Customer Release Notes

S-Series® and S-Series® Standalone
Firmware Version 8.22.01.0023
April 2014

INTRODUCTION:

This document provides specific information for version 08.22.01.0023 of firmware for the S155, S150 and S130 class of S-Series Modules and the S-Series Standalone (SSA) 1RU chassis. The S155/S150 and S130 modules may be installed in the S8, S6, S4 and S1 chassis. The 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:

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

Standalone Modules			
SSA-T1068-0652	SSA-T4068-0252	SSA-G1018-0652	

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/

S-Series and S-Series Standalone Customer Release Notes

Status	Firmware Version	Product Type	Release Date
Current Version	8.22.01.0023	Customer Release	April 2014
Previous Version	8.21.03.0003	Customer Release	February 2014
Previous Version	8.21.02.0002	Customer Release	December 2013
Previous Version	8.11.05.0007	Customer Release	December 2013
Previous Version	8.11.04.0006	Customer Release	October 2013
Previous Version	8.11.03.0006	Customer Release	August 2013
Previous Version	8.11.02.0002	Customer Release	July 2013
Previous Version	8.11.01.0015	Customer Release	June 2013
Previous Version	8.02.01.0012	Customer Release	March 2013
The current image does not support the S180/S140 and SSA180/SSA150A classes. An alternate image is available for mixed class configurations that include the S140/S180 and S130/S150/S155 classes and SSA180/SSA150A and SSA130/SSA150 classes.			

HIGH AVAILABILITY UPGRADE (HAU) FW COMPATIBILITY:

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

```
7865eaa158e33077c38921d35877ee78744fd212
```
(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.

INSTALLATION INFORMATION:

Installing an I/O or I/O Fabric Module

When installing a new S130/S150/S155 type 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.

S-Series and S-Series Standalone Customer Release Notes

Modules Minimum FW Version Required:

S155 Class		S150 Class		S130 Class	
SK5208-0808-F6	07.21.02.0002	SK1208-0808-F6	07.01.01.000X	ST4106-0348-F6	07.02.02.0002
ST5206-0848-F6		ST1206-0848-F6		ST4106-0248	
SG5201-0848-F6		SG1201-0848-F6		SG4101-0248	
		SK1008-0816			
		ST1206-0848			
		SG1201-0848			

Option Modules			
Series 1		Series 2	
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	8.02.01.0012

Multislot Chassis Minimum FW Version Required:

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

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

SSA Class	
SSA-T4068-0252	07.01.01.000X
SSA-T1068-0652	
SSA-G1018-0652	

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.

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/S3/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.

S-Series and S-Series Standalone Customer Release Notes

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	2,560 (2.5K) (S130 Class)
S3-Chassis-POE4	
S3-Chassis-A	
S3-Chassis-POEA	
S1-Chassis	2,048 (2K)
S1-Chassis-A	

S150/S155 Class modules Multi-User Capacities
Each of the S150 modules contributes 1024 users to the overall chassis capacity. Each S150/S155 class module has unrestricted access to the entire system user capacity. This allows for up to the entire system's user capacity to be consumed on a single port.

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

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

Uplink ports installed on the S130 modules, defined as modular SFP, 10 Gbps, and 100Mb FX ports, support up to 128 authenticated users per port. This includes modules SOK1208-0102, SOK1208-0104, SOK1208-0204, and SOG1201-0112. 802.3 LAG ports support 128 users.

SSA User Capacities:

Chassis Type	Maximum User Capacity
SSA-T4068-0252	512
SSA-T1068-0652	2048 (2K)
SSA-G1018-0652	2048 (2K)

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 and SSA-G1018-0652.

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 S130 class the default capacity will be increased from 512 to 1024 user per SSA.
- In the SSA150 class the default will be increased from 2048 to 4096 users per SSA.

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

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. (Licensed Feature)

S-Series and S-Series Standalone Customer Release Notes

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

Link Aggregation (LAG)
The S-Series chassis, S1/S3/S4/S8, supports a total of 190 LAGs per chassis with up to 64 ports per LAG. The SSA products support up to 62 LAGs per SSA with up to 64 ports per LAG.

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

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

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

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

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

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

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

Features, Scale and Capacity
Each release of S-Series firmware contains specific features and associated capacities or limits. The CLI command "show limits" provides a detailed description of the features and capacity limits available on your specific HW with its current licensing. Please use this command to get a complete list of capacities for this release.

F0615-O

Router Capacities (Brief)

ARP Entries (per router / per chassis)	32,000
Static ARP Entries	1,024
IPv4: Route Table Entries(S155)	1,600,000
IPv4: Route Table Entries(S130/S150, SSA130/SSA150)	100,000
IPv6: Route Table Entries(S155)	50,000
IPv6: Route Table Entries(S130/S150, SSA130/SSA150)	16,000
IPv4: Router interfaces	1024
IPv6: Router interfaces	256
OSPF Areas	16
OSPF LSA(s)	50,000
OSPF Neighbors	60
Static Routes	2048
RIP Routes	3,000
Configured RIP Nets	300
VRRP Interfaces	1024
Routed Interfaces	1024
ACLs	1,000
-Access Rules	5,000
-Access Rules – Per ACL	5,000
Policy Based Routing Entries	100
ECMP Paths	8
Static VRFs (S150/S155 Class, Licensed S130 Class)	128
Dynamic VRFs (S150/S155 Class, Licensed S130 Class)	16
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

IGMP/MLD Static Entries	64
IGMP/MLD *,G and S,G Groups[1]	64K
IGMP/MLD Snooping Flow Capacity(SSAs/S130 Class)	8K
IGMP/MLD Snooping Flow Capacity (S150/S155 Class)	16K
Multicast Routing (PIM/DVMRP flows) (SSAs/S130 Class)	8K
Multicast Routing (PIM/DVMRP flows) (S150/S155 Class)	16K
Multicast Routing (PIM/DVMRP flows) (S150/S155 Class) When Virtual Switch Bonded in a S4, S6 or S8 chassis	8K
IGMP/MLD Clients[2]	64K

   [1]   Group entries may be consumed for each egress VLAN of a routed flow.
   [2]   A client is defined as a reporter subscribing to a *, G or S, G group, or sourcing a multicast flow.

DHCP Capacities

DHCP Server Leases	5,000
DHCP Pools	100

TWCB Capacities

Bindings (SSA S150 Class)	131,072
Bindings (S150/S155 Class)	65,536
Caches	500
ServersFarms	50
WebCaches	50

LSNAT Capacities

LSNAT Bindings (SSA S150 Class)	131,072
LSNAT Bindings (S150/S155 Class)	65,536
SLB Real Server IPs (S150/S155)	500
SLB Real Server IPs (SSA150)	640
SLB Server Farms	320
VIP Addresses	1,000
SLB Virtual Servers	500
Sticky Entries (SSA S150 Class)	131,072
Sticky Entries (S150/S155 Class)	65,536

NAT Capacities

Bindings (SSA S150 Class)	131,072
Bindings (S150/S155 Class)	65,536
IP Addresses	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 functionality in the S150 class is included without the need for a license.)

The S-EOS-S150 license adds GRE tunnel 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.

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

1Gb Optics:

MGBICs	Description
MGBIC-LC01	1 Gb, 1000Base-SX, IEEE 802.3 MM, 850 nm Short Wave Length, 220/550 M, LC SFP
MGBIC-LC03	1 Gb, 1000Base-SX-LX/LH, MM, 1310 nm Long Wave Length, 2 Km, LC SFP
MGBIC-LC07	1 Gb, 1000Base-EZX, IEEE 802.3 SM, 1550 nm Long Wave Length, 110 Km, LC SFP (Extended Long Reach)
MGBIC-LC09	1 Gb, 1000Base-LX, IEEE 802.3 SM, 1310 nm Long Wave Length, 10 Km, LC SFP
MGBIC-MT01	1 Gb, 1000Base-SX, IEEE 802.3 MM, 850 nm Short Wave Length, 220/550 M, MTRJ SFP
MGBIC-02	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, 10 km LC SFP+
10GB-BX10-D	10Gb, Single Fiber SM, Bidirectional, 1330nm Tx / 1270nm Rx, 10 km SFP+
10GB-BX10-U	10Gb, Single Fiber SM, Bidirectional, 1270nm Tx / 1330nm Rx, 10 km SFP+
10GB-BX40-D	10Gb, Single Fiber SM, Bidirectional, 1330nm Tx / 1270nm Rx, 40 km SFP+
10GB-BX40-U	10Gb, Single Fiber SM, Bidirectional, 1270nm Tx / 1330nm Rx, 40 km SFP+
SFP+ Copper	**Description**
10GB-C01-SFPP	10Gb pluggable copper cable assembly with integrated SFP+ transceivers, **1 m**
10GB-C03-SFPP	10Gb pluggable copper cable assembly with integrated SFP+ transceivers, **3 m**
10GB-C10-SFPP	10Gb pluggable copper cable assembly with integrated SFP+ transceivers, **10 m**
SFP+ Laserwire	**Description**
10GB-LW-SFPP	SFP+ Laserwire Transceiver Adapter
10GB-LW-03	Laserwire Cable **3 m**
10GB-LW-05	Laserwire Cable **5 m**
10GB-LW-10	Laserwire Cable **10 m**
10GB-LW-20	Laserwire Cable **20 m**
10GB-F10-SFPP	10Gb, Active optical direct attach cable with 2 integrated SFP+ transceivers, **10m**
10GB-F20-SFPP	10Gb, Active optical direct attach cable with 2 integrated SFP+ transceivers, **20m**
SFP+ DWDM Optics	**Description**
10GB-ER21-SFPP	10GB-ER, DWDM CH21 SFP+
10GB-ER23-SFPP	10GB-ER, DWDM CH23 SFP+
10GB-ER24-SFPP	10GB-ER, DWDM CH24 SFP+
10GB-ER29-SFPP	10GB-ER, DWDM CH29 SFP+
10GB-ER31-SFPP	10GB-ER, DWDM CH31 SFP+
10GB-ER33-SFPP	10GB-ER, DWDM CH33 SFP+
SFP+ CWDM Optics	**Description**
10GB-LR271-SFPP	10Gb, CWDM SM, 1271 nm, 10 km, LC SFP+
10GB-LR291-SFPP	10Gb, CWDM SM, 1291 nm, 10 km, LC SFP+
10GB-LR311-SFPP	10Gb, CWDM SM, 1311 nm, 10 km, LC SFP+
10GB-LR331-SFPP	10Gb, CWDM SM, 1331 nm, 10 km, LC SFP+

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

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

S-Series and S-Series Standalone Customer Release Notes

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

SUPPORTED FUNCTIONALITY:

Features		
Multiple Authentication Types Per Port - 802.1X, PWA+, MAC	Layer 2 through 4 VLAN Classification	Entity MIB
Multiple Authenticated Users Per Port - 802.1X, PWA+, MAC	Layer 2 through 4 Priority Classification	IP Routing
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 Mirror	ICMP
GARP VLAN Registration Protocol (GVRP)	Flow Setup Throttling	Protocol Independent Multicast - Sparse Mode (PIM-SM)
802.1p – Traffic Class Expediting	MAC locking (Static/Dynamic)	Proxy ARP
802.1w – Rapid Reconfiguration of Spanning Tree	Node/Alias table	Basic Access Control Lists
802.1s – Multiple Spanning Trees	Policy-Based Routing	Extended ACLs
802.1t – Path Cost Amendment to 802.1D	SSH v2	Auto MDI-X Media Dependent Interface Crossover Detect (Enhanced for non auto negotiating ports)
802.3 – 2002	OSPF NSSA, equal cost multi-path	DHCP Server
802.3ad – Link Aggregation (128 users)	Audit trail logging	DHCP Relay w/option 82
802.3x – Flow Control	RADIUS Client	Jumbo Frame support
Load Share Network Address Translation (LSNAT)	FTP/TFTP Client	Directed Broadcast
Static Multicast Configuration	Telnet – Inbound/Outbound	Cisco CDP v1/2
Broadcast Suppression	Configuration File Upload/Download	CLI Management
Inbound and Outbound Rate Limiting	Text-based Configuration Files	DFE CPU and task Debugging
Strict and Weighted Round Robin Queuing	Syslog	RADIUS (Accounting, Snooping)
IGMP v1/v2/v3 and Querier support	Span Guard	Split RADIUS management and authentication
SMON Port and VLAN Redirect	RAD (Remote Address Discovery)	Link Flap detection
Spanning Tree Loop Protection	Cabletron Discovery Protocol (CDP)	Daylight Savings Time
TACACS+	NetFlow v5/v9	RFC 3580 with Policy support
Type of Service (ToS) Re-write	LLDP and LLDP-MED	Flex-Edge
NAT(Network Address Translation)	TWCB (Transparent Web Cache Balancing)	eBGP
iBGP	BGP Route Reflector	BGP 4 byte AS number
BGP Graceful Restart	BGP Route Refresh	BGP Extended Communities

F0615-O

Features		
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:

Feature Enhancements in 8.22.01.0023

Hardware Support Enhancements in 8.22.01.0023

Support for additional 10Gb active optical direct attach cable transceivers:

10GB-F10-SFPP 10Gb, Active optical direct attach cable with 2 integrated SFP+ transceivers, 10m

10GB-F20-SFPP 10Gb, Active optical direct attach cable with 2 integrated SFP+ transceivers, 20m

Captive Portal Re-direct Feature Enhancements in 8.22.01.0023

Captive Portal uses HTTP redirection to force a client's web browser to be redirected to a particular administrative web page. A network administrator can use this feature for authentication purposes (a user login and password), payment (i.e., at an airport hotspot), or usage-policy enforcement. This feature is an extension of the Policy infrastructure, where Policy Roles may be configured to force redirection of HTTP traffic.

OSPF Default Route Injection Feature Enhancements in 8.22.01.0023

Support for directly advertising a default route into OSPF has been added via the "default-information originate" command. There are two options available, advertise the default route into the OSPF domain, provided the advertising router already has a default route. Alternatively, advertise the default route regardless of whether the advertising router already has a default route. Option 2 is chosen by adding the "always" keyword to the "default-information originate" command.

BGP "Pass Through" Route Feature Target Support Enhancements in 8.22.01.0023

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

S-Series and S-Series Standalone Customer Release Notes

Problems Corrected in 8.22.01.0023

802.1x Problems Corrected in 8.22.01.0023	Introduced in Version:
802.1x may not require an 802.1x supplicant to wait the configured quiet period (set dot1x auth-config quietperiod <period> <port-string>) to start a new authentication after a failed authentication.	8.21.01

ARP Problems Corrected in 8.22.01.0023	Introduced in Version:
If system sends packet to a remote IP address, an ARP request for the remote IP address may be transmitted on a configured interface.	8.21.01

Auto-Negotiation Problems Corrected in 8.22.01.0023	Introduced in Version:
If "clear port advertise *.*.*" is executed on a system on which not all ports support auto-negotiation, the message "failed to set ifMauAutoNegCapAdvertisedBits on port x.y.z" will be displayed for each port that does not support auto-negotiation.	7.00.01
"Setting ifMauAutoNegRemoteFaultAdvertised (1.3.6.1.2.1.26.5.1.1.12) MIB value to offline(2) for a port brings the port down until reset, even if ifMauAutoNegRemoteFaultAdvertised value is changed to noError (1)."	5.11.21

Auto-Tracking Problems Corrected in 8.22.01.0023	Introduced in Version:
Auto-tracking radius-timeout-profile and radius-reject-profile per port configuration may allow profile ID configuration that is greater than allowed by the system.	8.01.01
Outputted log event from auto-tracking and quarantine-agent "Unable to set policy rule" port string is not user friendly.	8.01.01
If auto-tracking multiauth sessions are configured to be allowed on authentication required ports then unauthenticated traffic matching the auto-tracking multiauth session will be switched by the system.	8.01.01

BGP Problems Corrected in 8.22.01.0023	Introduced in Version:
When MPLS is disabled, established state with BGP peers are lost.	8.02.01
The S-Series router currently does not have a mechanism to replace the export route targets defined on a VRF with the existing route targets on an L3VPN route.	8.01.01
The 'show ip protocols' command output does not display the BGP max-as limit.	8.21.01
The "show ip bgp" output of the AS-Path will display incorrect AS numbers if the AS-Path is longer than 30 AS numbers.	7.20.01
The BGP Autonomous system number of 0 is accepted at the CLI even though the help indicates the minimum value is 1. In this case, "show configuration" output will not display the "router bgp <AS>" command.	7.20.01
The "show ip bgp <prefix> detail" command will display repeated instances of the same community and extended-community values in some cases.	7.91.01
The 'show ip bgp peer <ip> advertised-routes <prefix/length> detail' command does not always display the correct communities and extended-communities associated with the route.	7.20.01

BGP Problems Corrected in 8.22.01.0023	Introduced in Version:
Negating the BGP peer-group soft-reconfiguration command does not take effect. The show running-config output will indicate the command is negated, however the setting is not negated internally.	7.20.01
Redistribution of IS-IS into BGP under non-vrf address-family mode will result in show running-config output that is inconsistent with the required command syntax for the "match" path type options.	7.30.01
The following error message may occur if deleting an instance of a routing protocol which contains redistribution entries with multiple references to the same route-map: "Error decrementing route map <name> ."	7.20.01
The output of the "show ip bgp" command does not display any information under the AS-Path heading if the actual patch contains approximately 70 or more AS numbers.	7.22.01

Chassis Bonding Problems Corrected in 8.22.01.0023	Introduced in Version:
During a time of chassis instability, a module in a Bonded chassis may reset after logging an error with format similar to: "<0>Default[12.tBondProto]Assertion failed: hdr>reqGeneration == generation, file /firmware/common/chassisBond/01_06_16/src/chassis_bond_protoco l.cxx, line 599".	7.72.01
A module in a Bonded chassis may reset soon after power up. When this failure occurs, a message with format similar to: "Default[14.tBondProto]Assertion failed: hdr->reqGeneration == generation, file /firmware/common/chassisBond/01_05_15/src/chassis_bond_protocol.cxx, line 599" is logged.	7.72.01
Configure from file fails when enabling bonding. When this error occurs a message similar to "<2>System[12]Detected missing or reset module, aborting configure" is logged.	8.21.03
Bonding mode may be changed from software-assist to hardware at boot with mix of 8.12 and older firmware images in the chassis. When this occurs modules will reset and a message with following format will be logged: "Received Bonding mode = hardware from master. Rebooting ..."	8.11.01
Module in a software bonded system may reset while bonding ports are being enabled. A message similar to "1>DistServ[2.tDsBrdOk]serverWatchDog.6(PortInfo), client 106(Bonding) in recv for 6300 tics" is logged on this failure. A workaround is to wait 1 minute between bonding port enables.	8.21.01

CiscoDP Problems Corrected in 8.22.01.0023	Introduced in Version:
Cisco VTP packets are not forward when Cisco CDP is enabled.	7.91.01

ECMP Problems Corrected in 8.22.01.0023	Introduced in Version:
If an interface that is part of an equal cost multipath route goes down, host originated traffic to destinations in the route's subnet may temporarily fail.	8.21.01

Host Services Problems Corrected in 8.22.01.0023	Introduced in Version:

Some devices may reset after logging a message similar to the one listed below. This may occur intermittently on S140 modules during initialization of the onboard power controllers.
 	
Message   6/213 Exception PPC750 Info   08.11.04.0005 01/09/2014 11:14:31
  Exc Vector: DSI exception   (0x00000300)
  Thread Name: tRootTask
  Exc Addr: 0x0168ba70
  Thread Stack: 0x7dfffd10..0x7dfec7c0
  Stack Pointer: 0x7dfff4f0
  Traceback Stack:
 	
 GENERAL EXCEPTION INFO.......
	8.11.03

Messages like the following can be seen during slot resets in busy systems. Transmit errors(8) to slot # are preventing heartbeat checks.	7.72.01
No eligible master messages are misleading because slot/peer ID is mistaken for server ID.	7.03.05
"show system utilization storage" will report inaccurate size and available size for USB drives greater than 2G.	7.60.01
When updating to a new image that has microcode updates there are error messages displayed about the DOSFS/DOS volume. Example: ** This board is updating its hardware microcode. This process may take up to 3 minutes to complete. Please do not reset or disconnect power during this update. The board will be reset when this process is complete. ** Writing primary host controller : bank 1 block 0 Programming flash. Complete: 100% Verify succeeded. Writing alternate host controller : bank 0 block 0 Programming flash. Complete: 100% Verify succeeded. /flash0 Not a DOSFS Device /flash1 Not a DOSFS Device /flash2 Not a DOSFS Device <163>Dec 18 08:27:57 0.0.0.0 System[1]Resetting after host controller microcode update. <163>Dec 18 08:27:57 0.0.0.0 System[1]DOS Volume /flash0 cannot be set to read-only on unmount. <163>Dec 18 08:27:57 0.0.0.0 System[1]DOS Volume /flash1 cannot be set to read-only on unmount. <163>Dec 18 08:27:57 0.0.0.0 System[1]DOS Volume /flash2 cannot be set to read-only on unmount.	8.21.01

F0615-O

S-Series and S-Series Standalone Customer Release Notes

Host Services Problems Corrected in 8.22.01.0023	Introduced in Version:
Usually on a reboot after an uncontrolled reset (power-loss, board pull, exception, DSI, watchdog reset) you may see the following file system error during initialization: /flash2/ - disk check in progress ... "/flash2/usrroot/someFileName" too many clusters in file, adjusted. Errors detected. All corrections stored to disk and lost chains recovered.	7.30.01
Continuous poll of TCP or UDP mibs may result in the exhaustion of memory resulting in an out of memory reset action on a specific slot.	7.40.00
The "show running-config" command may not display all static ARP/ND entries that are configured.	7.00.01
Performing the "show vlan portinfo" CLI command under configurations where there are many VLANs in use may lead to the CLI becoming inoperable, or the system to reset.	8.21.01

In the unexpected event where resources needed to transmit a routed L3 Multicast packet failed to be obtained, a blade will reset, and leave a message in log similar to:
Message 9/333 Exception PPC750 Info   08.21.02.0002 12/21/2013 23:22:53
  Exc Vector: DSI exception   (0x00000300)
  Thread Name: tDispatch
  Exc Addr: 0x0191e77c
  Thread Stack: 0x06921000..0x06914000
  Stack Pointer: 0x06920f40
  Traceback Stack......
	7.00.01

Doing a set on a large range of data could cause a board reset. Example: cfm vlan-table primary 99 selector 1-98,100-4094. The syslog will show an error similar to below: <1>NonVol[1.tNVolCUp]cleanup:Remove() on store=0, fileIndex=2863311530 majorId=140 failed retval=8, write_file_num=50 (0x00d12590 0x00a79af4 0x00a81504 0x01686324 0x00000000) A core file will be generated.	8.21.01
Infrequently, when switch is adding (encaping) tunnel headers, a message similar to: <163>Dec 5 15:11:28 100.10.10.22 PiMgr[16.tDispatch]piMgrBindSystemPortAndHwPort(0,0):Port(s) are already bound. pimSystemPortToHwPort[0]=0x8000;pimHwPortToSystemPort[0]=0x100 <163>Dec 5 15:11:28 100.10.10.22 PiMgr[16.tDispatch]piMgrBindSystemPortAndHwPort(0,0):Port(s) are already bound. pimSystemPortToHwPort[0]=0x8000;pimHwPortToSystemPort[0]=0x100 <165>Dec 5 15:11:28 100.10.10.22 PiMgr[16.tDispatch]piMgrHwPortRxIcpu (131072,2,63,0,0x7eb82188,1052):piMgrBindSystemPortAndHwPort(0,0) failed;hwPort=0;portCount=43;tmpBufLen=700 may be logged.	7.40.00

IGMP Snooping Problems Corrected in 8.22.01.0023	Introduced in Version:
IGMP/MLD database entries (primarily, but not limited to IGMPv3/MLDv2 reporter state) do not age out correctly.	7.30.01
Legacy S-Series modules (S130/S150) with IGMP/MLD snooping enabled log messages similar to the following: <188>Jan 6 07:26:20 172.20.1.20 RfrmrHw[3.tDSrecv2]Invalid MCI - 1, for asic 0 <188>Jan 6 07:26:20 172.20.1.20 RfrmrApp[3.tDSrecv2]addPortReframing, Error: Could not convert mciIndex 8113 to UNTAGGED error status -2, then reset with a DSI exception in thread tDispatch.	8.21.01

S-Series and S-Series Standalone Customer Release Notes

IPV6 Forwarding Problems Corrected in 8.22.01.0023	Introduced in Version:
IPv6 packets destined to a remote subnet whose route has a link-local nexthop address and deferred to neighbor discovery for MAC address resolution may be transmitted with a destination MAC address of 00:00:00:00:00:00.	7.40.00
Some IPv6 addresses may remain in the tentative state when the master blade changes from one slot to another.	7.30.01
The IPv6 /128 host address of tunnel interfaces appears in output of 'show ipv6 route'.	8.21.01

IPv6 Neighbor Discovery Problems Corrected in 8.22.01.0023	Introduced in Version:
The router may not accept router advertisements to generate IPv6 addresses when the "ipv6 address autoconfig" command in applied to an interface.	8.21.01

LSNAT Forwarding Problems Corrected in 8.22.01.0023	Introduced in Version:
It is possible that while processing using sticky entries on a multiple blade system, that the sticky entry may not be deleted from all blades and subsequent sticky creations will fail causing a failure of processing LSNAT packets.	6.12.01
LSNAT FIN/RST timeout may not be properly set while running with multiple blades.	8.11.01
During debug session the command "/* rtr lsnat show data-plane bindings detail <ID>" caused a reset.	6.00.02

Mirroring Problems Corrected in 8.22.01.0023	Introduced in Version:
When the device acts as a Pseudowire tunneled endpoint the de-capsulated packet would not egress out a software bond port.	8.21.01
The blade may be reset (and continuously reset) with the following messages if the LAG used by IDS mirror has more than 2 ports: <3>Dune[5.tSlac]Err_id=0x16a1d3af: error in fap21v_sch_is_subflow_valid() ExitPlace (40) Params(0,0,0,0,0) <0>Dune[5.tSlac]Err_id=0x16a1d3af: error in fap21v_sch_is_subflow_valid() ExitPlace (40) Params(0,0,0,0,0).	8.11.01

MPLS Problems Corrected in 8.22.01.0023	Introduced in Version:
The command 'no mpls ip propagate-ttl [local]' did not affect packets originated by the local host. The packet's TTL was propagated to the MPLS label.	8.21.01
Learning the internet route table from BGP with LDP configured will exhaust system memory and cause a reset.	8.21.01
First hop may not respond when issuing traceroute from PE router across MPLS network.	8.21.01
The wrong MTU is specified in an ICMP Fragmentation Needed packet sent by an MPLS provider edge router for packets egressing an LSP and exceeding MTU of egress port.	8.12.01
Configuring LDP with the internet route table present in the system will exhaust memory and cause a system reset.	8.21.01
'mpls ip propagate-ttl' settings not stored in persistent storage.	8.21.01
Given MPLS/LDP enabled in the system, if it is disabled and enabled again, additional FECs are unnecessarily created for connected subnets.	8.21.01

MPLS Problems Corrected in 8.22.01.0023	Introduced in Version:
When both IPv4 and IPv6 prefixes were in use with label switched paths, display commands showing the mpls forwarding table contained invalid characters when trying to interpret the next hop addresses.	8.21.01

Multi-Auth Problems Corrected in 8.22.01.0023	Introduced in Version:
When multiauth sessions-unique-per-port is disabled and multiple multiauth agents are enabled a failure of one agent may cause additional agents to fail outputting the error message "Unable to set policy rule for mac XX-XX-XX-XX-XX-XX on system port 443".	8.01.01
Standardized multi-authentication session and idle timeout maximum values to be 172800 seconds.	8.01.01
With sessions-unique-per-port disabled and multiple authentication agents enabled and active, when a session moves from one slot to another, it may not session or idle timeout appropriately.	8.01.01
Multiauth Quarantine Agent sessions to not correctly apply policy if the policy maptable response is set to tunnel.	8.01.01

NAT Problems Corrected in 8.22.01.0023	Introduced in Version:
It is possible to for NAT to stop working due to running out of local buffers.	8.11.01

OSPF Problems Corrected in 8.22.01.0023	Introduced in Version:
If an OSPF area id is changed while an interface is transitioning to the DOWN state, an assert may occur in thread tRtrPtcls with the following log: "SMS assert in qopmmim2.c at line 1958 : is one of if_cb->repl.row_data.oper_status 4 qopm_mib_if_product_data.oper_states.down or qopm_mib_if_product_data.oper_states.act_failed".	8.11.01
If multiple OSPF processes learn the same route, metrics are not compared between them, both routes are installed in the route table as the administrative distance is the same and cannot be changed for an individual process.	7.00.01
An OSPF NSSA ABR configured as "transrole always" may not always be the translator.	8.01.01
OSPF log-adjacency cannot be removed with a no log-adjacency under router ospf <pid>.	8.01.01
If OSPF is configured to run BFD on a non-existent interface, the interface will not be displayed in show running. When the interface is created, the display will show, and BFD will run on that interface.	8.21.01
The display of an OSPF external LSA metric has the first byte truncated so the largest number displayed is 4095, though the real value may be up to 65535.	7.00.01
Using OSPF with a route-map for redistribution that sets the metric to a number greater than 65535 will result in an assert in thread tRtrPtcls with the following log in OSPFv2: "SMS assert in qodmbld3.c at line 471 : == (NBB_INT ((route_entry_ptr->path_cost) >> 24) 255 nbb_zero 0" and log "SMS assert in qod3bld2.c at line 214 : == (NBB_INT)((route_entry_ptr->path_cost) >> 24) 255 nbb_zero 0" for OSPFv3.	7.00.01
If the display of OSPF passive-interfaces in show running exceeded 80 characters, no interfaces are displayed.	8.21.01

F0615-O

OSPFv3 Problems Corrected in 8.22.01.0023	Introduced in Version:
If debug logging is turned on for OSPF, and filter route-maps are in use, the route-src is seen as 0.0.0.0 for local routes from our router ID.	8.01.01
When an OSPFv3 NSSA translator is configured to always be translating, it will not always translate if a higher router ID is also eligible.	8.01.01

PIM-DM Problems Corrected in 8.22.01.0023	Introduced in Version:
Enabling a PIM-DM upstream interface may result in an assert similar to "sms[1.tRtrPtcls]SMS assert in qptuftmr.c at line 1134 : (null) NTL_CLTIM_IN_LIST(&s_g->sg_join_timer) 0 (null) 0".	8.21.01
Changing the route to sources may result in an assert similar to "<0>sms[2.tRtrPtcls]SMS assert in qptuwsn2.c at line 669 : (null) QPTM_USM_S_G_GET_JDES(s_g) 0 (null) 0".	8.21.01
The use of IGMP V3 to PIM DM may cause crash.	8.21.01
The use of exclude mode in igmpv3 may result in a PIM DM assert.	8.21.01
Rebooting PIM DM source router may result in an assert similar to sms[1.tRtrPtcls]SMS assert in qptuwapi.c at line 602 : == 0 0 (s_g_i->sgi_flags & QPTM_DSM_SGI_WP_ALL) 4096.	8.21.01

PIM-SM Problems Corrected in 8.22.01.0023	Introduced in Version:
IGMPv3/MLDv2 source-specific reporter state is missing from layer3/router.	7.30.01
The internal IGMP/MLD database may be inconsistent across all modules after a bonded system is segmented, then re-joined. This can lead to incorrect multicast operation and/or inconsistent aging of entries.	7.30.01

Platform Problems Corrected in 8.22.01.0023	Introduced in Version:
Infrequently a board will not boot up and will end up in a halted state after a failure to read chassis type. The following message is output to the console when this error occurs: "A device within this chassis has encountered a hardware failure. Could not read chassis type. Please contact Support for the troubleshooting procedure to determine which device will possibly need to be repaired/replaced. Press <r> to reset board." This may not be a real hardware failure and a module reset will result in successful module initialization.	8.01.01
A watchdog timeout exception message may be logged (followed by a system reset) when a card with SFP+ ports is booted and one or more of these ports contain SFP modules.	8.21.01

Platform Problems Corrected in 8.22.01.0023	Introduced in Version:
If a macsource policy is applied, packet statistics from the following apps may not be valid: Smon stats Rmon Host/Matrix Router ACL Policy Routing Tunneling Policy In addition, if any SMON stats are enabled, messages similar to: SMON[6.tSmonCnt]getHwPrioStats(ge.6.3,0): packet count < previous 2/172401; detected 1 times, may be logged.	8.21.01
On a bonded system, a file may be left in an improper state which is identified and corrected by the file system verification and recovery tool that is run at each boot up. If a file in this state is detected, a set of messages like the following will be displayed during boot up. /flash2/ - disk check in progress ... "/flash2/usrroot/foobar672" too many clusters in file, adjusted. Errors detected. All corrections stored to disk and lost chains recovered. This state is recoverable and should have no effect on the normal operation of the file system.	7.60.01
10G port with 1G SFP doesn't propagate its advertised speed to link partner.	8.11.04
If a 1G SFP is inserted into one of the 10G ports on a SOTK2268-0212 or SOGK2218-0212 option module, the system will reset.	7.91.01
Doing a "dir" on a remote directory with a large number of files has a long delay before the output starts. Ex: A directory with 1000 files may take around 34 seconds before being displayed.	7.91.01
When doing a "dir" from CLI, if the directory is currently being modified (file being added/deleted)an incomplete listing can be returned.	7.91.01
During a system reset or a module reset, removal or insertion, it is possible to receive a DSI exception containing the text "DuneCB::RemoteSlot". This exception can be ignored once the system completes normal initialization.	8.11.01

PoE Problems Corrected in 8.22.01.0023	Introduced in Version:
POE redundancy shown as Not Supported after POE blade is reset and boots up.	7.60.01

Policy Problems Corrected in 8.22.01.0023	Introduced in Version:
Unable to clear all policy profiles with a single CLI command.	1.07.19
VLAN authorization commands allow for configuration and display of tunnel bridge ports although they are not supported port types for VLAN authorization.	8.21.01
Policy "macsource" rules configured with a mask less than 48 are not applied to traffic immediately upon configuration.	8.21.01

PWA Problems Corrected in 8.22.01.0023	Introduced in Version:
The "set pwa ipaddress <ip-address>" CLI command allows invalid values for the <ip-address> field.	4.00.50
PWA occasionally becomes unresponsive under heavy load. Device resets with this message in the log: <0>PWA[1.tPwaHtWD]pwaHttpReadWatchDog expired!	4.00.50

QOS Problems Corrected in 8.22.01.0023	Introduced in Version:
Default port group COS txq settings are applied to hardware VSB ports. No COS settings, default group or not, should ever be applied to hardware VSB ports.	8.11.01

RADIUS Problems Corrected in 8.22.01.0023	Introduced in Version:
RADIUS Server sticky sessions count may be inaccurate after session terminations.	8.11.01
If the radius algorithm is changed while multiauth sessions are active incorrect sticky session counters may be both displayed and used by the system.	8.11.01
RADIUS Dynamic Authorization responses cannot be sent in response to disconnect or change of authorization RADIUS Dynamic Authorization requests resulting in the error message "Unable to transmit the RADIUS frame" and retransmissions from the RADIUS server.	8.21.01

RMON Problems Corrected in 8.22.01.0023	Introduced in Version:
In rare instances, upon a blade reset in a multi-blade system with a large number of RMON alarms configured DSI exception, resets may occur.	5.01.58

Security Problems Corrected in 8.22.01.0023	Introduced in Version:
PWA will discard HTTP GET requests with HTTP headers that exceed 2048 bytes.	4.00.50

Spanning Tree Problems Corrected in 8.22.01.0023	Introduced in Version:
When the root port of a bridge receives a value for remainingHops greater than 63, there will be overflow when storing the value. For example, if the remainingHops value is 100, it will be stored as 36. This is because the field width is six bits. This is enough to hold the standard defined maximum value of 40. This is true for both cistRemainingHops and remainingHops for any MSTI. This only has a practical effect within an MST region. These values are not used external to the region. Note that values greater than 40 are non-conformant as of 802.1Q-2005 so are not likely to be seen.	8.21.01
In a multi-blade chassis or stack, when setting Spanning Tree stpmode to the value _none_, the non-master blades will still operate as if the mode were _ieee8021_ until those blades are reset.	8.21.01
When a device in a multi-blade chassis or a bonded setup fails, and that device contained the spanning tree root port for the bridge, the new root port, if there is one, may not take on its root role and therefore be stuck in a discarding state. If this does occur then a workaround for this is to disable the new root port (which will show a role of alternate port) and then reenable the port.	8.21.01

Static Routes Problems Corrected in 8.22.01.0023	Introduced in Version:
Static route leaking between non-global VRFs does not work. The routes are not promoted to the FIB.	8.21.01
A static host route whose address matches an LSNAT virtual server address will not be displayed in the router configuration if the LSNAT virtual server is up.	7.00.01

Tunneling Problems Corrected in 8.22.01.0023	Introduced in Version:
Host generated IPv6 packets that are encapsulated into an IP or GRE tunnel could have an incorrect DIP.	7.60.01
The software forwarding path was retrieving the GRE header when it was not part of the flow. This would sometimes cause the IP-in-IP to be translated as a L2 IP-in-IP flow.	7.62.02
When the device acts as a Pseudowire tunneled endpoint, the de-capsulated packet would not egress out a software bond port.	8.21.01
For pseudo-wire tunnels, the soft forwarding path was not adding the Chassis Bond header when going across a software bond.	8.21.01
The egress point of a Tagged IPv6-GRE(with GRE Keyword) tunnel would not decrement the inner IPv4 TTL or change the TOS due to hardware limitations.	8.21.01
L2 Tunnels across a Software Bond was not updating the L2 IP's total length field when adding the GRE header and Chassis Bondheader to the egress packet.	8.21.01
If tunnels are configured and at least one is up, then the connection database is no longer flushed when a route changes.	8.21.01
The ifMib returns a valid ifIndex row with no other valid leaves for internal ports that should be hidden.	8.21.01
Traceroute does not work from layer 3 VPN when configured over IP tunnels.	8.21.01
Infrequently, when switch is adding (encaping) tunnel headers, a message similar to: <0>chassis[9.tBcastStRx]powerSupplyComputeModuleConsumedPower :Invalid uplink number 0x00 detected in remote info table, may be logged	7.40.00
Infrequently, when switch is adding (encaping) tunnel headers, a message similar to: <3>chassis[1.tBcastStRx]remoteModuleInfoPowerUpdate(6,""):Unsupported board type found., may be logged.	7.40.00

VLAN Problems Corrected in 8.22.01.0023	Introduced in Version:
Performing the "show vlan portinfo" CLI command under configurations where there are many VLANs in use may lead to the CLI becoming inoperable, or the system to reset.	8.21.01

VRF Problems Corrected in 8.22.01.0023	Introduced in Version:
When clearing a vrf router config, "clear router vrf <NAME>" the error message "Error destroying BFD process 22185496: AMB_RC_NO_SUCH_OBJECT" is displayed, but has no adverse effect.	8.21.01
It is possible for show running to erroneously display "set router vrf vrf-management <VRFNAME> disable".	8.21.01

S-Series and S-Series Standalone Customer Release Notes

Problems Corrected in 8.21.03.0003

IGMP Snooping Problems Corrected in 8.21.03.0003	Introduced in Version:
Legacy S-Series modules (S130/S150) with IGMP/MLD snooping enabled log messages similar to the following: " <188>Jan 6 07:26:20 172.20.1.20 RfrmrHw[3.tDSrecv2]Invalid MCI -1, for asic 0" " <188>Jan 6 07:26:20 172.20.1.20 RfrmrApp[3.tDSrecv2]addPortReframing, Error: Could not convert mciIndex 8113 to UNTAGGED error status -2" then reset with a DSI exception in thread tDispatch.	8.21.01

Feature Enhancements in 8.21.02.0002

Virtual Private Ethernet Service Enhancements in 8.21.02.0002
L2VPN capability to connect Layer 2 networks transparently over a Switched or Routed IP core network using GRE or IP tunnels. With this feature, Layer 2 traffic within the switch (VLANs) can be switched into and out of the encapsulated tunnel to be transmitted across the network.

MPLS/BGP L3VPN over Native MPLS - RFC4364 Enhancements in 8.21.02.0002
Layer 3 VPN capability over MPLS transport. With this feature Layer 3 VPN traffic can be transported transparently over a native MPLS infrastructure.

Bi-directional Forwarding Detection (BFD) Enhancements in 8.21.01.0002
Support for BFD probe as a mechanism to detect a communications failure with an adjacent system forwarding plane. This version of BFD probe supports monitoring OSPF neighbors.

BGP Route-Flap Dampening Enhancements in 8.21.02.0001
Support for BGP Route-flap dampening to suppress routes that are being repeatedly advertised and withdrawn (flapping) due to mis-configuration or a badly behaving (i.e. rebooting or a link flapping) router.

PIM Dense Mode Enhancements in 8.21.02.0001
Support for PIM-DM to allow dense mode multicast distribution utilizing PIM-DM flood and prune mechanism to build source distribution trees for multicast flows.

LAG Enhancements in 8.21.02.0002
The LAG capacity has been increased to 190 LAGs in multislot S-Series chassis.

Tunnel Enhancements in 8.21.02.0002
The S-Series IP tunnel capacity has been increased to 62 tunnels.

Remote Port Mirroring Enhancements in 8.21.02.0001
Remote port mirroring is now included in the base firmware and does not need a feature license.
Remote port mirroring is now supported when user ports are used to create a "software" bonded VSB. Previously hardware VSB ports were required to bond chassis together AND use the remote port mirror feature.

CLI Enhancements in 8.21.02.0002
Show vlan portinfo CLI – CLI command has been added to display VLAN information regardless of forwarding state.
Added configuration to allow the UDP broadcast helper address to be configured to accept a classful network address.   Global configuration mode: 'ip forward-protocol allow-classful'
'show ipv6 interface' list all multicast groups the VLAN has joined.
A command to disable DHCP server logging has been added.
'show support', now includes 'show linkflap' status.
'show running bgp' has add a keyword to only display modal configuration.

Webview Enhancements in 8.21.02.0002
The left-hand WebView menu has been changed for better browser compatibility.

HOST Enhancements in 8.21.02.0002
Improved rate limiting and prioritization for Host traffic.

VLAN Enhancements in 8.21.02.0002
Support for 2 secondary VLANs per primary VLAN has been added.

Problems Corrected in 8.21.02.0002

802.1x Problems Corrected in 8.21.02.0002	Introduced in Version:
EAPOL frames may be switched when multiauth is in either forced-auth, auth-optional, or auth-required port mode.	7.00.01
802.1x global enable status may become enabled during a single board reset in multi-blade system.	8.11.01

Anti-Spoofing Problems Corrected in 8.21.02.0002	Introduced in Version:
'show config antispoof' may not display class names correctly.	8.01.01
IPv6 forwarding can be disabled on an interface that has IPv6 checkspoof configured.	7.31.02
Setting the antispoof notification interval to 0 and antispoof to enabled will consume all resources and cause the switch to be unresponsive.	8.01.01
Modifying the etsysAntiSpoofThresholdType MIB leaf to a value other than 1 (IPv4) will result in the following syslog: "Internal error: unknown remapping case (3) in make_error_pdu". The setting will not take affect as only the IPv4 Threshold Type is currently supported.	8.01.01

ARP Problems Corrected in 8.21.02.0002	Introduced in Version:
In very rare instances a module may complete it's boot process with ARP/ND entries that are present on all other blades but missing from the blade that just booted.	7.00.01
Occasionally syslog messages may appear indicating that a MAC address for an existing ARP or ND entry has changed from: ec-c1-e5-ec-c1-e5 to a different MAC address. The MAC in question is a special purpose MAC address and the message does not indicate anything has gone wrong.	8.11.01

F0615-O

ARP Problems Corrected in 8.21.02.0002	Introduced in Version:
The commands "show arp" and "show ipv6 neighbors" will print "(null)" in the port column when the MAC address for the ARP/ND entry is a static multicast MAC address.	7.00.01

ARP/ND Problems Corrected in 8.21.02.0002	Introduced in Version:
When populating the ARP/ND static ARP table (either via configuration or during the boot cycle) the router will display a message indicating the chassis is 50% full. The message implies that the dynamic ARP/ND entries are triggering the messages but the message actually refers to the static ARP/ND limit.	Unknown
Stale ARP/ND entries are not removed if a filter database entry exists for the MAC address of the ARP/ND entry.	7.71.02

Auto-Tracking Problems Corrected in 8.21.02.0002	Introduced in Version:
Help string for auto-tracking port radius-reject-profile command is incorrect.	8.01.01

BGP Problems Corrected in 8.21.02.0002	Introduced in Version:
The BGP network command for the default route (0.0.0.0/0) will not inject the route if a redistribution command exists which uses a route-map/access-list combination and this access-list does not permit the default route. To ensure that the route is injected the access-list must permit the default route.	7.20.01
BGP peering sessions may time out due to the deletion of internal connections when ACL's are deleted.	7.20.01
When redistributing loopbacks from one VRF to another on the same PE the loopbacks will not be redistributed.	7.91.01
A BGP ORF route-refresh message is not sent in a Layer-3 VPN network if the neighboring router reboots and the peering session supports graceful restart.	8.01.01
Filtering of BGP routes based on the length of the AS-PATH is not supported.	7.20.01
BGP ORF capability for multicast routes is not supported. However the CLI allows the user to enable this capability. Attempts to disable the capability will result in the following error message: "Error:Command Failed - :No such object: Setting orf admin_status".	7.20.01
The 'show ip bgp groups' command always displays the BGP peer-group address-family as IPV4/Unicast.	7.30.01
The "show ip protocols" command output is missing the following BGP related information: 1. The default values of peer based timer related variables are missing. 2. Redistribution of ISIS into BGP is not shown.	7.20.01
The 'debug ip bgp notification' log messages do not display text descriptions of the numerical error code and subcode for sent and received BGP notification messages.	7.20.01
The BGP "neighbor <IP address> clear-counters" command resets the counters the first time the command is issued for a given peer. Subsequent attempts do not clear the counters.	7.20.01
The output of the 'show ip bgp neighbors' command displays the established time in seconds instead of the "day,hours:minutes:seconds" format.	7.20.01
The BGP route-map match and set parameters will appear in the show running-config output with their default values after they are negated.	7.20.01
The BGP route-map "set extended-community ospf-route-type" command error message indicates the range of valid values is 1-7. The valid route-type values are actually 1,2,3,5, and 7.	7.20.01

BGP Problems Corrected in 8.21.02.0002	Introduced in Version:
Negating the "set community" or "set extended-community" clause from a BGP route-map yields two copies of the error message if the "action" keyword is incomplete. For example if the abbreviation for "remove" is entered for the action.	7.20.01
BGP prefix lists configured with sequence number 65535 are stored with sequence number 0 and can't be deleted.	7.20.01
The 'show ip bgp' command output does not display the value of the weight attribute.	7.20.01
The output of the 'show ip bgp neighbors' command displays the established time in seconds instead of the "day,hours:minutes:seconds" format.	7.20.01
A system reset may occur if an IPV6 BGP peer is disabled and BGP route aggregation is configured. The following error message will appear when the system resets: SMS assert in qbpmreca.c at line 233 : != *old_route 0x0x0 NULL 0x0x0.	7.30.01
The BGP neighbor activate command will not appear under the VPNv4 or VPNv6 address-family configuration if the address-family has not been enabled.	8.01.01
The 'show ip	ipv6 bgp summary' output display of the time in established state and time since last message received is in seconds instead of days,hours:minutes:seconds format.	7.20.01
The 'show running-config' and 'show config' output displays extra exclamation points in the BGP section.	7.20.01

Bonding Problems Corrected in 8.21.02.0002	Introduced in Version:
Under heavy traffic conditions, a bonded system may see "failed to send message - Buffer Full" in the message log.	7.61.02
If a VSB system is segmented, and the systems have different firmware versions, when the bond link is established between the two systems, a common image is not distributed and the system does not complete the bonding process.	Unknown
Cabling a VSB port to an ethernet port may cause modules to reset, and a message similar to" <0>Bond[14.tDispatch]getVsbInPort: learn inport:000037e3 outport:00000be7 binding failed" logged.	8.11.01
Configuration of RMON stats and history options on existing default entries will be lost on reboot.	8.11.01
Sometimes after a configure, bonding ports that are attached to a partner port are not activated for bonding. When this happens the following message is logged on the partner chassis for the partner port: "No Bond Partner found on port x.y.z, possible misconfiguration." To fix this issue disable the port for bonding then re-enable it.	7.72.01
In a bonded system, with logging for bonding set to debugging(8), messages with following format are infrequently logged: - Bonding[1]Starting inter-module communication to bonded slot <slot> - Bonding[2]Received first inter-module communication from bonded slot <slot>	7.72.01

Converged End Point (CEP) Problems Corrected in 8.21.02.0002	Introduced in Version:
Active Convergent End Point (CEP) entries will remain even if CEP is disabled globally or on a per-port basis.	6.02.04
CEP detection-id enabled/disabled state will not be displayed in 'show config' if set to disabled.	7.91.01

CFM Problems Corrected in 8.21.02.0002	Introduced in Version:
The CLI command "show cfm default-md VID <vid-number>" will display an incorrect selector type when attempting to display a single CFM Default MD.	7.91.03

CLI Problems Corrected in 8.21.02.0002	Introduced in Version:
On bonded systems while copying files from one blade to another or off the system and the bond link goes down, the master blade could reset/DSI.	7.60.01
Syslog message status is OK when setting port duplex and speed fails.	7.70.00
If the "set system lockout port" is enabled and a user fails to login via SSH the maximum allowed attempts, the user login gets locked but the port lockout fails to get locked.	7.40.01
The "show config quarantine-agent" command may leak memory.	8.01.01
The "show config dot1x" command may leak memory.	8.11.01
The "show config auto-tracking" command may leak memory.	8.01.01
Issuing a "show config" or "show config pwa" will cause a small amount of memory to leak per iteration.	8.11.01
The traceroute command only executes once inside a CLI 'loop'.	7.00.01
Memory leak executing CLI command "show snmp counters".	4.05.08

COS Problems Corrected in 8.21.02.0002	Introduced in Version:
COS ORL actions may be applied to the equivalent port on the receiving blade if the egress port is on a remote blade.	7.00.01
"processCosPortConfig" message log entry may occur if removing and showing COS configuration at the same time.	7.00.01
COS ORL rates less than or equal to 30pps may not work if the traffic is received on a remote blade.	7.00.01

DVMRP Problems Corrected in 8.21.02.0002	Introduced in Version:
DVMRP may get crashed when sending upstream prune after routes change.	7.60.01
"With DVMRP configured, the Management Module (Router) resets with a message similar to the following: "SMS assert in ntlcltim.c at line 547 : < duration -296"	7.00.01

ECMP Problems Corrected in 8.21.02.0001	Introduced in Version:
The CLI command to show the current setting of the IPv6 ECMP forwarding algorithm is missing.	7.00.01

Filter Data Base (FDB) Problems Corrected in 8.21.02.0001	Introduced in Version:
When the maximum amount of MAC entries is attempted to be set to 128K, but all blades in chassis do not have required 2G of memory, the CLI command correctly fails. However, the status returned is OK rather then ERROR.	7.91.01
When multiple static mac address (unicast and/or multicast) are configured, at boot time messages similar to: "FilterDb[2.tusrAppInit]fast_add restore (local) failed 14,60968" may be logged. There are no negative consequences, other then the messages being logged.	8.11.03

Filter Data Base (FDB) Problems Corrected in 8.21.02.0001	Introduced in Version:
If the source port of a static unicast MAC address is changed without first deleting exiting entry and recreating it, messages similar to: "FilterDb[2.tusrAppInit]restored duplicate(60126112,1 - 26-00-01-02-03-04.5 on 2" may be displayed at boot time. In addition, after reboot an entry may not restore with correct source port, or a deleted entry may re-appear.	7.00.01

GVRP Problems Corrected in 8.21.02.0002	Introduced in Version:
The ctDot1qVlanGvrpRestrictedStatus MIB object cannot be set and the "set gvrp vlan" CLI command is ignored in provider bridge mode.	7.91.01

High Availabilty Upgrade (HAU) Problems Corrected in 8.21.02.0002	Introduced in Version:
CLI does not reject out of range slot lists when configuring HAU upgrade groups. For example, "set boot high-availability group 1 1-256" should result in a CLI error, but instead the command is accepted and slots 1-N (where N is the highest slot in the system) are assigned to group "1".	7.60.01

IGMP Problems Corrected in 8.21.02.0002	Introduced in Version:
When using SSM with IGMP, SSM packet drop counters may be incorrect.	7.30.01
After a chassis segments and reforms, message of the form: "Error: Mis-Matching MCI chain data tag:1 v6:1 for MCI:131 tag:1 v6:0" are displayed and the IGMP database may become corrupted.	7.00.01
IGMP/MLD IP Multicast traffic only utilizes a single underlying physical port of a LAG in a Bonded system.	7.60.01
IGMP will not correctly update the drop counter for leaves with a bad group address.	8.11.01
While running IGMP v3 with 'include' source-list, a module crashes with a message containing to: "CIgmp::GroupTableAddPortToGroupEntry Src port mismatch".	7.30.01
IGMP ignores reports immediately after booting until the connected interfaces are populated in the Route Table.	7.31.02
IGMP running in v1 mode will drop queries for missing Router Alert.	8.11.01
When loading a configuration from a file that contains IGMP config which has "set igmp disable <x>" where x is the VLAN, any command set after this will re-enable the IGMP config for this VLAN.	7.00.01
It is possible for flows to continue egressing out a port that was removed from an IGMP static configuration.	7.91.01
IGMP and MLD frames ingressing at excessively high rates on VLANs that do not have IGMP/MLD enabled, but have a Layer 3 interface, can cause system instability including module/system resets.	7.00.01

IPv4 Frowarding Problems Corrected in 8.21.02.0001	Introduced in Version:
'ip checkspoof strict-mode' will no longer be applied to packets destined to host address configured on packet's ingress interface.	7.00.01

S-Series and S-Series Standalone Customer Release Notes

IPv6 Forwarding Problems Corrected in 8.21.02.0002	Introduced in Version:
Packets received on interfaces where IPv6 forwarding is disabled and destined to host address configured on a different interface are incorrectly delivered to the host.	7.00.01
An IPv6 address configured on a VLAN interface with a 128-bit mask is not reachable.	7.00.01
IPv4-mapped IPv6 addresses and IPv4 compatible addresses are not supported, but are accepted by the Command Line Interface. When entered an error occurs, but the address in some cases appears to be valid when in fact it is not working.	8.01.01
Route table updates may result in layer 3 VPN packets using VPN labels present before route updates occurred.	8.11.01

IPv6 Neighbor Discovery Problems Corrected in 8.21.02.0002	Introduced in Version:
Attempts to send packets from the host to a directly connected IPv6 link-local address will not work because the incorrect MAC address will be used as the destination MAC address of the destination Link-Local address.	8.11.01

CFM Problems Corrected in 8.21.02.0001	Introduced in Version:
The CLI command "show cfm default-md VID <vid-number>" will display an incorrect selector type when attempting to display a single CFM Default MD.	7.91.03

IS-IS Problems Corrected in 8.21.02.0002	Introduced in Version:
Cisco LSPs are sometimes displayed incorrectly.	8.01.01
'show isis hostname' for a level-1 router displays hostnames for level-2 router instances.	8.01.01
ISIS hostnames do not appear in LSP Summary database.	8.01.01
'show isis topology' does not display configured hostnames.	8.01.01
Within a VRF, a 32 bit summary address for ISIS is incorrectly displayed in the running config and cannot be negated.	7.73.01
When displaying ISIS LSP database information having Cisco routers, the multiple metric and IP information may be strung together.	8.01.01
When redistributing RIP into ISIS, the wrong metric is displayed.	8.01.01
When displaying the ISIS LSP database, the wide metric values shown are incorrect.	8.01.01
ISIS LSP database wide metrics are displayed incorrectly.	8.01.01
Deleting 'isis lsp-gen-interval' will result in the default value (in seconds) showing up in show running config.	8.11.01
'show running isis' may display additional blank line between authentication statements.	8.11.01
ISIS LSP databases with metrics associated with multiple IPs get displayed incorrectly.	8.01.01
When the connection between ISIS neighbors is tunneled and over-subscribed, the neighbors periodically lose connectivity.	8.01.01
ISIS database display shows an incorrect format after metric IS.	8.01.01
After 'no ip router isis' is done on an interface, hellos may continue to be sent.	8.01.01
When redistributing into ISIS the show running command displays command twice.	8.11.01

F0615-O

Jumbo Problems Corrected in 8.21.02.0002	Introduced in Version:
Invalid sized non-tagged packets of size 1519 to 1522 bytes and tagged packets of size 1523 to 1526 received on non-jumbo enabled ports are correctly dropped. However, the SA MAC is incorrectly learned in MAC table.	7.00.01
Invalid sized non-tagged packets of size 10240 to 10243 bytes and tagged packets of size 10244 to 10247 received on jumbo enabled ports are correctly dropped. However, the SA MAC is incorrectly learned in MAC table.	7.00.01
For some flows that require reframing, if any one of the first few packets in flow are jumbo sized, those packets could be dropped (and not forwarded).	7.60.01

L3VPN Problems Corrected in 8.21.02.0002	Introduced in Version:
Using L3VPNs when BGP is attached to a route-reflector client will result in routes not propagated to the VRFs.	8.01.01
Creating domain-ID (primary or secondary) with an invalid <6 octet domain id> creates on with FF instead of returning an error message.	8.11.01

LACP Problems Corrected in 8.21.02.0002	Introduced in Version:
LACP marker response not within frame rate limitation constraint for slow protocols.	1.07.19
In some instances, LACP is not setting, collecting, and distributing bits to false after a partner PDU change, resulting in the port not leaving the LAG as it should.	1.07.19
A set of a lag port attribute may fail without a message at the console.	1.07.19
Distribution of traffic over the ports in a LAG could vary over 10% port-to-port from a uniform distribution when an odd number of ports are in the LAG	7.30.01
In rare instances, a port that joins a LAG briefly then stays down/is removed from the lag may still be considered an available egress port for a few percent of LAG traffic which would be undelivered. A subsequent change of state of any of the ports in the LAG or the addition/removal of a port in the LAG will clear the condition.	5.01.58

LLDP Problems Corrected in 8.21.02.0002	Introduced in Version:
MIB lldpStatsRemTablesAgeouts is not incremented when a neighbor ages out.	7.00.01
lldpStatsRxPortTLVsDiscardedTotal may not increment for non-support LLDP TLVs.	7.00.01
LLDP Management Address TLV has incorrect interface index.	7.00.01
LLDP Link Aggregation TLV was using a format that was deprecated in IEEE 802.1AB-2009.	7.00.01
The PoE TLV in a transmitted LLDP packet correctly shows a TLV length of 12, but the extended information shows an incorrect Type/Source/Priority (TSP) field, PD requested power value, and PSE allocated power value.	8.11.01
Occasionally while under heavy processing load, LLDP may cause the system to crash.	7.62.00

LSNAT Problems Corrected in 8.21.02.0002	Introduced in Version:
In a previous release access to a VIP server from a VRF via a route was not allowed without the 'all_vrfs' configuration command option defined on the VIP server. Access to a VIP from a VRF via a route leak is now allowed.	7.00.01
In a previous release after modifying the global NAT SLB or TWCB binding limit, it may cause runtime issues while processing bindings.	6.12.05

LSNAT Problems Corrected in 8.21.02.0002	Introduced in Version:
"show config", "show running", and "show ip slb info" will not display the "real server access client" configuration lines.	6.12.08

MAC AUTH Problems Corrected in 8.21.02.0002	Introduced in Version:
Setting the authallocated macauthentication field (''set macauthentication authallocated <port string>'') to a value of 0 does not correctly result in an outputted error although the value is not set.	5.01.58

Mirroring Problems Corrected in 8.21.02.0002	Introduced in Version:
Infrequently, a chassis module with port mirrors configured resets. On this failure a message similar to "PortMirr[3.tDSrecv1]processMirrorDestination(1,61013): index mismatch detected: smon=2, mirror=5" is logged.	7.41.02
Port mirroring may reset and log a message similar to "PortMirr[12.tDSrecv1]setMirrorIndex(72028,162024): invalid mirror index transition 2->1".	7.00.01
The "clear port mirroring orl" command does not disable mirror outbound rate-limiting.	8.11.01
When mirroring, the physical loopback port does not go down when the tunnel goes operationally down.	8.11.03
The "clear port mirroring" and "set port mirrorring [enable|disable]" commands do not set the lower numbered destination ports if the destination port-string is in descending order (i.e. tg.4.3;tg.4.2). These commands function properly when the destination port-string lists the lowered ports first.	7.91.01
In a software VSB chassis, if mirrors are active, very infrequently a blade may reset, leaving a message similar to: "<1>DistServ[12.tDsBrdOk]serverWatchDog.1(Config), client 88(mirrorMgr) in recv for 6721 tics".	7.60.01

MSDP Problems Corrected in 8.21.02.0002	Introduced in Version:
Multiple removing MSDP by 'no ip msdp originator-id' causes a crash similar to: 'SMS assert in qptujms.c at line 257 : == msdp_mj_cb-> j_mj_cb.mj_row_data.oper_status 8 AMB_NPG_OPER_STATUS_GOING_UP 3 '.	8.02.02

MULTI AUTH Problems Corrected in 8.21.02.0002	Introduced in Version:
Executing 'show multiauth session port <port-string>' might result in an error.	7.30.01
The 'show multiauth station port' command displays multiple entries for each provisioning agent type.	5.01.58
If 'multiauth sessions-unique-per-port' is disabled and CEP multiauth sessions are moving from one port to another RADIUS accounting data may be output inconsistently for that session.	8.11.01
Multiauth sessions that port roam may not session timeout at the expected time.	8.11.01
Quarantine Agent Multiauth Sessions may not idle timeout as expected when port roaming.	8.11.01
Clearing multiauthentication stations using the etsysMultiAuthStationClearUsers MIB leaf may cause the multiauthentication software to treat the clearing as a failure for both logging and trap purposes.	7.72.01

MULTI AUTH Problems Corrected in 8.21.02.0002	Introduced in Version:
When multiauthentication traps for authentication success, authentication rejection or port termaination are enabled and are being sent they result in duplicate notice level log events that indicate the same or similar information.	7.00.01
Networks utilizing multiauth session or idle timeouts greater than 65535 may have sessions that timeout inaccurately.	6.11.01

NAT Problems Corrected in 8.21.02.0001	Introduced in Version:
No counter for NAT packets and NAT drop packets are included in "debug packet show-statistics".	1.07.19
It is possible that once the "ip nat log translations" config has been entered that it will remain persistent even when a "no ip nat log translations" command is entered.	6.00.02
It is possible when upgrading from 08.02.xx or downgrading to 08.02.xx that the following NAT config may be lost: ip nat ftp-control-port ip nat log translation ip nat inspect dns ip nat translation max-entries ip nat translation 'timeouts'	8.11.01

NETFLOW Problems Corrected in 8.21.02.0002	Introduced in Version:
When NetFlow is enabled, very infrequently, an error message similar to: "<3>netflow[4.tNetflow]netflow_record_processing_task - unexpected error taking semaphore"may be displayed. When that message is logged, a single frame, which can consist of any where from 1 to 30 NetFlow records, is dropped and will not be delivered to NetFlow collectors.	8.01.01
When NetFlow export-data higher-layer is enabled, messages similar to: "PiMgr[7.tMcnxPer]generateIfIndex():retval=7;mediaType(0);mediaPos(8)" may be displayed. For each message generated, a NetFlow record with an invalid destination interface will be sent.	8.01.01
When the 'set default-nexthop[-v6]' option is used in route-maps and Netflow is in use, a reset may occur when route updates are being processed.	7.20.01
Very infrequently, when NetFlow export data higher layer is enabled, messages similar to: "PiMgr[7.tMcnxPer]generateIfIndex():retval=0;owner(1);mediaType(7);mediaPos(0)" may be logged. For every message logged, a NetFlow record would be generating with invalid source and/or destination interfaces.	8.01.01

OAM Problems Corrected in 8.21.02.0002	Introduced in Version:
Disabling OAM on a port does not clear the OAM or ULD operstatuscause	7.30.01

OSPF Problems Corrected in 8.21.02.0002	Introduced in Version:
When running OSPF, and using the passive-interface default command, an assert could occur in thread tRtrPtcls with the following log, "SMS assert in qopmmim5.c at line 879 : (null) AVLL_IN_TREE(if_cb->active_if_tree_node) 0 (null) 0".	7.00.01

S-Series and S-Series Standalone Customer Release Notes

OSPF Problems Corrected in 8.21.02.0002	Introduced in Version:
A tunnel interface running OSPF will default to network type POINT_TO_POINT. If it is explicitly configured as POINT-TO-POINT and then removed, it defaults to BROADCAST instead.	7.41.02
If an OSPF interface running over a tunnel is explicitly configured as point-to-point this is displayed in the config even though it is the default.	7.41.02
OSPFv2 will accept the configuration of an invalid nssa-range and display it incorrectly.	7.00.01
Configuring an OSPF cost metric outside the range results in an unclear message error.	8.11.01
When issuing a "clear ip ospf process" and multiple OSPF processes exist, the ambiguous message "Resetting the OSPF process" is seen multiple times.	7.00.01
When running OSPFv2 or v3 with auto-cost reference bandwidth and tracked objects, it is possible with multiple cost changes to have the router LSA not reflect the cost seen on the interface.	8.11.01
If OSPF logging is enabled, and multiple OSPF processes are in use, an abundance of messages are seen about each process when the reference bandwidth is changed in a single process.	8.01.01
sham is spelled "shaml" on both OSPFv2 and OSPFv3 debug.	8.11.01
The wrong dead interval range was displayed in the help section of the CLI for sham links.	8.11.01
'show running config' for the sham link authentication would not be displayed.	8.11.01
When looking at the debug syslog, sham-link interval mismatch messages do not decode ifindex to text strings.	8.11.01
The 'show ip ospf interface vlan.0.x' command may show additional space at the end if multiple addresses are configured on that interface that are not running OSPF.	8.01.01
On bonded systems while copying files from one blade to another or off the system and the bond link goes down, the master blade could reset/DSI.	8.11.04
When using a route map when entering the 'redistribute bgp global' command, the route type will be changed to E2 and not use the correct domain id type.	8.11.01
OSPF has no warning message when the calculated cost metric for an interface due to an auto-cost reference bandwidth change results in a too large metric.	8.11.01

OSPFv3 Problems Corrected in 8.21.02.0002	Introduced in Version:
If an OSPF vlan interface is configured to be POINT_TO_POINT, then the configuration is removed with "no ip ospf network point-to-point", the interface network type is POINT_TO_POINT instead of reverting to the default type of BROADCAST.	7.41.02
If an OSPF auto-cost reference bandwidth is configured that causes the interface to calculate a cost greater than the maximum, the cost remained based on the previous auto-cost reference bandwidth value.	8.11.01

PIM-SM Problems Corrected in 8.21.02.0002	Introduced in Version:
IGMP/Multicast in a bonded chassis appears to take longer for some events than an identically configured single chassis.	7.61.02
"ip pim multipath" configuration is not cleared after executing a "clear router vrf <vrfname>."	8.01.01
PIM configuration for ipv4 is accepted after removal of L3 license.	7.00.01

S-Series and S-Series Standalone Customer Release Notes

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

Platform Problems Corrected in 8.21.02.0002	Introduced in Version:
Reading a file from another blade (Ex: 'show file' or 'configure') could cause a DSI/reset, usually if remote file is being updated, or remote connection goes away (other blade resets or bonding goes away).	7.00.01
Running "chkdsk repair" could cause a reset. This command is only available from debug, or during boot if filesystem corruption is detected.	7.00.01
Performing a configuration operation via the command line interface may result in the old configuration remaining due to file access errors.	7.70.01
The following syslog messages may be seen on bonded systems when the remote blade is under heavy load or the remote system is resetting/booting: 'NonVol[1.tNvBulk]nonvol_copy: Copying of redundant store will need to retry (FIOSYNC of outFd failed errno(errno = 0x300005))' These are log level 5(NOTICE) messages and usually only an issue if persistent.	7.30.01
Setting port speed on 1G fiber port to a speed the SFP does not support, and disabling auto negotiation causes the port to go down and stay down.	7.91.01
Releases before 8.20.01 may fail at initialization time with large configurations causing continuous resets caused by denoted rdyToSwitch fault.	8.11.01
"set port mdix" CLI commands may display "Error: failed to set mdix configuration of swap on port tg.5.2." on RJ45 ports that don't support mdix commands.	8.01.01
Unsupported Option Module will halt the board and not allow software upgrade until removed.	8.11.01
System may log a message similar to <1>DistServ[1.tDsBrdOk]serverWatchDog.5(Host), client 26(Emanate) in recv for 6446 tics and then reset.	7.60.01
If a SFP+ Direct Attach cable assembly is used to connect two 1000Base-X ports (an unsupported configuration), the message "Incompatible pluggable module" will be logged on behalf of each port, but physical link will not be forced down.	7.00.01
10GB-LRW-SFPP Laser Wire transceivers not acquired through Extreme Networks that do not have a "-EN" part number will display "sfpDataAccess: CI2CBus::Access() failed to write transceiver data for slave 0xa2 on NIM x port y" message.	8.11.01
SFP pluggable failure messages are not as user friendly as they should be.	8.01.01
Port advertisement settings are not persistent when auto negotiation is disabled.	7.91.01
100M SFP inserted into 1G port shows default speed and negotiation disable in 'show config'.	7.00.01

F0615-O

Platform Problems Corrected in 8.21.02.0002	Introduced in Version:
"Core files might not be generated for defects which result in stack corruption whenever a DSI or ISI exception occurs, the system logs the original exception to NONVOL then attempts to generate a core file (i.e., /slot<x>/cores/<xxxx>.core.gz) which will include a stack trace of the offending task. If the stack is corrupted, then the process of printing the stack trace to the core file will itself create a new DSI. This new DSI prevents core file generation from completing and being saved to disk."	7.00.01
'show port status' is missing speed and duplex for 10G copper ports without a link.	8.01.01
No message is logged indicating a reason for board shutdown due to over temperature on S-Series.	7.00.01
"At boot a board could get into reset loop with the following syslog output: 'NonVol[1.tusrAppInit]Nonvol reached max fileIdx 4080, storeNum 1, major 1'. Sets will be dropped until space if freed."	7.30.01
Unsupported speed of 100M is allowed to be set for 1G SFP.	7.91.01

Policy Problems Corrected in 8.21.02.0002	Introduced in Version:
Rules to drop GVRP or MVRP packets are ignored.	7.00.01
IP addresses in "set policy rule" would be treated as octal if a preceding "0" is present.	6.00.02
The CLI command 'show vlanauthorization' will not display the vlan authorization status of all ports in the system.	6.00.02
Multiauth failure traps may be output for port roaming sessions that roam to ports with insufficient per port multiauth number of users to support the new session.	7.72.01

PWA Problems Corrected in 8.21.02.0002	Introduced in Version:
PWA set portcontrol CLI commands do not output an error if wildcarding is used for a port string which contains no valid ports.	5.42.04

RADIUS Problems Corrected in 8.21.02.0002	Introduced in Version:
8.11 RADIUS Enhancement documentation does not clearly indicate that round-robin handling occurs on a per blade basis.	8.11.01
RADIUS authentication server max-sessions configuration is not output as part of "show config" or "show config all" commands.	8.11.01
RADIUS authentication server realm is not displayed as part of the "show config all" command if it is set to the default of any.	8.11.01
RADIUS authentication, authorization, and accounting server configuration may be lost upon upgrade from any release prior to 7.40 to any release post 8.02.	8.11.01

RADIUS-SNOOPING Problems Corrected in 8.21.02.0002	Introduced in Version:
If multiple CLI sessions are concurrently accessing RADIUS Snooping information the system may crash or provide inaccurate results.	6.11.01
Show config of the RADIUS Snooping, auto tracking and quarantine provisioning agents displays default port parameters whenever at least one port field is set to a non-default setting.	6.11.01
Setting radius-snooping port configuration for unsupported ports may not result in proper CLI error messaging.	6.11.01

RMON Problems Corrected in 8.21.02.0002	Introduced in Version:
Heavy use of RMON alarm and RMON event may result in a system reset and the log message "memPartFree: invalid block 0x3257c710 in partition 0x59a0a78 <memSysPartition>".	5.01.58
Configuration of RMON etherStats may return an incorrect value upon using an invalid index as input.	1.07.19
MIB leaf historyControlStatus can be set directly to under creation with non-existent index.	1.07.19
Configuration of an RMON function with an out of range index does not always return error.	5.01.58
"show rmon alarm" will show a negative value for alarm variables that are counters(unsigned), specifically for values between 2147483647 and 4294967294(rollover).	5.01.58

Routing Problems Corrected in 8.21.02.0002	Introduced in Version:
Negating interface checkspoof setting without a keyword returns an error when checkspoof loose-mode is configured.	7.00.01

SMON Problems Corrected in 8.21.02.0002	Introduced in Version:
Infrequently, a chassis module with port mirrors configured resets. On this failure a message similar to "setMirrorIndex(103201,122024): invalid mirror index transition 2->1" is logged.	7.00.01
Polling the SMON Vlan Statistics MIB smonVlanidStatsCreateTime object returns an error.	7.91.01

SNMP Problems Corrected in 8.21.02.0002	Introduced in Version:
For snmp view configuration, snmp view mask values entered as single byte hexidecimal values (without a colon) that are less than 0x7f (and are printable ascii characters) appear as printable ascii characters instead of hexidecimal values, and result in missing configuration lines.	4.00.50

SYSLOG Problems Corrected in 8.21.02.0002	Introduced in Version:
Messages that should be logged to the console as part of the shut down process are not seen.	7.80.01
client 22(Syslog) not ready in <number_of_tics> tics' message seen on console.	5.51.xx

Tracked Objects Problems Corrected in 8.21.02.0002	Introduced in Version:
Taking a tracked object out of service ('no inservice' sub-mode command) while a state change is in progress does not remove the state change action from the delay queue. If the tracked object is put back into service ('inservice' sub-mode command) prior to the state change action expiring from the delay queue, the new state change action is not queued. The new state change action is triggered when the old state change action expires from the delay queue.	7.60.01

Tunneling Problems Corrected in 8.21.02.0002	Introduced in Version:
When either a tunnel probe or the GRE keepalive is down, the tunnel is held down. This has been changed. If either the probe or keepalive is up or neither are configured, then the tunnel will be operationally up given other conditions are correct.	8.11.01

F0615-O

Tunneling Problems Corrected in 8.21.02.0002	Introduced in Version:
The range check on a tunnel keepalive period prevented the user from entering anything larger than 255.	8.11.01
A GRE keepalive nested within another GRE tunnel would be dropped.	8.11.01
HW connections may be incorrectly installed to drop virtual private port flows that include nested GRE packets with the protocol=0x6558	8.11.01
IPv6 encapsulated flow of an IPv4 flow was using the IP version from the Transformation. It now uses the IP Version from the Ingress Flow.	8.01.01

VRF Problems Corrected in 8.21.02.0002	Introduced in Version:
When using the maximum length VRF name, it insists on a context, but when one is specified, it takes the VRF name and discards the extra characters.	7.62.02
From device command line, a ping to device's address configured in another VRF fails even though VRF route leaking is provided by static routes.	8.11.01

VRRP Problems Corrected in 8.21.02.0002	Introduced in Version:
After repeated enable/disable of VRRP accept-mode, the system may not reply to ICMP echo requests to the VRRP Virtual IP address.	8.11.01
Host routes added by host mobility may age out during first age pass after they are added.	8.11.01

Feature Enhancements in 8.11.05.0007

Transceiver Enhancements in 8.11.05.0007
CWDM support: 10GB-LR271-SFPP - 10Gb, CWDM SM, 1271 nm, 10 km, LC SFP+ 10GB-LR291-SFPP - 10Gb, CWDM SM, 1291 nm, 10 km, LC SFP+ 10GB-LR311-SFPP - 10Gb, CWDM SM, 1311 nm, 10 km, LC SFP+ 10GB-LR331-SFPP - 10Gb, CWDM SM, 1331 nm, 10 km, LC SFP+
Additional DWDM support: 10GB-ER21-SFPP - 10GB-ER, DWDM CH21 SFP+ 10GB-ER24-SFPP - 10GB-ER, DWDM CH24 SFP+ 10GB-ER31-SFPP - 10GB-ER, DWDM CH31 SFP+ 10GB-ER33-SFPP - 10GB-ER, DWDM CH33 SFP+

Problems Corrected in 8.11.05.0007

ACLs Problems Corrected in 8.11.05.0007	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

F0615-O

S-Series and S-Series Standalone Customer Release Notes

Auto-config Problems Corrected in 8.11.05.0007	Introduced in Version:
On a chassis with 6 or more filled slots running with no/default configuration, if you do a "set configuration" command, during the reset you may see the following messages in the log: "<163>Sep 13 14:12:03 0.0.0.0 autoConfig[4.tDSrecv7]setConfigAtDefaultsBySlot: Unable to send nonvol change to msgQ inslot(6) value(0)" "<163>Sep 13 14:12:03 0.0.0.0 autoConfig[4.tIpAddrCb]autoConfig_IfEventCallback: Unable to send IF_DELETED-event(6), id(1) myid(0) to msgQ"	8.11.01

IGMP Problems Corrected in 8.11.05.0007	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 8.11.05.0007	Introduced in Version:
When inserting a new blade into the system the new blade may end up with an interface in the "stalled" state which indicates that the IPv6 addresses have not passed Duplicate Address Detection. The interface will not forward IPv6 packets until the interface is bounced (the operational status goes down then back up).	7.41.02

LLDP Problems Corrected in 8.11.05.0007	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 8.11.05.0007	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 8.11.05.0007	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 8.11.05.0007	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]generateIfIndex() :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 8.11.05.0007	Introduced in Version:
Under rare circumstances, the "ctAliasControlTable" will not return all valid entries.	7.91.01

Node Alias Problems Corrected in 8.11.05.0007	Introduced in Version:
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

OSPF Problems Corrected in 8.11.05.0007	Introduced in Version:
An assertion failure and reset occurs and is recorded in message log as; "SMS assert in qoamlsts.c at line 1218"	7.00.01
When running OSPFv2 and flapping the passive value on an interface, an assert can occur in thread tRtrPtcls with the following message; "SMS assert in qopmmim5.c at line 879 : (null) AVLL_IN_TREE(if_cb->active_if_tree_node) 0 (null) 0 "	8.11.01
When running OSPF a DSI can occur in thread tRtrPtcls, message displayed is: "SMS assert in ntlavll.c at line 644 : != AVL3_IN_TREE(*node) 0 0 0"	8.11.01

PWA Problems Corrected in 8.11.05.0007	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 8.11.05.0007	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 8.11.05.0007	Introduced in Version:
Precision Time Protocol (PTPv1) UDP broadcast port 139, when being forwarded through switch, may not function reliably.	1.07.19

Feature Enhancements in 8.11.04.0006

Tranceiver Enhancements in 8.11.04.0006
Auto negotiation support for 1Gb SFP GBICs installed in SFP+ sockets.

Problems Corrected in 8.11.04.0006

CLI Problems Corrected in 8.11.04.0006	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 8.11.04.0006	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]CIgmpEtsc::DistGrpTblRecvDistributedAdd Recv base index out of range baseidx:xxx flowIdx:xxx	8.11.01

L3 VPN Problems Corrected in 8.11.04.0006	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 8.11.04.0006	Introduced in Version:
Querying the ctAliasInterface table may not return all entries on a given interface.	8.11.01
Querying the ctAliasInterface table may not return all entries on a given interface in multislot systems.	8.11.01

NONVOL Problems Corrected in 8.11.04.0006	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.tusrAppInit]validate_files: Unknown record type;store=1,offset=4105,file=0.80, type=0,rawMaj=0,rawMin=0,rawLen=0 <163>Sep 19 14:46:02 0.0.0.0 NonVol[1.tusrAppInit]validate_files: file=1/0.80 rewinding over incomplete record. Truncating to size 4105 <163>Sep 19 14:46:02 0.0.0.0 NonVol[1.tusrAppInit]nvFilePtrMgr::fFlush(5) fflush(0x72b03b0) retval=-1, errno=9 Configuration could have been lost due to file corruption and should be verified.	3.00.33
The nonvol cleanup task can write incomplete files to the nonvol store that will not be detected until a reboot: NonVol[1.tusrAppInit]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 8.11.04.0006	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
"Some devices may reset after logging a message similar to the one listed below because memory requires an adjustment to the 1.0V power controller. <163>Apr 7 15:05:51 0.0.0.0 Dune[5.tRootTask]PETRA[0] failed to initialize DRAM (0x65535). "	8.01.01
Some devices may reset after logging a message similar to the one listed below because memory requires an adjustment to the 1.0V power controller. <163>Mar 27 03:06:57 192.168.100.18 Dune[2.dTcmTask]Petra[0] Received Interrupt PB_IPT_CRC_ERR_PKT instance 0, count 1, value= 0x1	8.01.01
System logs the message "bcmStrat[1.tNimIntr]MEM_FAIL_INT_STAT=0x00200000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000" and resets.	7.70.01
System logs the message "bcmStrat[1.tNimIntr]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000001, IP3_INTR_STATUS=0x00000000" and resets.	7.70.01
System logs the message "bcmStrat[2.tNimIntr]MEM_FAIL_INT_STAT=0x00040000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000" and resets.	7.70.01
System logs the message "bcmStrat[1.tNimIntr]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000002" and resets.	7.70.01

PoE Problems Corrected in 8.11.04.0006	Introduced in Version:
'set inlinepower management class' configuration might not be persistent.	8.01.01

RADIUS Problems Corrected in 8.11.04.0006	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 8.11.04.0006	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 8.11.04.0006	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

TWCB Problems Corrected in 8.11.04.0006	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 8.11.04.0006	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 8.11.03.0006

Automated Deployment Feature Enhancements in 8.11.03.0006
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 8.11.03.0006

ACL Problems Corrected in 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	Introduced in Version:
Displaying FIB history via debug CLI may block BGP from maintaining connection to peers.	7.00.01
"Negating a BGP route-map ""match extended-community as-route-target"" command may result in a system reset. The following error message will appear at the CLI: SMS assert in qbmlrex3.c at line 414 : >= string_len 0 (2 * QB_LEN_EXT_COMMUNITY) 16 "	8.01.01
A system reset may occur when running BGP with the full Internet routing table and resetting or changing the export policy of a neighboring router. The following error message will appear: SMS assert in qbdcnhr.c at line 959 : || (old_loc_route == ari_route->loc_route) 0 (QBRA_CHECK_FLAG(ari_route->loc_route->flags, QBRA_LOC_FLAG_REMOVAL_DONE)) 0	8.11.01
Multiprotocol BGP peering with third party products may not establish if received update messages contain out of order path attributes such that AS-PATH is the last attribute.	7.30.01
A system reset may occur if peering is attempted with a router supporting multisession BGP. The reset will occur on receipt of a Notification message with the error code of 2 (Open message error) and subcode 8 (grouping conflict). The following error message will appear: SMS assert in qbnmpd.c at line 141 : (null) INVALID BRANCH 0 (null) 0	8.11.01

Bonding Problems Corrected in 8.11.03.0006	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
Bonded chassis may segment after a slot reset.	7.70.00
Modules in a hardware bonded chassis may reset when a VSB port is connected to a front panel port. A message similar to "<0>Bond[13.tDispatch]getVsbInPort: learn inport:000033eb outport:00002bef binding failed (0x00c77d1c 0x00574058 0x015830e4 0x015756f4 0x0157ebec 0x01830ea0 0xeeeeeeee)" is logged on this error.	8.11.01
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 8.11.03.0006	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 8.11.03.0006	Introduced in Version:
DHCPv6 server responds to DHCPv6 request on interfaces that do not have 'ipv6 dhcp server' configured.	8.11.01

FDB Problems Corrected in 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	Introduced in Version:
Occasionally running the show neighbor command will display a neighbor multiple times.	7.91.01

MAC Authentication Problems Corrected in 8.11.03.0006	Introduced in Version:
MAC-Authenication auth-mode may be set to radius-username when upgrading from older firmware versions.	8.11.01

Multi User Authentication Problems Corrected in 8.11.03.0006	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 8.11.03.0006	Introduced in Version:
It is possible for a NAT Static reserved binding to age out.	8.11.01

NAT Problems Corrected in 8.11.03.0006	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 8.11.03.0006	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 Alais Problems Corrected in 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	Introduced in Version:
"System logs the message ""bcmStrat[2.tNimIntr]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000000, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000010, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000"" and resets.	7.70.01
System logs the message "bcmStrat[5.tNimIntr]MEM_FAIL_INT_STAT=0x00000000, EP_INTR_STATUS=0x00000080, IP0_INTR_STATUS=0x00000000, IP1_INTR_STATUS=0x00000000, IP2_INTR_STATUS=0x00000000, IP3_INTR_STATUS=0x00000000" and resets.	7.70.01
Some devices may reset after logging a message similar to the one listed below: <163>Mar 27 03:06:57 192.168.100.18 Dune[2.dTcmTask]Petra[0] Received Interrupt PB_IPT_CRC_ERR_PKT instance 0, count 1, value= 0x1	8.01.01
Some devices may reset after logging a message similar to the one listed below because memory requires an improved initialization sequence.<163>Apr 7 15:05:51 0.0.0.0 Dune[5.tRootTask]PETRA[0] failed to initialize DRAM (0x65535).	8.01.01

S-Series and S-Series Standalone Customer Release Notes

Platform Problems Corrected in 8.11.03.0006	Introduced in Version:
S chassis reporting an incorrect ambient temperature of -3C.	7.60.01
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
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 8.11.03.0006	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 8.11.03.0006	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

SSH Problems Corrected in 8.11.03.0006	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 8.11.03.0006	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 8.11.03.0006	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

VRRP Problems Corrected in 8.11.03.0006	Introduced in Version:
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 8.11.02.0002

Upgrade Problems Corrected in 8.11.02.0002	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 8.11.01.0015

Application Policy Feature Enhancement in 8.11.01.0015
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 8.11.01.0015
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 8.11.01.0015
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 8.11.01.0015
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 8.11.01.0015
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 8.11.01.0015
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 8.11.01.0015
SSH enhancement to support Public Key Authentication as an additional client authentication method.

RMON Stats and History Feature Enhancement in 8.11.01.0015
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 8.11.01.0015
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 8.11.01.0015
Allows the MAC Authentication password to use the configured password or the username as password.

IPv6 DHCP Server Feature Enhancement in 8.11.01.0015
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 8.11.01.0015
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 8.11.01.0015
Enhancement to support configuring OSPF reference bandwidth, allowing for more granular auto-costing of OSPF links.

OSPF RFC 4577 Support Feature Enhancement in 8.11.01.0015
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 8.11.01.0015
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 8.01.01.0012
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

S-Series and S-Series Standalone Customer Release Notes

IP Service Level Agreements Feature Enhancements in 8.02.01.0012

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

Tracked Objects Feature Enhancements in 8.02.01.0012

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

L3VPN over GRE Feature Enhancements in 8.02.01.0012

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

User Tracking and Control Feature Enhancements in 8.02.01.0012

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

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

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

Anti-Spoofing Suite Feature Enhancements in 8.02.01.0012

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

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

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

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

DHCP Feature Enhancements in 8.02.01.0012

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

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

Port Mirror Feature Enhancements in 8.02.01.0012

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

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

Network Address Translation Feature Enhancements in 8.02.01.0012

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

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

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

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

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

Multicast Feature Enhancements in 8.02.01.0012

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

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

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

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

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

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

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

MVRP - IEEE 802.1ak Feature Enhancements in 8.02.01.0012

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

F0615-O

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 liklihood 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:**
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. *For example:* 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	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 qodmnssa.c" when user adds and all zeros NSSA route	7.00.01
When gracefully restarting a Designated Router, OSPF may not send hellos with itself as the DR.	8.01.01
A blade may reset repeatedly logging a DSI exception for thread 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

F0615-O

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 "stpcombatible" and entering or leaving a topology change condition.	7.00.01

SYSLOG Problems Corrected in 8.02.01.0012	Introduced in Version:
Messages sent to syslog servers could contain unprintable control characters in the middle of the messages.	7.11.01

VLAN Problems Corrected in 8.02.01.0012	Introduced in Version:
A VLAN interface based mirror will continue to mirror traffic after the VLAN interface is removed from the config with the clear command.	1.07.19

VRF Problems Corrected in 8.02.01.0012	Introduced in Version:
When doing a fail over, then a show running config, some limit commands will show up even though they were not set.	7.70.01

KNOWN RESTRICTIONS AND LIMITATION:

It is not possible to mix S130/S150/S155 fabrics and the S180 fabric class in the same chassis.
S140 and S180 class modules require the use of S180 class fabrics when used in the S4/S6 and S8 chassis. S150/S130 class I/O can be used with any fabrics class.
MPLS/LPD/L3VPNs will not function over an IPv6 core.   This will be added in a later release.
When upgrading to 8.11.05, it is possible that some IPv6 interface configuration will be lost. This has been observed in bonded systems when doing a HAU upgrade.
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 VSB several features are resized or restricted:
  LAG capacities are reduce to 126 for chassis, 61 for SSAs,
  GRE Tunnels 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-SFPP (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 requires the SSA-AC-PS-1000W power supplies. (The SSA-AC-PS-625W must not be used in the S1-Chassis.) The Fabrics/Option Modules and optics along with the Fans can exceed the power available in the 625W supply during the startup and when the fans operate at full speed.
The "script" command should not be used. Its use will result in memory corruption and reset or other undesired behavior.
When an SFP (1G) module is inserted or removed from an SFP+ (10G capable) port, all ports on the associated MAC chip are reset. This results in a momentary loss of link and traffic on affected ports and forces topology protocols to process a link bounce. On SSA all 10G ports are in the same group.   All ports on a 10G Option Module are grouped together. For S blades shipping with factory configured ports the groups are: tg.x.1-4, tg.x.5-8, tg.x.9-12, tg.x.13-16.
MGBIC-100BT doesn't support automatic detection of MDIX (Medium Dependent Interface Crossover) or Auto-negotiation.
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

RFC No.	Title
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
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

F0615-O

RFC No.	Title
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
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

RFC No.	Title
RFC2697	A Single Rate Three Color Marker
RFC2710	Multicast Listener Discovery (MLD) for IPv6
RFC2711	IPv6 Router Alert Option
RFC2715	Interop Rules for MCAST Routing Protocols
RFC2740	OSPF for IPv6
RFC2763	Dynamic Hostname Exchange Mechanism for IS-IS
RFC2784	GRE
RFC2787	VRRP MIB
RFC2796	BGP Route Reflection
RFC2819	RMON MIB
RFC2827	Network Ingress Filtering
RFC2858	Multiprotocol Extensions for BGP-4
RFC2863	IF-MIB
RFC2864	IF-INVERTED-STACK-MIB
RFC2865	RADIUS Authentication
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

RFC No.	Title
RFC3415	SNMP-VIEW-BASED-ACM-MIB
RFC3417	SNMPv2-TM
RFC3418	SNMPv2 MIB
RFC3446	Anycast RP mechanism using PIM and MSDP
RFC3484	Default Address Selection for IPv6
RFC3493	Basic Socket Interface Extensions for IPv6
RFC3509	Alternative Implementations of OSPF ABRs
RFC3513	RFC 3513 IPv6 Addressing Architecture
RFC3542	Advanced Sockets API for IPv6
RFC3562	Key Mgt Considerations for TCP MD5 Signature Opt
RFC3567	IS-IS Cryptographic Authentication
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 signalling 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 No.	Title
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
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

RFC No.	Title
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
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

F0615-O

RFC No.	Title
RFC5602	Pseudowire (PW) over MPLS PSN MIB
RFC5643	OSPFv3 MIB
RFC5798	Virtual Router Redundancy Protocol (VRRP) V3
RFC6104	Rogue IPv6 RA Problem Statement
RFC6105	IPv6 Router Advertisement Guard
RFC6106	IPv6 RA Options for DNS Configuration
RFC6164	Using 127-Bit IPv6 Prefixes on Inter-Router Links
RFC6296	IPv6-to-IPv6 Network Prefix Translation
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-vrrp-unified-spec-03.txt

EXTREME NETWORKS PRIVATE ENTERPRISE MIB SUPPORT:

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

F0615-O

Title	Title	Title
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
	ENTERASYS-PTOPO-MIB-EXT-MIB	RSTP-MIB
ENTERASYS-ETH-OAM-EXT-MIB	ENTERASYS-PWA-MIB	U-BRIDGE-MIB
ENTERASYS-IEEE8021-BRIDGE-MIB-EXT-MIB	ENTERASYS-RESOURCE-UTILIZATION-MIB	USM-TARGET-TAG-MIB
ENTERASYS-IEEE8021-SPANNING-TREE-MIB-EXT-MIB	ENTERASYS-RIPv2-EXT-MIB	ENTERASYS-TWCB-MIB
ENTERASYS-IEEE8023-LAG-MIB-EXT-MIB	ENTERASYS-RMON-EXT-MIB	ENTERASYS-NAT-MIB
ENTERASYS-IETF-BRIDGE-MIB-EXT-MIB	VSB-SHARED-SECRET-MIB	ENTERASYS-LSNAT-MIB
ENTERASYS-IETF-P-BRIDGE-MIB-EXT-MIB	ENTERASYS-SNTP-CLIENT-MIB	ENTERASYS-VRRP-EXT-MIB DEFINITIONS
ENTERASYS-IF-MIB-EXT-MIB	ENTERASYS-RADIUS-ACCT-CLIENT-EXT-MIB	SNMP-RESEARCH-MIB
ENTERASYS-IP-SLA-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	ospfIfStateChange ospfVirtIfStateChange ospfNbrStateChange ospfVirtNbrStateChange ospfIfConfigError 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 No.	Title
RFC 2787	vrrpTrapNewMaster vrrpTrapAuthFailure
RFC 3621	pethPsePortOnOffNotification pethMainPowerUsageOnNotification pethMainPowerUsageOffNotification
RFC4268	entStateOperEnabled entStateOperDisabled
Enterasys-mac-locking-mib	etsysMACLockingMACViolation
Cabletron-Traps.txt	boardOperational boardNonOperational wgPsInstalled wgPsRemoved wgPsNormal wgPsFail wgPsRedundant wgPsNotRedundant fanFail fanNormal boardInsertion boardRemoval
	etsysPseChassisPowerRedundant etsysPseChassisPowerNonRedundant etsysPsePowerSupplyModuleStatusChange
Power-ethernet-mib	pethPsePortOnOffNotification pethMainPowerUsageOnNotification pethMainPowerUsageOffNotification
Enterasys-link-flap-mib	etsysLinkFlapViolation
Enterasys-ietf-bridge-mib-ext-mib	etsysIetfBridgeDot1qFdbNewAddrNotification etsysIetfBridgeDot1dSpanGuardPortBlocked etsysIetfBridgeDot1dBackupRootActivation etsysIetfBridgeDot1qFdbMovedAddrNotification etsysIetfBridgeDot1dCistLoopProtectEvent
Enterasys-flow-limiting-mib	etsysFlowLimitingFLowCountActionLimit1 etsysFlowLimitingFLowCountActionLImit2
Enterasys-notification-auth-mib	etsysMgmtAuthSuccessNotificiation etsysMgmtAuthFailNotificiation
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

RFC No.	Title
Enterasys-mstp-mib	etsysMstpLoopProtectEvent
Ctron-environment-mib	chEnvAmbientTemp
	chEnvAmbientStatus

RADIUS ATTRIBUTE SUPPORT:

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

RADIUS AUTHENTICATION AND AUTHORIZATION ATTRIBUTES:

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

RADIUS ACCOUNTING ATRRIBUTES:

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

GLOBAL SUPPORT:

By Phone: **603-952-5000**

1-800-872-8440 (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@enterasys.com**

By Web: **www.extremenetworks.com/support/**

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

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

F0615-O