement in 8.11.01.0014

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 8.11.01.0014

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 8.11.01.0014

SSH enhancement to support Public Key Authentication as an additional client authentication method.

RMON Stats and History Feature Enhancement in 8.11.01.0014

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 8.11.01.0014

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.

MAC Authentication Feature Enhancement in 8.11.01.0014

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

IPv6 DHCP Server Feature Enhancement in 8.11.01.0014

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 8.11.01.0014

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 8.11.01.0014

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

OSPF RFC 4577 Support Feature Enhancement in 8.11.01.0014

Enhancement to allow OSPF to be used as the routing protocol between provider edge and customer edge devices when deployed in a BGP/MPLS L3VPN environment.

Neighbor Discovery Enhancement Feature Enhancement in 8.11.01.0014

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

Feature Enhancements in 8.02.01.0012**HW Feature Enhancements in 8.02.01.0012**

This image supports the hybrid TripleSpeed PoE/SFP+ option module part number; SOTK2268-0212, S-Series Option Module (Type2) - 10 Ports 10/100/1000BASE-T via RJ45 with PoE and 2 ports 10GBASE-X via SFP+ (Compatible with Type2 option slots)

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

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.

L3VPN over GRE Feature Enhancements in 8.02.01.0012

This feature adds support for creating L3VPNs transparently over an IP core network using GRE or IP tunnels. With this feature core network routers do not need to be VRF aware or carry knowledge of the specific routes.

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

Sampled 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.

Network Address Translation Feature Enhancements in 8.02.01.0012

NAT Cone with hair pinning support – Enhancement to existing NAT functionality to allow connections to be initiated from external devices once the internal device has primed the NAT engine with an internal/external binding. With hair pinning, multiple devices on the internal network will not be routed externally regardless of the fact they may only have knowledge of external IP addresses. When NAT is in use, traffic like XBOX live requires the use of this feature.

Network Address Translation – Feature enhancement to support network address translation (NAT) for IPv6 to IPv6 addresses.

Load Sharing NAT – Feature enhancement to support load sharing network address translation (LSNAT) for IPv4 to IPv6, IPv6 to IPv4 as well as IPv6 to IPv6 addresses.

Transparent Web Cache Balancing (TWCB) – Feature enhancement to support Transparent Web Cache Balancing for IPv6 clients to IPv6 destination addresses.

Proxy-Web – This feature is an enhancement to TWCB that leverages NAT functionality so that web cache servers do not need to be local to the router performing TWCB. Web cache servers can be distributed throughout the network if desired. This feature enhancement is applicable to both IP4 and IPv6 implementations of TWCB. In addition the feature allows for a proxy environment without the need to configure user end stations.

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 BGP, 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 50,000 routes for the S155 module class.

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.)

LSNAT Feature Enhancements in 8.02.01.0012

'show running slb' now displays additional information.

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></p> <p>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
BGP Problems Corrected in 8.02.01.0012	Introduced in Version:
System may log a "BGP SMS assert in qbmlpar3.c" message and reset.	7.00.01
Config Problems Corrected in 8.02.01.0012	Introduced in Version:
Configs not cleared when moving modules to new chassis in the same slots.	7.60.01
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
NAT Problems Corrected in 8.02.01.0012	Introduced in Version:
An "ICMP Port Unreachable" message being NATted to an overloaded List rule will no longer generate a log "Failed to allocate ip address (Global IP addresses exhausted for pool) reported x times" but will be silently discarded.	6.12.08
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
When gracefully restarting a Designated Router, OSPF may not send hellos with itself as the DR.	8.01.01
A blade may reset repeatedly logging a DSI exception for thread tDSync5.	8.01.01
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
SSA may report multiple fan insert/removal messages when a single insert or removal occurs.	UNTARGETED
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
SNMP Problems Corrected in 8.02.01.0012	Introduced in Version:
S-Series returns no interface speed value for vtap interface.	1.07.19

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:

It is not possible to mix S130/S150/S155 fabrics and the S180 fabric class in the same chassis.
S140 and S180 class modules require the use of S180 class fabrics when used in the S4/S6 and S8 chassis. S150/S130 class I/O can be used with any fabrics class.
MPLS/LPD/L3VPNs will not function over an IPv6 core. This will be added in a later release.
When upgrading to 8.11.05 it is possible that some IPv6 interface configuration will be lost. This has been observed in bonded systems when doing a HAU upgrade.
<p>The S140 modules are shipped with factory only version firm ware 7.99.06. As shipped, this module is compatible with S3 Chassis systems running 7.99.06 or newer (factory only firmware) or 8.11.01 or newer (customer firmware). When installing this module in an S3 chassis, as shipped this module is not compatible with an S3 chassis running customer firmware version 8.01 or 8.02. This module will run on an 8.02 system once the module is upgraded to 8.02.</p> <p>If the chassis this module is installed in is running firmware 8.02 or less, use the following instructions to upgrade the module firmware:</p> <p>Installing in an S3 Chassis Currently Running 8.01 Firmware</p> <p>If you are installing this Module in an S3 chassis that is currently running 8.01 Firmware, the chassis firmware must be upgraded:</p> <ol style="list-style-type: none"> 1. Load and boot the desired firmware on existing modules in the chassis 2. If you wish to operate the chassis with FW version 8.02, you must follow the instructions in section "Installing in an S3 Chassis Currently Running 8.02 Firmware " for install in chassis running 8.02 firmware <p>Installing in an S3 Chassis Currently Running 8.02 Firmware</p> <p>If you are installing this Module in an S3 Chassis that is currently running 8.02 firmware, you must:</p>

<ol style="list-style-type: none"> 1. Install this module. After an extended boot time the module will remain isolated from other modules in the chassis, but becomes operational. 2. Attach a console cable to the chassis' com port associated with this module's slot. 3. Log in using username: Admin and password <Enter> (null password). 4. Use a USB storage device inserted in the chassis' USB port associated with this module to copy the desired 8.02 firmware onto the module. 5. Set the boot firmware version using the set boot system 8.02-firmware-name command. 6. Reset the module using the reset slot-number command
The S180 class fabrics require the S1-Chassis-A and are not supported in the S1-Chassis. The S1-Chassis-A supports all of the S-Series fabrics module classes.
Adjacent 40Gb QSFP+ ports must operate in the same mode. Upon release, adjacent ports (1/2, 3/4, 5/6) must run in the same mode, 4x10Gb or 40Gb. This restriction will be removed in a future release.
Only Extreme sourced 40 Gigabit optical transceivers are supported. Use of any other transceiver types will result in an error.
The 10GB-LRM-SFP transceiver is not supported when plugged into a QSFP+ port via a QSFP-SFP-ADPT.
MGBIC-100BT doesn't support automatic detection of MDIX (Medium Dependent Interface Crossover) or Auto-negotiation.
Only Series 2 option modules may be used with the S140/S180 Class modules. These include model numbers: SOK2208-0102, SOK2208-0104, SOK2208-0204, SOG2201-0112, SOT2206-0112, SOGK2218-0212, SOTK2268-0212, SOK2209-0204.
The VSB HW expansion module; SOV3008-0404, S-Series VSB Expansion Module - 4 port VSB Module can only be used on the S180 I/O modules, SL8013-1206, SK8008-1224, SK8009-1224
Mixing S140 class and S130 class in the same S3 chassis is not supported. The S3 chassis must be populated with only S140 or S130 classes.
The following interface configuration command introduced in 8.01.01, ip ospf <pid> area <x.x.x.x> can cause a DSI and reset. Continue to use the <i>network</i> command under OSPF configuration mode. The <i>network</i> command is the preferred and in previous releases, the only way to enable OSPF on an interface.
When using VSB the number of configured bonding ports should be limited to no more than 16 on each physical chassis. Exceeding this limit may result in delays processing bond port link events.
When using HW VSB, the IDS mirror feature is not supported.
When using SW VSB several features are resized or restricted: <ul style="list-style-type: none"> LAG capacities are reduce to 189 for chassis, 61 for SSAs, IP tunnels including VXLAN and GRE Tunnels are not supported, (Remote port mirrors are supported) Port Mirroring support for 5 mirrors, - 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.
Systems with the NAT/LSNAT/etc family of features enabled should not populate slot 16 in a VSB chassis.
The S1-Chassis and S1-Chassis-A requires the SSA-AC-PS-1000W power supplies. (The SSA-AC-PS-625W must not be used in the S1-Chassis.) The Fabrics/Option Modules and optics along with the Fans can exceed the power available in the 625W supply during the startup and when the fans operate at full speed.
The "script" command should not be used. Its use will result in memory corruption and reset or other undesired behavior.
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. On SSA all 10G ports are in the same group. All ports on a 10G

Option Module are grouped together. For S blades shipping with factory configured ports the groups are: tg.x.1-4, tg.x.5-8, tg.x.9-12, tg.x.13-16.

The S130 Class of blades supports Jumbo Frames on a maximum of 12 ports simultaneously. These ports can be any combination of the fixed 48 ports found on the module.

Route-map (PBR) counters may not display correctly, causing them to appear as though the counts are not changing.

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

IEFT STANDRDS 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
RFC1265	BGP Protocol Analysis
RFC1266	Experience with the BGP Protocol

RFC No.	Title
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
RFC1518	CIDR Architecture
RFC1519	Classless Inter-Domain Routing (CIDR)
RFC1542	BootP: Clarifications and Extensions
RFC1624	IP Checksum via Incremental Update
RFC1657	Managed Objects for BGP-4 using SMIv2
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
RFC1771	A Border Gateway Protocol 4 (BGP-4)
RFC1772	Application of BGP in the Internet
RFC1773	Experience with the BGP-4 protocol
RFC1774	BGP-4 Protocol Analysis
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
RFC1930	Guidelines for creation, selection, and registration of an Autonomous System (AS)
RFC1966	BGP Route Reflection
RFC1981	Path MTU Discovery for IPv6
RFC1997	BGP Communities Attribute
RFC1998	BGP Community Attribute in Multi-home Routing
RFC2001	TCP Slow Start
RFC2003	IP in IP Tunneling
RFC2012	TCP-MIB
RFC2013	UDP-MIB
RFC2018	TCP Selective Acknowledgment Options
RFC2030	SNTP

RFC No.	Title
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
RFC2260	Support for Multi-homed Multi-prov
RFC2270	Dedicated AS for Sites Homed to one Provider
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
RFC2385	BGP TCP MD5 Signature Option
RFC2391	LSNAT
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)
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 Autoconfiguration
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

RFC No.	Title
RFC2519	A Framework for Inter-Domain Route Aggregation
RFC2545	BGP Multiprotocol Extensions for IPv6
RFC2547	BGP/MPLS VPNs
RFC2548	Microsoft Vendor-specific RADIUS Attributes
RFC2553	BasicSocket 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
RFC2663	NAT & PAT (NAPT)
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
RFC2796	BGP Route Reflection
RFC2819	RMON MIB
RFC2827	Network Ingress Filtering
RFC2858	Multiprotocol Extensions for BGP-4
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
RFC2918	Route Refresh Capability for BGP-4
RFC2922	PTOPO-MIB
RFC2934	PIM MIB for IPv4
RFC2966	Prefix Distribution with Two-Level IS-IS
RFC2973	IS-IS Mesh Groups

RFC No.	Title
RFC2991	Multipath Issues in Ucast & Mcast Next-Hop
RFC3022	Traditional NAT
RFC3056	Connection of IPv6 Domains via IPv4 Clouds
RFC3065	Autonomous System Confederations for BGP
RFC3069	VLAN Aggregation for Efficient IP Address Allocation
RFC3101	The OSPF Not-So-Stubby Area (NSSA) Option
RFC3107	Carrying Label Information in BGP-4
RFC3137	OSPF Stub Router Advertisement
RFC3162	RADIUS and IPv6
RFC3273	HC-RMON-MIB
RFC3291	INET-ADDRESS-MIB
RFC3315	DHCPv6
RFC3031	Multiprotocol Label Switching Architecture
RFC3032	MPLS Label Stack Encoding
RFC3345	BGP Persistent Route Oscillation
RFC3359	TLV Codepoints in IS-IS
RFC3373	Three-Way Handshake for IS-IS
RFC3376	Internet Group Management Protocol, Version 3
RFC3392	Capabilities Advertisement with BGP-4
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
RFC3562	Key Mgt Considerations for TCP MD5 Signature Opt
RFC3567	IS-IS Cryptographic Authentication
RFC3579	RADIUS Support For Extensible Authentication Protocol (EAP)

RFC No.	Title
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
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)
RFC4023	Encapsulation of MPLS in IP or GRE
RFC4026	Provider Provisioned VPN Terminology
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
RFC 4250	The Secure Shell (SSH) Protocol Assigned Numbers
RFC 4251	The Secure Shell (SSH) Protocol Architecture
RFC 4252	The Secure Shell (SSH) Authentication Protocol
RFC 4253	The Secure Shell (SSH) Transport Layer Protocol (no support diffie-hellman-group14-sha1)

RFC No.	Title
RFC 4254	The Secure Shell (SSH) Connection Protocol
RFC 4256	Generic Message Exchange Authentication for the Secure Shell Protocol (SSH)
RFC4264	BGP Wedgies
RFC4265	Definition of Textual Conventions for VPN Mgt
RFC4268	ENTITY-STATE-MIB
RFC4268	ENTITY-STATE-TC-MIB
RFC4271	A Border Gateway Protocol 4 (BGP-4)
RFC4272	BGP Security Vulnerabilities Analysis
RFC4273	Managed Objects for BGP-4 using SMIv2
RFC4274	BGP-4 Protocol Analysis
RFC4275	BGP-4 MIB Implementation Survey
RFC4276	BGP-4 Implementation Report
RFC4277	Experience with the BGP-4 protocol
RFC4291	IP Version 6 Addressing Architecture
RFC4292	IP Forwarding MIB
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
RFC4360	BGP Extended Communities Attribute
RFC4364	BGP/MPLS IP Virtual Private Networks (VPNs)
RFC4365	Applicability Statement for BGP/MPLS IP VPNs
RFC4382	MPLS/BGP L3VPN MIB
RFC4384	BGP Communities for Data Collection
RFC 4419	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
RFC4451	BGP MULTI_EXIT_DISC (MED) Considerations
RFC4456	BGP Route Reflection
RFC4486	Subcodes for BGP Cease Notification Message
RFC4541	IGMP Snooping
RFC4541	MLD Snooping

RFC No.	Title
RFC4552	Authentication/Confidentiality for OSPFv3
RFC4560	DISMAN-PING-MIB
RFC4560	DISMAN-TRACEROUTE-MIB
RFC4560	DISMAN-NSLOOKUP-MIB
RFC4577	OSPF as PE/CE Protocol for BGP L3 VPNs
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)
RFC4659	BGP-MPLS IP VPN Extension for IPv6 VPN
RFC4668	RADIUS Client MIB
RFC4670	RADIUS Accounting MIB
RFC4673	RADIUS Dynamic Authorization Server MIB
RFC 4716	The Secure Shell (SSH) Public Key File Format
RFC4724	Graceful Restart Mechanism for BGP
RFC4750	OSPFv2 MIB
RFC4760	Multiprotocol Extensions for BGP-4
RFC4835	CryptoAlgorithm Requirements for ESP and AH
RFC4836	MAU-MIB
RFC4836	IANA-MAU-MIB
RFC4861	Neighbor Discovery for IPv6
RFC4862	IPv6 Stateless Address Autoconfiguration
RFC4878	OAM Functions on Ethernet-Like Interfaces
RFC4878	DOT3-OAM-MIB
RFC4884	RFC 4884 Extended ICMP Multi-Part Messages
RFC4893	BGP Support for Four-octet AS Number Space
RFC4940	IANA Considerations for OSPF
RFC5036	LDP Specification
RFC5059	Bootstrap Router (BSR) Mechanism for (PIM)
RFC5060	PIM MIB
RFC5065	Autonomous System Confederations for BGP
RFC5095	Deprecation of Type 0 Routing Headers in IPv6
RFC5132	IP Multicast MIB
RFC5176	Dynamic Authorization Extension to RADIUS

RFC No.	Title
RFC5186	IGMPv3/MLDv2/MCAST Routing Protocol Interaction
RFC5187	OSPFv3 Graceful Restart
RFC5240	PIM Bootstrap Router MIB
RFC5250	The OSPF Opaque LSA Option
RFC5291	Outbound Route Filtering Capability for BGP-4
RFC5292	Address-Prefix-Outbound Route Filter for BGP-4
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
RFC5396	Textual Representation AS Numbers
RFC5398	AS Number Reservation for Documentation Use
RFC5492	Capabilities Advertisement with BGP-4
RFC5519	MGMD-STD-MIB
RFC5601	Pseudowire (PW) MIB
RFC5602	Pseudowire (PW) over MPLS PSN MIB
RFC5643	OSPFv3 MIB
RFC5798	Virtual Router Redundancy Protocol (VRRP) V3
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
RFC6296	IPv6-to-IPv6 Network Prefix Translation
RFC6329	IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging
RFC6549	OSPFv2 Multi-Instance Extensions
RFC6565	OSPFv3 as PE/CE Protocol for BGP L3 VPNs
RFC7348	Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks
Drafts	draft-ietf-idr-bgp4-mibv2 (Partial Support)
Drafts	draft-ietf-idr-bgp-identifier
Drafts	draft-ietf-idr-as-pathlimit
Drafts	draft-ietf-idr-mrai-dep (Partial Support)
Drafts	draft-ietf-isis-experimental-tlv (Partial Support)

RFC No.	Title
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

EXTREME 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-VLAN-INTERFACE-MIB
CTRON-ENVIROMENTAL-MIB	ENTERASYS-MAU-MIB-EXT-MIB	IANA-ADDRESS-FAMILY-NUMBERS-MIB
CTRON-MIB-NAMES	ENTERASYS-MGMT-AUTH-NOTIFICATION-MIB	IEEE8021-PAE-MIB
CTRON-OIDS	ENTERASYS-MGMT-MIB	IEEE8023-LAG-MIB
DVMRP-MIB	ENTERASYS-MIB-NAMES DEFINITIONS	IEEE8021-BRIDGE-MIB
CTRON-Q-BRIDGE-MIB-EXT	ENTERASYS-MIRROR-CONFIG	IEEE8021-CFM-MIB
CISCO-CDP-MIB	ENTERASYS-MSTP-MIB	IEEE8021-CFM-V2-MIB
CISCO-NETFLOW-MIB	ENTERASYS-MULTI-AUTH-MIB	IEEE8021-MSTP-MIB
CISCO-TC	ENTERASYS-MULTI-TOPOLOGY-ROUTING-MIB	IEEE8021-Q-BRIDGE-MIB
ENTERASYS-FLOW-LIMITING-MIB	ENTERASYS-MULTI-USER-8021X-MIB	IEEE8021-SECY-MIB
ENTERASYS-AAA-POLICY-MIB	ENTERASYS-NETFLOW-MIB (v5 & v9)	IEEE8021-SPANNING-TREE-MIB
ENTERASYS-CLASS-OF-SERVICE-MIB	ENTERASYS-OIDS-MIB DEFINITIONS	IEEE8023-DOT3-LLDP-EXT-V2-MIB
ENTERASYS-CONFIGURATION-MANAGEMENT-MIB	ENTERASYS-OSPF-EXT-MIB	LLDP-MIB
ENTERASYS-CONVERGENCE-END-POINT-MIB	ENTERASYS-PFC-MIB-EXT-MIB	LLDP-EXT-MED-MIB
ENTERASYS-DIAGNOSTIC-MESSAGE-MIB	ENTERASYS-PIM-EXT-MIB	LLDP-EXT-DOT1-MIB

Title	Title	Title
ENTERASYS-DNS-RESOLVER-MIB	ENTERASYS-POLICY-PROFILE-MIB	LLDP-EXT-DOT3-MIB
ENTERASYS-DVMRP-EXT-MIB	ENTERASYS-POWER-ETHERNET-EXT-MIB	LLDP-EXT-DOT3-V2-MIB
ENTERASYS-ETH-OAM-EXT-MIB	ENTERASYS-PTOPO-MIB-EXT-MIB	LLDP-EXT-DOT3-V2-MIB (IEEE 802.3-2009) ETS Admin table read only
ENTERASYS-IEEE8021-BRIDGE-MIB-EXT-MIB	ENTERASYS-PWA-MIB	RSTP-MIB
ENTERASYS-IEEE8021-SPANNING-TREE-MIB-EXT-MIB	ENTERASYS-RESOURCE-UTILIZATION-MIB	U-BRIDGE-MIB
ENTERASYS-IEEE8023-LAG-MIB-EXT-MIB	ENTERASYS-RIPv2-EXT-MIB	USM-TARGET-TAG-MIB
ENTERASYS-IETF-BRIDGE-MIB-EXT-MIB	ENTERASYS-RMON-EXT-MIB	ENTERASYS-TWCB-MIB
ENTERASYS-IETF-P-BRIDGE-MIB-EXT-MIB	VSB-SHARED-SECRET-MIB	ENTERASYS-NAT-MIB
ENTERASYS-IF-MIB-EXT-MIB	ENTERASYS-SNTP-CLIENT-MIB	ENTERASYS-LSNAT-MIB
ENTERASYS-IP-SLA-MIB	ENTERASYS-RADIUS-ACCT-CLIENT-EXT-MIB	ENTERASYS-VRRP-EXT-MIB DEFINITIONS
ENTERASYS-ENTITY-SENSOR-MIB-EXT-MIB	ENTERASYS-RADIUS-AUTH-CLIENT-MIB	SNMP-RESEARCH-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	ospflfStateChange ospfVirtIfStateChange ospfNbrStateChange ospfVirtNbrStateChange ospflfConfigError ospfVirtIfConfigError ospfMaxAgeLsa ospfOriginateLsa
RFC 1907	Cold Start Warm Start Authentication Failure
RFC 4133	entConfigChange
RFC 2668	ifMauJabberTrap
RFC 2819	risingAlarm fallingAlarm

RFC No.	Title
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

RFC No.	Title
Enterasys-spanning-tree-diagnostic-mib	etsysMstpLoopProtectEvent etsysStpDiagCistDisputedBpduThresholdExceeded etsysStpDiagMstiDisputedBpduThresholdExceeded
Lldp-mib	lldpNotificationPrefix (IEEE Std 802.1AB-2004)
Lldp-ext-med-mib	lldpXMedTopologyChangeDetected (ANSI/TIA-1057)
Enterasys-class-of-service-mib	etsysCosIrlExceededNotification
Enterasys-policy-profile-mib	etsysPolicyRulePortHitNotification
Enterasys-mstp-mib	etsysMstpLoopProtectEvent
Ctron-environment-mib	chEnvAmbientTemp chEnvAmbientStatus

RADIUS ATTRIBUTE SUPPORT:

This section describes the support of RADIUS attributes on the S-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.