

Customer Release Notes

S-Series® and S-Series® Standalone

Firmware Version 8.12.04.0003

July 2016

INTRODUCTION:

This document provides specific information for version 08.12.04.0003 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			
SL8013-1206-F8	SK8008-1224-F8	SK8009-1224-F8	ST8206-0848-F8
SG8201-0848-F8	SL8013-1206	SK8008-1224	SK8009-1224

S140 Class Modules			
ST2206-0848	SG2201-0848	SK2008-0832	SK2009-0824

S155 Class Modules			
SK5208-0808-F6	ST5206-0848-F6	SG5201-0848-F6	

S150 Class Modules			
SK1208-0808-F6	ST1206-0848-F6	SG1201-0848-F6	SK1008-0816
ST1206-0848	SG1201-0848		

S130 Class Modules			
ST4106-0248	SG4101-0248	ST4106-0348-F6	

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

SSA Models			
SSA-T1068-0652A	SSA-T8028-0652	SSA-G8018-0652	SSA-G1018-0652
SSA-T1068-0652	SSA-T4068-0252		

S-Series and S-Series Standalone Customer Release Notes

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

For the latest firmware versions, visit the download site at:
www.extremenetworks.com/support/enterasys-support/

PRODUCT FIRMWARE SUPPORT:

Status	Firmware Version	Product Type	Release Date
Current Version	8.12.04.0003	Customer Release	July 2014
Previous Version	8.12.03.0001	Customer Release	May 2014
Previous Version	8.12.02.0006	Customer Release	April 2014
Previous Version	8.12.01.0011	Customer Release	January 2014
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

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).

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.

HIGH AVAILABILITY UPGRADE (HAU) FW COMPATIBILITY:

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

943d080eb98cfae754ac2b6b9f26aee9c26bf2d1
(The HUA key is reported using the CLI command 'dir images').

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.

S-Series and S-Series Standalone Customer Release Notes

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:

S180 Class		S155 Class		S140 Class	
ST8206-0848-F8	08.12.02.0006	SK5208-0808-F6	07.21.02.0002	ST2206-0848	08.01.01.0016
SG8201-0848-F8		ST5206-0848-F6		SG2201-0848	
SK8009-1224-F8		SG5201-0848-F6		SK2008-0832	
SL8013-1206-F8	08.11.01.0014	S150 Class		SK2009-0824	08.02.01.0012
SK8008-1224-F8		SK1208-0808-F6	S130 Class		
SL8013-1206		ST1206-0848-F6	ST4106-0348-F6	07.02.02.0002	
SK8008-1224		SG1201-0848-F6	ST4106-0248		
SK8009-1224		SK1008-0816	SG4101-0248		
			ST1206-0848		
		SG1201-0848			

Option Modules			
Series 1 (Compatible with S130/S150/S155 only)		Series 2 (Compatible with all classes)	
SOK1208-0102	07.01.01.000X	SOK2208-0102	07.72.01.0021
SOK1208-0104		SOK2208-0104	
SOK1208-0204		SOK2208-0204	
SOG1201-0112		SOG2201-0112	
SOT1206-0112		SOT2206-0112	
		SOGK2218-0212	
		SOTK2268-0212	08.02.01.0012
		SOV3208-0202	8.11.01.0014

Expansion Module	
SOV3008-0404	8.11.01.0014

S-Series and S-Series Standalone Customer Release Notes

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	07.72.01.0021
S3-Chassis-A	
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

Matrix S Standalone Series (SSA) Modules Minimum FW Required:

SSA 180 Class	
SSA-T8028-0652	08.01.01.0016
SSA-G8018-0652	
SSA 150 Class	
SSA-T1068-0652A	08.01.01.0016
SSA-T1068-0652	07.01.01.000X
SSA-G1018-0652	
SSA 130 Class	
SSA-T4068-0252	07.01.01.000X

Matrix S Power Supplies Series:

S-AC-PS	07.01.01.000X
S-AC-PS-15A	07.42.02.0002

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.

S-Series and S-Series Standalone Customer Release Notes

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) (S140 Class)
S3-Chassis-POE4	
S3-Chassis-A	
S3-Chassis-POEA	
S1-Chassis	2,048 (2K)
S1-Chassis-A	

S180/150/S155 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

S-Series and S-Series Standalone Customer Release Notes

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:

SOK1208-0102/SOK2208-0102,
SOK1208-0104/SOK2208-0104,
SOK1208-0204/ SOK2208-0204,
SOG1201-0112/SOG2201-0112,
SOGK2218-0212 and SOTK2268-0212

802.3 LAG ports support 128 users.

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
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.

S-Series and S-Series Standalone Customer Release Notes

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.

Multi-User Capacities Licensing

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.

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

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

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

S-Series and S-Series Standalone Customer Release Notes

- 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 127 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.

Multi-User 802.1X

Authentication of multiple 802.1X clients on a single port is supported. This feature will only operate correctly when the intermediary 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.

S-Series and S-Series Standalone Customer Release Notes

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.)

S-Series and S-Series Standalone Customer Release Notes

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 Control 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/ SSA150A	SSA130
ARP Entries (per router / per chassis)	32,000					
Static ARP Entries	1,024					
IPv4: Route Table Entries	1.6M	100,000	100,000	1.6M	100,000	100,000
IPv6: Route Table Entries	50,000	25,000	25,000	50,000	25,000	25,000
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					

S-Series and S-Series Standalone Customer Release Notes

	S180/S140/ S155	S150	S130	SSA180	SSA150/ SSA150A	SSA130
ECMP Paths	8					
Static VRFs	128	128	*Licensed	128	128	*Licensed
Dynamic VRFs	16	16	*Licensed	16	16	*Licensed
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					

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/ SSA150A	SSA130
DHCP Server Leases	5,000					
DHCP Pools	100					

TWCB Capacities

	S180/S140/ S155	S150	S130	SSA180	SSA150/ SSA150A	SSA130
Bindings	128K	64K	64K	128K	128K	64K
Caches	500					
Servers Farms	50					
WebCaches	50					

LSNAT Capacities

	S180/S140/ S155	S150	S130	SSA180	SSA150/ SSA150A	SSA130
LSNAT Bindings	64K	64K	-	128K	128K	-
SLB Real Server IPs	500	500	-	640	640	-
SLB Server Farms	320					
VIP Addresses	1,000					
SLB Virtual Servers	500					
Sticky Entries	64K	64K	-	128K	128K	-

S-Series and S-Series Standalone Customer Release Notes

NAT Capacities

	S180/S140/ S155	S150	S130	SSA180	SSA150/ SSA150A	SSA130
Bindings	64K	64K	-	128K	128K	-
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					

Some of these limits may **not** be enforced by the firmware and may cause unknown results if exceeded.

License Features

The S-EOS-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-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 the 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.

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)

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	5.0 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

S-Series and S-Series Standalone Customer Release Notes

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)
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

S-Series and S-Series Standalone Customer Release Notes

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
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+
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
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
10GB-4-C03-QSFP	10Gb, Copper DAC Fan out, 4xSFP+ to QSFP+, 3m
Adapters/Cables	Description
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.

Only Extreme sourced (SR4/LR4) 40 Gigabit optical transceivers are supported. Use of any other transceiver types will result in an error.

Example Message for 40G cables that are unrecognized or unauthenticated

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

S-Series and S-Series Standalone Customer Release Notes

Example port hold-down message for unauthenticated 40G optical transceiver

- System[1]port fg.1.4 will remain down because the pluggable module('manufacturer'/'part no.') is not supported

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.

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
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

S-Series and S-Series Standalone Customer Release Notes

Features		
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)
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

FIRMWARE CHANGES AND ENHANCEMENTS:

NOTICE: Minimum FW Revision Support Change

The ST8206-0848-F8, SG8201-0848-F8 fabric modules minimum FW version has changed. You may be required to upgrade FW in order to continue using these modules. Two High Availability Upgrade (HAU) FW tracks have been updated to provide support:
 8.12.02.0006 or greater 8.12.XX versions will support the fabrics,
 8.22.01.0022 or later versions will support the fabrics

S-Series and S-Series Standalone Customer Release Notes

Feature Enhancements in 8.12.04.0003

Loop Protect Enhancements in 8.12.04.0003
<p>LoopProtect is enhanced to coordinate more closely with changes in point-to-point status, mainly through the MultiSource function. LoopProtect does not allow a non-point-to-point port to forward. The MultiSource function detects the reception of BPDUs from multiple sources on a given port. If this occurs and the AdminPoinToPoint value indicates that auto-detection is used to determine the operational point-to-point value (which is the default administrative value), it sets the operational point-to-point value to false.</p> <p>Prior function of LoopProtect would use the operational point-to-point value of false to prevent the refresh of the LoopProtect BPDUs reception timer for a designated port. On expiry, a LoopProtect event would be indicated and the port state would become discarding.</p> <p>The new function is to force timer expiry immediately for a port on transition to non-point-to-point status for both designated and root port roles. Rather than waiting seconds for the port to become discarding, it discards immediately. For source address comparison, the MultiSource function has also been updated to use the MAC address of the transmitting bridge embedded in the BPDU, rather than the BPDU packet header source MAC address. This removes any delays in MultiSource detection allowed for LAG ports due to the fact that a LAG may change its port MAC address based on changes to port membership. Also, a change in point-to-point status due to port duplex change would result in the same quicker response.</p>

Problems Corrected in 8.12.03.0001

POE Problems Corrected in 8.12.03.0001	Introduced in Version:
<p>Occasional failure to communicate with PoE power supplies may result in a momentary drop of POE power for the entire module.</p> <p>Messages similar to the following might be observed during the failure:</p> <p><3>System[4]PoE power supply 4 information is unavailable. <3>System[4]PoE power supply 3 information is unavailable. <3>System[4]PoE power supply 2 information is unavailable. <3>System[4]PoE power supply 1 information is unavailable.</p>	8.11.01
Platform Problems Corrected in 8.12.03.0001	Introduced in Version:
S180 and S140 modules may exhibit a MAC chip XGXS error during initialization.	8.11.01

Problems Corrected in 8.12.02.0006

ARP Problems Corrected in 8.12.02.0006	Introduced in Version:
The "show running-config" command may not display all static ARP/ND entries that are configured. ST8206-0848-F8 SG8201-0848-F8	7.00.01
LSNAT Problems Corrected in 8.12.02.0006	Introduced in Version:
It is possible that while processing 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

S-Series and S-Series Standalone Customer Release Notes

NONVOL Problems Corrected in 8.12.02.0006	Introduced in Version:
<p>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 m ajorId=140 failed retval=8, write_file_num=50 (0x00d12590 0x00a79af4 0x 00a81504 0x01686324 0x00000000) A core file will be generated.</p>	8.11.04

Platform Problems Corrected in 8.12.02.0006	Introduced in Version:
<p>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)</p>	8.11.01

TCP Problems Corrected in 8.12.02.0006	Introduced in Version:
<p>Continuous polling 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

Problems Corrected in 8.12.01.0011

802.1x Problems Corrected in 8.12.01.0011	Introduced in Version:
<p>The "show config dot1x" command may leak memory.</p>	8.11.01

Auto Tracking Problems Corrected in 8.12.01.0011	Introduced in Version:
<p>The "show config auto-tracking" command may leak memory.</p>	8.01.01

Chassis Bonding Problems Corrected in 8.12.01.0011	Introduced in Version:
<p>In a stable 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></p>	7.72.01

CLI Problems Corrected in 8.12.01.0011	Introduced in Version:
<p>For a S180 dual speed 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.</p>	8.11.01
<p>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 all the entire range of quad ports associated with the port tg.x.x-y change.</p>	8.11.01

S-Series and S-Series Standalone Customer Release Notes

CLI Problems Corrected in 8.12.01.0011	Introduced in Version:
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 ends up in 4x10g mode.	8.11.01
Issuing a "show config" or "show config pwa" will cause a small amount of memory to leak per iteration.	8.11.01
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
High Availability Upgrade (HAU) Problems Corrected in 8.12.01.0011	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
IGMPv3 Problems Corrected in 8.12.01.0011	Introduced in Version:
While running IGMPv3 with 'include' source-list, a module crashes with a message containing to "Clgmp::GroupTableAddPortToGroupEntry Src port mismatch".	7.30.01
LAG Problems Corrected in 8.12.01.0011	Introduced in Version:
Source Port Exclusion not applied to LAG w/16 + links.	8.01.01
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
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
NAT Problems Corrected in 8.12.01.0011	Introduced in Version:
It is possible when upgrading from 08.02.xx or downgrading to 08.02.xx that the following NAT config may be missing. <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 8.12.01.0011	Introduced in Version:
Netflow cannot be enabled on 40G Ethernet ports.	8.01.01
OSPF Problems Corrected in 8.12.01.0011	Introduced in Version:
An assertion failure and reset occurs and recorded in message log as "SMS assert in qoamlsts.c at line 1218".	7.00.01

S-Series and S-Series Standalone Customer Release Notes

OSPF Problems Corrected in 8.12.01.0011	Introduced in Version:
If an OSPF interface is configured as passive and has a secondary address, the secondary address will be announced even if it is not within a configured OSPF network.	8.11.01
PKI Problems Corrected in 8.12.01.0011	Introduced in Version:
When configuring an X.509 certificate via the "set pki certificate <pki-cert-list>" command a warning is displayed if the same certificate already exists on the list, and the user is prompted as to whether or not they want to accept the new certificate. The user can avoid this prompt (to avoid breaking automated scripts) by specifying the "no-confirm" option on the command line. The "no-confirm" option should suppress the duplicate certificate warning as well as suppressing the prompt.	8.11.01
If a configuration file which contains PKI data is modified by an external text editor and that editor adds control characters (such as '\r' 0x0D), then sourcing the modified config file may not restore very large certificates (on the order of 10K PEM characters, which is the maximum allowed by the device).	8.11.01
Platform Problems Corrected in 8.12.01.0011	Introduced in Version:
Port advertisement settings are not persistent when auto negotiation is disabled.	7.91.01
Traffic performance may be severely degraded between S130/S150/S155 class IO blades and S180 class fabric blades, especially when it is near oversubscribed level.	8.01.01
If the rate of traffic destined to a front-panel port on a S140 Data Center switch exceeds the port's operating speed (i.e., oversubscribed egress), throughput may be degraded on other ports of the switch.	8.01.01
An unsupported Option Module will halt the board and not allow software upgrade until removed.	8.11.01
If the rate of traffic destined to a front-panel port on a S180 Data Center switch exceeds the port's operating speed (i.e., oversubscribed egress), throughput may be degraded on other ports of the switch.	8.11.01
When a high level of traffic is being received on the front-panel ports of an S140/S180 module, Spanning Tree may spuriously report that topology change events are occurring. This is a consequence of Spanning Tree BPDUs not being given higher priority than "best-effort" data traffic.	8.11.01
An unauthenticated 40GB-SR4-QSFP will show link if it is inserted into a 40G port after board boots when connected to another 40GB-SR4-QSFP.	8.11.01
A watchdog timeout reset occurs when boards are initializing and one or more boards in the chassis are experiencing out of band communication issues.	8.11.01
Quarantine Problems Corrected in 8.12.01.0011	Introduced in Version:
The "show config quarantine-agent" command may leak memory.	8.01.01
RADIUS Problems Corrected in 8.12.01.0011	Introduced in Version:
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

S-Series and S-Series Standalone Customer Release Notes

Tunneling Problems Corrected in 8.12.01.0011	Introduced in Version:
The tunnel manager connection setup firmware impacted the rate of non-tunnel connection setup.	8.02.01

Feature Enhancements in 8.11.05.0006

Transceiver Enhancements in 8.11.05.0006
<p><i>CWDM support:</i> 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+</p> <p><i>Additional DWDM support:</i> 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+</p>

Problems Corrected in 8.11.05.0006

ACLs Problems Corrected in 8.11.05.0006	Introduced in Version:
When the platform connection look-up level has been raised from L3 to L4 by application of an ACL, removing the ACL does not cause the look-up level to be reduced to L3.	7.40.01
When adding entries to an access-list, duplicates of existing entries are no longer accepted.	7.00.01

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.tDSrcv7]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

S-Series and S-Series Standalone Customer Release Notes

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]generateIndex() :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

S-Series and S-Series Standalone Customer Release Notes

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

Feature Enhancements in 8.11.04.0005

Tranceiver 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]ClgmpEtscc::DistGrpTbIRecvDistributedAdd 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

S-Series and S-Series Standalone Customer Release Notes

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 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]nvFilePtrMgr::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.tusrApplnit]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.tusrApplnit]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.tusrApplnit]nvFilePtrMgr::flush(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.tusrApplnit]nvFilePtrMgr::verify(0) checksum failure. store=4, fileIdx=0.37, udpSum=0x8f8dd5a, sumCount=65527	3.00.33
The nonvol cleanup task can cause a DSI reset: Exc Vector: DSI exception (0x00000300) Thread Name: tNVolCUp	3.00.33
The nonvol cleanup task can become stuck causing high system utilization: debug utilization show -i NAME TID PRI STATUS 5sec 1min 5min Got tid = 1 from successful call to getNextTaskId(). tNVolCUp 240412704 195 READY 99.37 99.28 99.27	3.00.33
PLATFORM Problems Corrected in 8.11.04.0005	Introduced in Version:
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
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

S-Series and S-Series Standalone Customer Release Notes

PLATFORM Problems Corrected in 8.11.04.0005	Introduced in Version:
System logs the message "bcmStrat[1.tNimIntr]MEM_FAIL_INT_STAT=0x00200000, EP_INTR_STATUS=0x00000000, IPO_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, IPO_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, IPO_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, IPO_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

S-Series and S-Series Standalone Customer Release Notes

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

S-Series and S-Series Standalone Customer Release Notes

BGP Problems Corrected in 8.11.03.0005	Introduced in Version:
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 output:00002bef binding failed (0x00c77d1c 0x00574058 0x015830e4 0x015756f4 0x0157ebec 0x01830ea0 0xeeeeeeee)" is logged on this error.	8.11.01
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

S-Series and S-Series Standalone Customer Release Notes

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 the 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 disable resources would be dropped until it aged out of hardware.	8.11.01
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

S-Series and S-Series Standalone Customer Release Notes

NAT Problems Corrected in 8.11.03.0005	Introduced in Version:
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
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

S-Series and S-Series Standalone Customer Release Notes

Platform Problems Corrected in 8.11.03.0005	Introduced in Version:
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
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.tusrAppInIt]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

S-Series and S-Series Standalone Customer Release Notes

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>

S-Series and S-Series Standalone Customer Release Notes

SK8009-1224 S-Series S180 Class I/O Module -24 Ports 10GBASE-T via RJ45, VSB expansion slot
(Used in S4/S6/S8)

Support for HW VSB models including:

SOV3208-0202 S-Series Option Module (Type2)- 2 port VSB Option Module (Compatible with Type2 option slots on S140/S180 modules only)

SOV3008-0404 S-Series VSB Expansion Module - 4 port VSB Module (Compatible with S180 Class 10Gb/40Gb I/O modules only)

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.

S-Series and S-Series Standalone Customer Release Notes

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.

S-Series and S-Series Standalone Customer Release Notes

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.

S-Series and S-Series Standalone Customer Release Notes

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.

S-Series and S-Series Standalone Customer Release Notes

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

S-Series and S-Series Standalone Customer Release Notes

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 lldpStatsRxPortAgeoutsTotal 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 qodmnsa.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 tDSsync5.	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

S-Series and S-Series Standalone Customer Release Notes

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:

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.

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.

If the chassis this module is installed in is running firmware 8.02 or less, use the following instructions to upgrade the module firmware:

Installing in an S3 Chassis Currently Running 8.01 Firmware

If you are installing this Module in an S3 chassis that is currently running 8.01 Firmware, the chassis firmware must be upgraded:

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

S-Series and S-Series Standalone Customer Release Notes

Installing in an S3 Chassis Currently Running 8.02 Firmware

If you are installing this Module in an S3 Chassis that is currently running 8.02 firmware, you must:

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).

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.

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 `network` command under OSPF configuration mode. The `network` 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:

- LAG capacities are reduce to 126 for chassis, 61 for SSAs,
- GRE Tunnels are not supported,
- Remote Port mirrors are not supported,
- Mirror N Packet mirrors are not 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.

S-Series and S-Series Standalone Customer Release Notes

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
RFC1323	TCP Extensions for High Performance
RFC1349	Type of Service in the Internet Protocol Suite
RFC1350	TFTP

S-Series and S-Series Standalone Customer Release Notes

RFC No.	Title
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
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
RFC2138	RADIUS Authentication
RFC2233	The Interfaces Group MIB using SMIv2
RFC2236	Internet Group Management Protocol, Version 2
RFC2260	Support for Multi-homed Multi-prov

S-Series and S-Series Standalone Customer Release Notes

RFC No.	Title
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
RFC2519	A Framework for Inter-Domain Route Aggregation
RFC2545	BGP Multiprotocol Extensions for IPv6
RFC2547	BGP/MPLS VPNs
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
RFC2618	RADIUS Client MIB
RFC2620	RADIUS Accounting 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

S-Series and S-Series Standalone Customer Release Notes

RFC No.	Title
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
RFC2865	RADIUS Accounting
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
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
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

S-Series and S-Series Standalone Customer Release Notes

RFC No.	Title
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
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 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

S-Series and S-Series Standalone Customer Release Notes

RFC No.	Title
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
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	MSDPDeployment Scenarios
RFC4624	MSDP MIB
RFC4632	Classless Inter-Domain Routing (CIDR)
RFC4659	BGP-MPLS IP VPN Extension for IPv6 VPN

S-Series and S-Series Standalone Customer Release Notes

RFC No.	Title
RFC4668	RADIUS Client MIB
RFC4670	RADIUS Accounting 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
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
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

S-Series and S-Series Standalone Customer Release Notes

RFC No.	Title
RFC6106	IPv6 RA Options for DNS Configuration
RFC6164	Using 127-Bit IPv6 Prefixes on Inter-Router Links
RFC6296	IPv6-to-IPv6 Network Prefix Translation
RFC6549	OSPFv2 Multi-Instance Extensions
RFC6565	OSPFv3 as PE/CE Protocol for BGP L3 VPNs
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)
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-rrrp-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-ENVIRONMENTAL-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-SPANNING-TREE-MIB
ENTERASYS-AAA-POLICY-MIB	ENTERASYS-NETFLOW-MIB (v5 & v9)	IEEE8023-DOT3-LLDP-EXT-V2-MIB
ENTERASYS-CLASS-OF-SERVICE-MIB	ENTERASYS-OIDS-MIB DEFINITIONS	LLDP-MIB
ENTERASYS-CONFIGURATION-MANAGEMENT-MIB	ENTERASYS-OSPF-EXT-MIB	LLDP-EXT-MED-MIB
ENTERASYS-CONVERGENCE-END-POINT-MIB	ENTERASYS-PFC-MIB-EXT-MIB	LLDP-EXT-DOT1-MIB
ENTERASYS-DIAGNOSTIC-MESSAGE-MIB	ENTERASYS-PIM-EXT-MIB	LLDP-EXT-DOT3-MIB
ENTERASYS-DNS-RESOLVER-MIB	ENTERASYS-POLICY-PROFILE-MIB	LLDP-EXT-DOT3-V2-MIB
ENTERASYS-DVMRP-EXT-MIB	ENTERASYS-POWER-ETHERNET-EXT-MIB	LLDP-EXT-DOT3-V2-MIB (IEEE 802.3-2009) ETS Admin table read only

S-Series and S-Series Standalone Customer Release Notes

Title	Title	Title
ENTERASYS-ETH-OAM-EXT-MIB	ENTERASYS-PTOPO-MIB-EXT-MIB	RSTP-MIB
ENTERASYS-IEEE8021-BRIDGE-MIB-EXT-MIB	ENTERASYS-PWA-MIB	U-BRIDGE-MIB
ENTERASYS-IEEE8021-SPANNING-TREE-MIB-EXT-MIB	ENTERASYS-RESOURCE-UTILIZATION-MIB	USM-TARGET-TAG-MIB
ENTERASYS-IEEE8023-LAG-MIB-EXT-MIB	ENTERASYS-RIPv2-EXT-MIB	ENTERASYS-TWCB-MIB
ENTERASYS-IETF-BRIDGE-MIB-EXT-MIB	ENTERASYS-RMON-EXT-MIB	ENTERASYS-NAT-MIB
ENTERASYS-IETF-P-BRIDGE-MIB-EXT-MIB	VSB-SHARED-SECRET-MIB	ENTERASYS-LSNAT-MIB
ENTERASYS-IF-MIB-EXT-MIB	ENTERASYS-SNTP-CLIENT-MIB	ENTERASYS-VRRP-EXT-MIB DEFINITIONS
ENTERASYS-IP-SLA-MIB	ENTERASYS-RADIUS-ACCT-CLIENT-EXT-MIB	SNMP-RESEARCH-MIB
	ENTERASYS-RADIUS-AUTH-CLIENT-MIB	

Extreme Networks Private Enterprise MIBs are available in ASN.1 format from the Extreme Networks web site at: <http://www.extremenetworks.com/support/enterasys-support/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 2863	linkDown linkup
RFC 2922	ptopoConfigChange
RFC 2787	vrrpTrapNewMaster vrrpTrapAuthFailure
RFC 3621	pethPsePortOnOffNotification pethMainPowerUsageOnNotification pethMainPowerUsageOffNotification

S-Series and S-Series Standalone Customer Release Notes

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

S-Series and S-Series Standalone Customer Release Notes

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

S-Series and S-Series Standalone Customer Release Notes

GLOBAL SUPPORT:

By Phone: +1 800-998-240 (toll-free in U.S. and Canada)

For the Extreme Networks Support toll-free number in your country:

www.extremenetworks.com/support/

By Email: support@extremenetworks.com

By Web: www.extremenetworks.com/support/

By Mail: Extreme Networks, Inc.
145 Rio Robles
San Jose, CA 95134 (USA)