

Customer Release Notes

7100-Series®

Firmware Version 8.41.01.0004

September 2015

INTRODUCTION:

This document provides specific information for version 8.41.01.0004 of firmware for the Enterasys 7100-Series products:

7100-Series Chassis			
71K11L4-48	71K11L4-24	71K91L4-48	71K91L4-24
71G21K2L2-48P	71G21K2L2-24P24	71G11K2L2-48	

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

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

PRODUCT FIRMWARE SUPPORT:

Status	Firmware Version	Product Type	Release Date
Current Version	8.41.01.0004	Customer Release	September 2015
Previous Version	8.32.02.0008	Customer Release	May 2015
Previous Version	8.32.01.0024	Customer Release	March 2015
Previous Version	8.31.03.0003	Customer Release	January 2015
Previous Version	8.31.02.0014	Customer Release	December 2014
Previous Version	8.31.01.0006	Customer Release	September 2014
Previous Version	8.22.03.0006	Customer Release	July 2014
Previous Version	8.22.02.0012	Customer Release	June 2014
Previous Version	8.21.03.0001	Customer Release	January 2014
Previous Version	8.21.01.0002	Customer Release	December 2013
Previous Version	7.91.03.0007	Customer Release	July 2013

7100-Series Customer Release Notes

Status	Firmware Version	Product Type	Release Date
Previous Version	7.91.02.0006	Customer Release	March 2013
Previous Version	7.91.01.0001	Customer Release	December 2012

HIGH AVAILABILITY UPGRADE (HAU) FW COMPATIBILITY:

HAU Key for this release: 059657e5e969558808eb8cfb67ca147d67daeaf6

The HAU 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:

System Minimum FW Version Required:

7100-Series Chassis			
Model	Minimum FW Version	Model	Minimum FW Version
71K11L4-48	07.91.01.0001	71G21K2L2-48P	08.21.01.0002
71K11L4-24		71G21K2L2-24P24	
71K91L4-48		71G11K2L2-48	08.22.02.0012
71K91L4-24			

It is recommended that the latest version of firmware be downloaded and the system be upgraded to the latest version of firmware prior to installation.

System Behavior

7100G - 71G21K2L2-48P / 71G21K2L2-24P24 / 71G11K2L2-48 Supported Port Configurations

The 7100G-Series models (71G21K2L2-48P, 71G21K2L2-24P24, and 71G11K2L2-48) do not support all combinations of front panel 10/100 Mb, Gigabit, 10 Gigabit, and 40 Gigabit port configurations. The dual QSFP+ ports must both be configured as 40Gb ports or either both as 4 x 10Gb Ethernet ports. When the two QSFP+ ports are configured as 4 x 10Gb Ethernet ports, the two SFP+ ports are not available for use and will be reported as not present. QSFP+ ports can each be individually configured as either Ethernet or VSB ports when in 40Gb mode.

7100-Series Customer Release Notes

Supported 7100G port configurations are shown in the table below.

7100G-Series Model	RJ45 Triple Speed PoE+ Ports	SFP 100Mb/1Gb Ports	SFP+ 1/10Gb Ethernet Ports	QSFP+	
				10Gb Ethernet Ports (4x10Gb Mode)	40Gb Ethernet or VSB Ports
71G21K2L2-48P	48	-	2	-	2
	48	-	-	8	-
71G21K2L2-24P24	24	24	2	-	2
	24	24	-	8	-
71G11K2L2-48	-	48	2	-	2
	-	48	-	8	-

MACsec Support

7100-Series MACsec Capable Ports

7100G MACsec capable ports – 10/100/1000Mb Base-T ports and 1Gb/10Gb SFP+ ports

7100K MACsec capable ports – 1Gb/10Gb BASE-T ports, and 1Gb/10Gb SFP+ ports

A MACsec license is required per unit to enable MACsec.

71A-EOS-GMACSEC - MACsec 7100G system License to enable MACsec

71A-EOS-KMACSEC - MACsec 7100K system License to enable MACsec

MACsec Limitations:

100Mb/1Gb SFP ports and 40Gb QSFP+ ports are not MACsec capable.

The MGBIC-02 copper SFP transceiver cannot be used with MACsec enabled in SFP+ ports.

Supported 10GBASE-T Port Speeds

10GBASE-T ports on 71K91L4-24 and 71K91L4-48 support 1Gb/10Gb speeds. With 8.41.01, 100Mb port speed is not supported on 10GBASE-T ports.

Half-Duplex Port Operation

The 7100-Series does not support half-duplex port configuration at any speed.

7100-Series Policy Capacities

Up to 63 policy profiles are supported by the 7100-Series.

Each 7100-Series chassis has a maximum authenticated user capacity of 512 MAC or port addresses with tci-overwrite enabled on all admin policy rules. A VSB stack of 8 7100s has a maximum authenticated user capacity of 4096 (8x512) MAC or port addresses.

7100-Series User Capacities:

Chassis Type	Maximum Authenticated MAC Address Capacity
71K11L4-48	512
71K11L4-24	512
71K91L4-48	512
71K91L4-24	512
71G21K2L2-48P	512

7100-Series Customer Release Notes

Chassis Type	Maximum Authenticated MAC Address Capacity
71G21K2L2-24P24	512
71G11K2L2-48	512

On the 7100-Series, there are 1024 hardware credits available for admin policy rules that are used for authenticated users or port addresses. The hardware cost for each type of admin rule is:

- 1 admin rule pointing to a profile with tci-overwrite enabled: 2 credits
- 1 admin rule pointing to a profile with tci-overwrite disabled: 3 credits

This means:

If all admin rules use tci-overwrite enable behavior up to 512 authenticated users are supported
 If all admin rules use tci-overwrite disable behavior up to 341 authenticated users are supported

A combination of the 2 types of admin rules is supported with the cost structure outlined above.

Combined tci-overwrite enabled and tci-overwrite disabled admin policy configuration example:

257 admin rules pointing to a profile with tci-overwrite enabled = 514 hardware credits
 170 admin rules pointing to a profile with tci-overwrite disabled = 510 hardware credits
 427 Total admin rules/authenticated users combined: Total hardware credits used = 1024 credits

Policy Resource Allocation Profile - The user can configure the policy resource allocation limits by selecting a profile from a predefined profile list using the “set limits resource-profile” command. The predefined profiles are “default” and “router1”. The “router1” profile allows for ingress ACL/PBR support.

```
TOR(su)->set limits resource-profile ?
  default                Default allocation profile
  router1                Router1 allocation profile
```

Policy Rule Traffic Classification Group	Maximum Policy Rule Capacity per Group: Default profile	Maximum Policy Rule Capacity per Group: Router1 profile
macsource macdest	121	0
ipv6dest	128	0
ipsourcesocket ipdestsocket ipfrag udpsourceportIP udpdestportIP tcpsourceportIP tcpdestportIP ipttl iptos iptype	250	249

7100-Series Customer Release Notes

Policy Rule Traffic Classification Group	Maximum Policy Rule Capacity per Group: Default profile	Maximum Policy Rule Capacity per Group: Router1 profile
Ethertype port	182	175

7100-Series Virtual Switch Bonding (VSB) Implementation Guidelines

Up to 8 7100-Series systems can be bonded using VSB, in any mix of chassis types.

VSB Support on Port Types - Only 40 Gigabit ports can be used as VSB interconnect ports. 10 Gigabit and 1 Gb ports can only be used as LFR ports. LFR is supported for VSB virtual stacks up to 8 systems.

Any port configured for VSB or LFR should only have bonding related configuration applied.

A closed ring VSB interconnect is not required, but if you do not close the ring and an interconnect or a system failure occurs, the remaining systems could be divided, causing two systems to reside in your network with the same IP address. LFR is highly recommended if a closed ring VSB topology is not used.

When replacing a system in a VSB stack you can restore the port level configuration by appending the configuration with the configuration from a previously stored configuration file when the chassis was operational within the stack, using the **configure filename append** command.

Port Mirroring

The 7100-Series device supports traffic mirroring for a maximum of 2 destination ports for mirrors.

A mirror could be a:

- "One-to-one" port mirror
- "One-to-many" port mirror
- "Many-to-one" port mirror

This allows configurations like: (a) up to two one-to-one mirrors, (b) up to two many-to-one mirrors, or (c) a single one-to-two mirror.

For the "one-to-many" there can be up to 2 destination ports.

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

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

LAG ports can not be used as the source port for a mirror.

Mirror destinations can be physical ports or LAGs, including ones on other switches in the same stack. Mirror destinations can not be VLANs.

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

Mirroring egress traffic results in the mirrored traffic always having an 802.1Q VLAN tag. The VLAN and priority values are the ones used for transmission of the original packet.

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.

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. They are not allowed to be changed as these are the default 802.1Q mappings for priority to queue.

Class of Service Support:

- Supports up to 256 Classes of Service
- ToS rewrite

7100-Series Customer Release Notes

- 802.1D/P Priority
- 9 Transmit Queues per port (8 customer and 1 internal reserved for control-plane traffic)
 - Queues support Strict, WFQ, ETS, and Hybrid Arbitration
 - All queues support rate-shaping
- 16 Inbound-Rate-Limiters per port
- Support for Flood-Limiting controls for Broadcast, Multicast, and Unknown Unicast per port.
- Management
 - Support for Enterasys CoS MIB

No support for Outbound-Rate-Limiters

Link Aggregation (LAG)

The 7100-Series chassis supports a total of 64 LAGs per chassis with up to 8 ports per LAG.

Multi-User 802.1X

Authentication of multiple 802.1X clients on a single port is supported. This feature 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.

This version of firmware does not support retrying MAC address authentication for failed stations, or renewing MAC address authentications for successful ones.

RMON Statistics:

Oversized packets are not counted on a port that is not enabled for jumbo frames.

If this oversized packet has an invalid CRC, it will be considered a jabber packet rather than an oversized packet.

RMON Packet Capture:

RMON packet capture is supported on the 7100-Series with the following limitations:

1. The 7100-Series only support one RMON channel and one RMON packet capture at a time.
2. The 7100-Series captures ingress packets, but not egress packets.
3. The 7100-Series will capture up to 100 consecutive full-size packets (size 1522 bytes or less), or up to 200 consecutive small packets (size 768 bytes or less). These are raw ingress packets. If there are non-trivial RMON filters, the number of packets selected for the RMON capture buffer may be a less.
4. The 7100-Series will automatically shut off hardware packet capture after the RMON packet buffer fills up, or after RMON sees 200 packets, whichever comes first.

7100-Series Customer Release Notes

SMON Guidelines:

The 7100-Series does not support port-VAN or LAG ports for SMON statistics collection.

Flash File System:

If for any reason the flash file system become seriously corrupted and nonfunctional the flash file system can be reformatted and the firmware image reloaded. Call Enterasys support.

Scale and Capacity Limits

Each release of 7100-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. Please use this command to get a complete list of capacities for this release.

	7100-Series	
ARP Entries (per router / per chassis)	4K	
Static ARP Entries	512	
IPv4: Route Table Entries	12000	
IPv6: Route Table Entries (/64)	6000	
IPv4: Router interfaces	256	
IPv6: Router interfaces	256	
OSPF Areas	8	
OSPF LSA(s)	12000	
OSPF Neighbors	60	
Static Routes	1024	
RIP Routes	2500	
Configured RIP Nets	300	
VRRP Interfaces	256	
ACLs	Resource Profile - default	Resource Profile - router1
IPv4 Ingress Access-Group Rules	0	128
IPv4 Egress Access-Group Rules	256	256
IPv6 Ingress Access-Group Rules	0	128
IPv6 Egress Access-Group Rules	256	256
Policy Based Routing (PBR) Entries (IPv4 only)	0	50
IPv4 Route-Map (Rules for all PBR entries)	0	128
ECMP Paths	8	
Static VRFs	128	
Dynamic VRFs	64	
Secondaries per Interface	128	
Total Primary + Secondary Interfaces per Router	512	
IP Helper addresses (per router/ per interface)	5120 / 20	
SPBv (constrained by 4094 VLANs)	Up to 100 VLANs mapped as base VIDs	Up to 100 SPBv nodes in SPB region

7100-Series Customer Release Notes

Multicast Capacities

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

¹ 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	5000
DHCP Pools	100

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

Advanced Routing License Feature

The 7100-Series Advanced Routing License license adds routing features to the 7100-Series.

7100-Series Chassis	Advanced Routing License	Licensed Features
71K11L4-48	71A-EOS-ADVL3	OSPFv2/v3, PIM-SM, PIM-SMv6, PIM-DM, PIM-SSM, PIM-SSMv6, BGP, ISIS, Fabric Routing, VRF
71K11L4-24		
71K91L4-48		
71K91L4-24		
71G21K2L2-48P	71A-EOS-G-ADVL3	
71G21K2L2-24P24		
71G11K2L2-48		

An advanced routing license is required per chassis in a VSB stack if Advanced Routing features are to be supported.

MACsec Licenses

In support of MACsec, there are two 7100-Series licenses. A MACsec license is required per unit to enable MACsec.

71A-EOS-GMACSEC - MACsec 7100G system License to enable MACsec on 10/100/1000Mb Base-T ports, and 1Gb/10Gb SFP+ ports

71A-EOS-KMACSEC - MACsec 7100K system License to enable MACsec on 100Mb/1Gb/10Gb BASE-T ports, and 1Gb/10Gb SFP+ ports

Virtual Switch Bonding (VSB)

No License is required for VSB support in the 7100-Series.

7100-Series Customer Release Notes

NETWORK MANAGEMENT SOFTWARE:

NMS	Version No.
NetSight Suite	6.1 or greater

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

PLUGGABLE PORTS SUPPORTED:

100Mb Optics: Supported on 7100G SFP ports only – 71G21K2L2-24P24 & 71G11K2L2-48

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

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-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)

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-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+
10GB-SRSX-SFPP	10Gb/1Gb Dual Rate, MM 850nm 10GBASE-SR / 1000BASE-SX, LC, SFP+
10GB-LRLX-SFPP	10Gb/1Gb Dual Rate, SMF 1310nm 10GBASE-LR / 1000BASE-LX, LC, SFP+
10GB-LR271-SFPP	10G Gb, CWDM SM, 1271 nm, 10 km, LC SFP+

7100-Series Customer Release Notes

10GB-LR291-SFPP	10G Gb, CWDM SM, 1291 nm, 10 km, LC SFP+
10GB-LR311-SFPP	10G Gb, CWDM SM, 1311 nm, 10 km, LC SFP+
10GB-LR331-SFPP	10G Gb, CWDM SM, 1331 nm, 10 km, LC SFP+
SFP+ Direct Attach Copper Cables	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
SFP+ Direct Attach Active Optical Cables	Description
10GB-F10-SFPP	10Gb Active optical direct attach cable with integrated SFP+ transceivers, 10m
10GB-F20-SFPP	10Gb Active optical direct attach cable with integrated SFP+ transceivers, 20m

40Gb Transceivers:

QSFP+ Optics	Description
40GB-SR4-QSFP	40Gb, 40GBASE-SR4, MM 100m OM3, MPO QSFP+ Transceiver
40GB-ESR4-QSFP	40Gb, Extended Reach SR4, MM, 300m OM3, MPO QSFP+
40GB-LR4-QSFP	40Gb, 40GBASE-LR4, SM 10km LC QSFP+ Transceiver
10326	40Gb, QSFP+ Parallel Single Mode (PSM), MPO connector, 10km SMF
10327	MPO to 4xLC SMF 10m patch cord (for use with 10326)
QSFP+ Direct Attach	Description
40GB-C0.5-QSFP	40Gb, Copper DAC with integrated QSFP+ transceivers, 0.5m
40GB-C01-QSFP	40Gb, Copper DAC with integrated QSFP+ transceivers, 1m
40GB-C03-QSFP	40Gb, Copper DAC with integrated QSFP+ transceivers, 3m
40GB-C07-QSFP	40Gb, Copper DAC with integrated QSFP+ transceivers, 7m
40GB-F10-QSFP	40Gb, Active Optical DAC with integrated QSFP+ transceivers, 10m
40GB-F20-QSFP	40Gb, Active Optical DAC with integrated QSFP+ transceivers, 20m
10318	40Gb, Active Optical DAC with integrated QSFP+ transceivers, 100 m
10GB-4-C03-QSFP	10Gb, Copper DAC Fan out, 4xSFP+ to QSFP+, 3m
10GB-4-F10-QSFP	10Gb, Active Optical DAC, 4xSFP+ to QSFP+, 10m
10GB-4-F20-QSFP	10Gb, Active Optical DAC, 4xSFP+ to QSFP+, 20m
QSFP+ Adapter	Description
QSFP-SFPP-ADPT	QSFP+ to SFP+ Adapter

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

Only the above listed Extreme Networks 40 Gigabit optical transceivers are supported. Use of any other optical transceiver types results in an error.

7100-Series Customer Release Notes

Example Message for 40G cables that are unrecognized or unauthenticated

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

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

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

Auto Configuration of 4 x 10Gb Mode

The 7100-Series will recognize a 10GB-4-xxx-QSFP cable when inserted in a QSFP+ port and reconfigure a QSFP+ port to 4 x 10 Gigabit Ethernet. A system reset is required for the port speed change to take effect.

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

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

QSFP-SFPP-ADPT transceiver support:

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

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

Gigabit Support on QSFP+ ports:

When using the QSFP-SFPP-ADPT adapter on the 7100-Series, Gigabit port speed can be configured and a single Gigabit SFP can be used. When configured for Gigabit port speed, only the MGBIC-LC01 and MGBIC-LC09 Gigabit SFP transceivers are supported with the QSFP-SFPP-ADPT.

SFP and SFP+ Dual speed operation:

The SFP+ ports support the use of SFP+ transceivers and SFP transceivers. (10Gb/1Gb) SFP ports on the 7100G-Series models support the use of SFP transceivers and 100Mb transceivers. (1Gb/100Mb)

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

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

To disable auto-negotiation on an S-Series/7100-Series 40Gb port:

```
set port negotiation fg.x.y disable
```

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

7100-Series Customer Release Notes

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	ICMP
SNTP	Dynamic VLAN/Port Egress Configuration	Auto MDI-X Media Dependent Interface Crossover Detect (Enhanced for non auto negotiating ports)
Web-based configuration (WebView)	Ingress VLAN Tag Re-write	DHCP Server
Multiple local user account management	VLAN-to-Policy Mapping	Jumbo Frame support
Denial of Service (DoS) Detection	RMON – Statistic, History, Alarms, Events,	RMON Packet Capture
802.1X – Authentication	SMON – VLAN and Priority Statistics	CLI Management
Directed Broadcast	Cisco CDP v1/2	RADIUS (Accounting, Snooping)
802.1D – 1998	Distributed Chassis Management (Single IP Address)	Split RADIUS management and authentication
802.1Q – Virtual Bridged Local Area Networking	SNMP v1/v2c/v3	Port Mirroring
GARP VLAN Registration Protocol (GVRP)	IEEE 802.1ak MVRP (Multiple VLAN Registration Protocol)	Link Flap detection
802.1p – Traffic Class Expediting	MAC locking (Static/Dynamic)	Daylight Savings Time
802.1w – Rapid Reconfiguration of Spanning Tree	Node/Alias table	RFC 3580 with Policy support
802.1s – Multiple Spanning Trees	SSH v1/v2	IPv6 Node Alias Support
802.1t – Path Cost Amendment to 802.1D	Audit trail logging	RADIUS Client
802.3 – 2002		
802.1AX-2008 Link Aggregation (formerly 802.3ad)	FTP/TFTP Client	Virtual Switch Bonding (VSB) with Link Failure Response (LFR) links
802.3x – Flow Control	Telnet – Inbound/Outbound	Unidirectional Link Detection (ULD)
Broadcast Suppression	Configuration File Upload/Download	Configurable login banner
Ingress Rate Limiting	Text-based Configuration Files	High Availability FW Upgrades
Transmit queue shaping	Syslog	Type of Service (ToS) Re-write
Strict and Weighted Round Robin Queuing	Span Guard	802.3-2008 Clause 57 (Ethernet OAM – Link Layer OAM)
IGMP v1/v2/v3 and Querier support	Cabletron Discovery Protocol (CDP)	Path MTU Discovery
SMON Port and VLAN Redirect ?	LLDP and LLDP-MED	Secure Copy Protocol (SCP)
Spanning Tree Loop Protection	MLDv1/MLDv2	TACACS+
Data Center Bridging 802.1Qaz Enhanced Transmission Selection (ETS), Data Center Bridging Exchange Protocol (DCBx), Application Priority	Data Center Bridging 802.1Qbb Priority Flow Control (PFC)	Data Center Bridging 802.1Qau Congestion Notification (CN)
IP Routing	DVMRPv3	OSPF/OSPFv3
Static Routes	RIP ECMP, CIDR configuration	OSPF ECMP
Protocol Independent Multicast - Sparse Mode (PIM-SM) IPv4/v6	Virtual Router Redundancy Protocol (VRRP) v2/v3	OSPF Alternate ABR
RIP v2	Policy-Based Routing	Graceful OSPF Restart (RFC 3623)
Proxy ARP	DHCP Server	Passive OSPF support
Basic Access Control Lists	DHCP Relay w/option 82	OSPF NSSA, equal cost multi-path

7100-Series Customer Release Notes

Features		
Extended ACLs	Static Multicast Configuration	Bidirectional Forwarding Detection (BFD)
iBGP	BGP Route Reflector	eBGP
BGP Graceful Restart	BGP Route Refresh	BGP 4 byte AS number
IPv6 Policy Based Routing	IPv6 Static Routing	BGP Extended Communities
PIM-SSM IPv4/v6	PIM-DM IPv4/v6	IPv6 ACLs
Tracked Objects	IPv6 DHCP Relay	IP Source Guard
VLAN Provider Bridging (Q-in-Q)	IPsec support for OSPFv3	RIPng
Anti-spoofing User Tracking and Control	ISIS	IPv6 Node Alias Support
IEEE 802.1Q-2011 Connectivity Fault Management (CFM)	Fabric routing	ISIS Graceful Restart
Dynamic Arp Inspection (DAI)	DHCP Snooping	IP Service Level Agreements (IPSLA)
Virtual Routing and Forwarding (VRF)	Remote Port Mirroring	RADIUS Server Load Balancing
IEEE 802.1aq SPBv Shortest Path Bridging	Transceiver extended digital diagnostics	Network load balanced servers (NLB)
802.1AE-2006 802.1X-2010 MACsec		IEEE 802.3az Energy Efficient Ethernet (EEE)

FIRMWARE CHANGES AND ENHANCEMENTS:

Feature Enhancements in 8.41.01.0004

Feature Enhancements in 8.41.01.0004

MACsec – MACsec is defined by IEEE802.1AE-2006 and 802.1X-2010 and can be used to provide hardware-based Point-to-Point link layer security using authentication and encryption using pre-shared key exchange between two MACsec-capable devices. MACsec licenses are required to enable MACsec. See System Behavior: MACsec Capable Ports for definition of ports capable of supporting MACsec and MACsec license definitions. Introduction of MACsec capability in 8.41.01 restricts 10GBASE-T port speed support to 1Gb/10Gb only.

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

Dual rate SFP+ support – 10GB-SRSX-SFPP and 10GB-LRLX-SFPP.

RMON packet capture and filter groups. The 7100-Series supports one RMON packet capture at a time with a maximum of up to 200 ingress packets on a port. For RMON packet capture limitations, see System Behavior section.

SPBv configuration of per port hello intervals and multiplier parameters:

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

Support for "show vlan fid <fid>" command filters the show vlan output and only displays the VLANs that match <fid>.

7100-Series Customer Release Notes

Feature Enhancements in 8.41.01.0004

Enabled AES CTR Ciphers

Three (3) new ciphers have been added to SSH:

- aes128-ctr AES in Counter mode, with 128-bit key
- aes192-ctr AES in Counter mode, with 192-bit key
- aes256-ctr AES in Counter mode, with 256-bit key

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

- hmac-sha1-etm@openssh.com SHA-1 with 20-byte digest and key length, encrypt-then-mac
- hmac-md5-etm@openssh.com MD5 with 16-byte digest and key length, encrypt-then-mac
- hmac-ripemd160-etm@openssh.com RIPEMD-160 algorithm with 20-byte digest length, encrypt-then-mac
- hmac-sha1-96-etm@openssh.com SHA-1 with 20-byte key length and 12-byte digest length, encrypt-then-mac
- hmac-md5-96-etm@openssh.com MD5 with 16-byte key length and 12-byte digest length, encrypt-then-mac

Problems Corrected in 8.41.01.0004

L1 Problems Corrected in 8.41.01.0004	Introduced in Version:
The command "show port status" does not report the correct type for a 4x10g-c05 copper hydra cable.	8.32.01
The command "show port status" may display improper type for 10318 100m active optical cable in 4 x 10Gb mode. "4x10g-f100" should appear.	8.32.01
A 10GBASE-T port sometimes does not establish link with a Intel Quad i340-T4 system. When connecting/disconnecting the cable repeatedly after a number of interactions (not always the same from 2 to 15,) the port no longer links. The link can be recovered by disabling/enabling negotiation on the system. This problem has been addressed by implementing the: set port low-power-mode <port-string> disable command. Execute this command to prevent the problem from happening. Use of this command disables auto power-down mode on the port that is linked to the Intel Quad i340-T4 system.	7.91.03
When setting speed on a port, and then disabling autoneg the speed may not be applied.	7.91.03
Occasionally, traffic with a priority greater than 3 causes LLDP neighbor entries to be lost.	7.00.01
LLDP sends incorrect requested and allocated power values in the 802.3 power via MDI TLV.	TBD
On 7100G devices with many port bounces over a period of time, the CPU utilization of the system goes to nearly 100% and causes a crash.	8.21.01
Module might reset during shutdown with possible "EDR Record" message.	7.91.01
Module might reset during initialization with possible "DR Record" message.	7.91.01
During initialization a 7100G may unexpectedly reset, producing an error log message similar to the following:	8.21.01

7100-Series Customer Release Notes

L1 Problems Corrected in 8.41.01.0004	Introduced in Version:
<p>Message nn/mmm Syslog Message vv.vv.vv.vvv dd/mm/yyyy hh:mm:ss</p> <p>Full Version: <version> <0>bcmStrat[5.tusrApplnit]bcm_cosq_gport_attach(0 (50), 240C8309, 34017B 32, 9) failed: Operation timed out (0x00ddae08 0x00ed1564 0x00ec72a0 0x0043ccdc 0x02714760 0x02716824 0x02718aa0 0x02722c68 0x00b5b2ec 0x02fdd0 c0 0x02720af4 0x0270e9d4 0x027214dc 0x009328b8 0x017970e4 0x00000000)</p>	
Disabling auto-negotiation on 1G transceiver, inserted in 10G SFP+ port, is not persistent.	7.91.01

L2 Problems Corrected in 8.41.01.0004	Introduced in Version:
<p>When bridge mode is changed and there are 10's of thousands of FDB entries, a message similar to the following may appear:</p> <p>"Unit 0: CLEAR_RESTORE: L2_ENTRY_2[2018] blk: ism0 index: 5536 : [0][60000000]"</p> <p>The system may also reset.</p>	7.91.03
If a MAC SA is learned on the ingress member port of a LAG, where the LAG spans multiple slots, the L2 entries programmed on the other LAG member slots for this MAC address may timeout causing flooding of traffic received on these other slots that is destined to the specified MAC address.	7.91.03
When running a bonded setup with macauth learned entries, a port bounce may cause the following message to appear, followed by a system crash: "soc_tr3_l2_bulk_age_stop: thread will not exit".	8.32.01
When large amounts of VLANs are deleted together (for example, created by MVRP/GVRP), the following message may appear, followed by a core file and reset: "Exc Vector: Machine Check exception(0x00000200)".	7.91.03
GVRP may fail to propagate dynamic VLANs on a LAG after a topology change. The result is that the switch on the remote side of the LAG fails to add the LAG to the tagged VLAN egress list. The only way to recover from this failure is to disable, and then re-enable, the LAG.	7.00.01
VLAN egress registered dynamically by MVRP may bounce when the system is in a steady state.	7.91.01
The CPU utilization may increase up to 99% indefinitely due to MVRP. The system may crash or require you to reset it.	8.31.01
MVRP may fail to propagate SPB Base VLANs on ports that are forwarding in the CIST context after disabling SPB on a device.	8.31.01
Spanning tree consumes all packets with the destination address for the IEEE Bridge Group Address/Nearest Customer Bridge group address. This has two effects. First, other applications for which the PDU is intended do not receive it. Second, a PDU which is not a BPDU is processed by spanning tree and marked as an invalid BPDU.	7.30.01

7100-Series Customer Release Notes

L2 Problems Corrected in 8.41.01.0004	Introduced in Version:
Spanning tree debug counters are incorrect for RSTP.	8.20.02
Connecting an SPB device in customer bridge mode to a bridge running in provider mode can produce malformed adjacencies with other devices, leading to network instability and spanning tree ports in "listening" state.	8.31.01
On boot up, in a device with multiple connections to root, there may be an initial delay of up to 10 seconds for the root port to reach the forwarding state and pass traffic.	7.63.01
Clearing a VLAN created through "set vlan create" occasionally causes traffic destined to a GVRP- or MVRP-configured port to be lost on the cleared VLAN in a multi-module system.	7.91.01
Multicast is not forwarded over LAGs correctly in bonded sys after MFM del.	8.32.01

Shortest Path Bridging Problems Corrected in 8.41.01.0004	Introduced in Version:
Traffic traversing an SPBV network does not egress out access ports. Filter database entries indicate traffic is not received on the correct internal ports. If the filter database is cleared, traffic correctly egresses out the access ports.	8.31.01
SPBV port may not become internal to the region even though ISIS adjacency is indicated.	8.31.01
MVRP may propagate SPBV Base-VID registrations on ports within the SPBV domain.	8.31.01
In Shortest Path Bridging, an SNMP query with a context of getNext on the ieee8021SpbTopNodeTable table causes the device to stop responding. The system ID index passed into the getNext query actually exists in the topology, which is the underlying problem. This effect may also occur when issuing SPB CLI commands to show topology information, such as "show spb neighbors".	8.31.01
Configuring multiple bridges with different SPBV SPVID allocation modes can produce high CPU utilization.	8.32.01
There is no user-evident notification that SPB ports go operationally down when setting that port's spantree adminPathCost to a value greater than 16777213.	8.32.02
Resetting a device in a stacked chassis, which isolates another device from the rest of the stack, causes the isolated device to reset.	8.32.01
After clearing and re-creating a static multicast MAC address, traffic destined through a Shortest Path Bridging network is dropped.	8.32.01
When backuproot is enabled for the CIST on a device that is part of an SPB region and the directly connected root bridge is external to the region, and backuproot is triggered for that device due to failure of the root bridge, the new topology resulting from the change in bridge priority is not affected. This results in a loss of connectivity. Spanning tree modifies the CIST bridge priority, but fails to convey the change to ISIS-SPB, which is responsible for calculating the topology within the SPB region.	8.32.01

7100-Series Customer Release Notes

L3 Problems Corrected in 8.41.01.0004	Introduced in Version:
When running in provider bridge mode, IGMP queries are not be transmitted properly.	8.32.01
An existing IPv6 NAT binding may continue to be used after the NAT outside interface has been de-configured.	7.91.01
For 7100 platforms, IPv4 ESP encrypted frames traversing a routed interface can cause stability issues with OSPF.	8.01.01
PIM may drop neighbor adjacencies when running with large number of PIM neighbors.	7.00.01
PIM Bootstrap messages are sent out that are slightly greater than MTU, requiring unnecessary IP fragmentation .	7.00.01
Multicast flows are not correctly forwarded after disabling, and then re-enabling, PIM on an interface.	8.21.01
When running PIM or DVMRP to route multicast traffic, errors similar to the following are appear: "RtrMc[1.tRMcEvt]Error deleting tmpFlow from TmpDb (2,723,1.1.1.1,225.1.1.1) = notFound" "[1.tRMcPkt]Hash find - flow vrfIds don't match (0,2)"	8.31.01
Multicast flows that are blocked by an ACL are constantly added and deleted by the multicast cache manager.	8.21.01
It is possible to enter an IPv6 address as a VRRP address when the VRID is a VRRPv3 IPv4 VRID. The address entered becomes a seemingly random IPv4 address in the configuration.	8.01.01

Platform Problems Corrected in 8.41.01.0004	Introduced in Version:
When the FDB table has 40K+ entries and they are deleted by disabling a port, a message similar to the following appear, resulting in a slot reset: "Unit 0: CLEAR_RESTORE: L2_ENTRY_2[2018] blk: ism0 index: 5536 : [0][60000000]"	7.91.03
System instability might occur with messages similar to the following appearing: "Interhost Unit 1 no rx space in Net Pool".	6.00.02
A stack that includes one or more 7100Gs may segment with messages similar to the following appearing: "FtmLi[2.bcmATP-RX]heartbeat rx on slot 2: from slot (3) != origin slot (1)." The segmentation may be triggered by a reset of a module or system, or by enabling or inserting a link on a stacking port. The segmentation and messages persist until the system is reset.	8.21.01
Stacking port in Spanning Tree may block when it should not.	7.91.01

7100-Series Customer Release Notes

Security Problems Corrected in 8.41.01.0004	Introduced in Version:
When ACL logging is enabled on a Policy ACL, the Policy ACL specific field "set-dscp <value>" does not appear in the log message.	8.32.01
Packets with multicast and broadcast source MAC addresses might get authenticated with CEP.	7.91.01
Number of MAC authentication sessions being used might be greater than the hardware can support.	8.21.01
When doing mac auth with a large number of users in a short period of time, the following message may appear, resulting in that mac being unable to authenticate.: "UPN[1]Policy 0(NoAuth) assignment by rule [MacSrc 00:00:55:55:01:92 ge.1.38] failed (exceeded blade hardware limits)".	7.91.03
When a mac auth to RADIUS is executed shortly after bootup, the following message may appear, and the mac may not being able to authenticate: "AAA[1]Unable to send frame because we are unable to obtain the default IPV4".	7.91.03
etsysPolicyRuleAttributeMaxCreatable rows were displaying system-wide rule limits for each rule type, instead of the more specific limits applicable to particular rule type groups.	7.91.01
etsysPolicyRulesMaxEntries reports more rules than actually supported.	7.91.01
Under higher frame rates, some source MAC addresses may not be authenticated, and may instead receive the default port policy.	8.22.01
Netsight Policy Manager might be unable to enable tci-overwrite on the profile.	7.91.01

Host Services Problems Corrected in 8.41.01.0004	Introduced in Version:
Configuration might fail to load due to device discovery timeouts and messages similar to the following may appear: "<2>System[6]Detected missing or reset module, aborting configure".	8.21.01
Module might reset with messages similar to the following may appear: "Chassis coherency timeout exceeded".	7.62.07
After a denial of service attack, in a multi-slot configuration, the 'dir' command only produces a list of the files on a single slot.	8.20.02
Chassis might experience stability/distribution issues during DoS LAN attack.	8.20.02
Denial of service (DoS) attack results in warning messages: "this server has been invalidated" printed to the console.	1.07.19
The help text for the "set/clear license { I3-7100k, I3-7100g " commands (on the 7100-Series) and the "set/clear license vsb" commands (on the S-Series) shows <CR> as an alternative to the "chassis" qualifier.	7.91.01
The help text for entering a slot or chassis number changed from <value> to <slot_number> or <chassis_number>.	

7100-Series Customer Release Notes

Host Services Problems Corrected in 8.41.01.0004	Introduced in Version:
EDR memory in free list error appears while setting snmpTargetAddrTDomain to a value other than snmpUDPDomain without changing snmpTargetAddrTAddress to match the domain type.	4.11.17
Updated CLI engine to make TAB-key function as '?' whenever command cannot be completed.	1.07.19
"set port negotiation <port> disable" command executed on 40G copper port may show a "not supported" error, or it may not be persistent after a reset of the system	7.91.01
When starting up, 7100-Series switches may display "semGive failed" error messages for sysLedMutex (for example, "Default[4.tDSSync3]sysLedMutex semGive failed!"). These messages are harmless.	7.91.01
During initialization, a 7100G may unexpectedly reset producing an error log message similar to the following: Message nnn/mmm Syslog Message vv.vv.vv.vvvv mm/dd/yyyy hh:mm:ss <0>bcmStrat[1.tEmanate1]bcm_cosq_gport_bandwidth_set(0 (14), 4, 240380D4, 0, 0, 0) failed: Invalid parameter (0x00d1d3e4 0x00e0978c 0x00e01d68 0x00439108 0x024e8de4 0x024e9f60 0x024da7f8 0x024f003c 0x025016b4 0x024f9940 0x02510830 0x006b4d08 0x006a0580 0x006c5948 0x006b9fe8 0x006bbf7c 0x00697798 0x016cd3a4 0x00000000)	8.21.01
If you issue a "set rmon stats" command that includes an owner name, but does not include a port string, the command uses the owner name as the port string, and then indicates that it is invalid.	1.07.19
Priority 0 might be slow to display correct values when using "show smon priority" command.	1.07.19
With SNTP unicast client configured, after 497 days, SNTP time requests may stop being sent.	4.05.08
Logging server list identifiers are translated incorrectly between releases causing logged messages to be directed to the incorrect logging server, console, file, or secure file.	7.40.00

Problems Corrected in 8.32.02.0008

ACL Problems Corrected in 8.32.02.0008	Introduced in Version:
When ACL logging is enabled on a Policy ACL, the Policy ACL specific field "set-dscp <value>" was not displayed in the log message.	8.32.01

Host Problems Corrected in 8.32.02.0008	Introduced in Version:
During initialization a 7100G may unexpectedly reset producing an error log message similar to the following: Message nnn/mmm Syslog Message vv.vv.vv.vvvv mm/dd/yyyy hh:mm:ss <0>bcmStrat[1.tEmanate1]bcm_cosq_gport_bandwidth_set(0 (14), 4, 240380D4, 0, 0, 0) failed: Invalid parameter (0x00d1d3e4 0x00e0978c 0x00e01d680x00439108 0x024e8de4	8.21.01

7100-Series Customer Release Notes

Host Problems Corrected in 8.32.02.0008	Introduced in Version:
0x024e9f60 0x024da7f8 0x024f003c 0x025016b4 0x024f9940 0x02510830 0x006b4d08 0x006a0580 0x006c5948 0x006b9fe8 0x006bbf7c 0x00697798 0x016cd3a4 0x00000000)	
IGMP Problems Corrected in 8.32.02.0008	Introduced in Version:
When IGMP/MLD snooping is not disabled on a 7100 stack, IGMP/MLD reports may not be correctly flooded out LAG ports.	8.21.01
When running in provider bridge mode, IGMP queries will not be transmitted properly.	8.32.01
Multicast Problems Corrected in 8.32.02.0008	Introduced in Version:
After clearing and recreating a static multicast MAC, traffic destined through a Shortest Path Bridging network will be dropped.	8.32.01
PIM-SM Problems Corrected in 8.32.02.0008	Introduced in Version:
Protocol neighbor adjacencies may briefly go down when starting lots of multicast traffic/reporters in a L2-meshed environment.	8.21.01
Shortest Path Bridging Problems Corrected in 8.32.02.0008	Introduced in Version:
Resetting a device in a stacked chassis, which isolates another device from the rest of the stack, will cause the isolated device to reset.	8.32.01
Spanning Tree Problems Corrected in 8.32.02.0008	Introduced in Version:
On boot up, in a device with multiple connections to root, there may be an initial delay of up to 10 seconds for the root port to reach the forwarding state and pass traffic.	7.63.01
VRF Problems Corrected in 8.32.02.0008	Introduced in Version:
DHCP relay agent does not work over L3VPN.	8.01.01

Problems Corrected in 8.32.01.0024

802.1d Filter Database Problems Corrected in 8.32.01.0024	Introduced in Version:
In configurations consisting of several multiple authentication (multiauth) sessions, MAC addresses are sometimes not aged out from the filter database.	8.22.01
Auto Negotiation Problems Corrected in 8.32.01.0024	Introduced in Version:
A 10GBASE-T port on a 7100K- or an S-Series system sometimes will not establish link with an Intel Quad i340-T4 system. When connecting/disconnecting the cable repeatedly after a number of interactions (not always the same from 2 to 15), the port no longer links. The link can be recovered by disabling/enabling negotiation on the 7100K- or S-Series system.	8.22.01
This problem has been addressed by implementing:	

7100-Series Customer Release Notes

Auto Negotiation Problems Corrected in 8.32.01.0024	Introduced in Version:
<p>set port low-power-mode <port-string> disable</p> <p>Execute this command to prevent the problem from happening. Use of this command disables auto power-down mode on the 7100K- or S-Series system port that is linked to the Intel Quad i340-T4 system.</p>	
Boot Config Problems Corrected in 8.32.01.0024	Introduced in Version:
Configuration might fail to load due to device discovery timeouts with messages similar to: "<2>System[6]Detected missing or reset module, aborting configure".	8.21.01
Data Center Bridging Problems Corrected in 8.32.01.0024	Introduced in Version:
A message similar to "DCB[1]HW: getCnCpHwIndexDD:CPHwIndex mapping was unable to map port 258, priority 2 to a txQueue" may be seen at startup on a 7100G with a CNPV enabled.	8.31.01
Distributed Services Problems Corrected in 8.32.01.0024	Introduced in Version:
Module might reset with messages similar to: "Chassis coherency timeout exceeded".	7.91.001
Denial of Service (DOS) attack results in warning messages "this server has been invalidated" printed to the console.	7.91.001
Hardware Problems Corrected in 8.32.01.0024	Introduced in Version:
A 1-Gigabit port with a copper SFP inserted on a 71G21K2L2-24P24 or 71G11K2L2-48 system may not achieve link when the system is booted even though both the local and remote port are enabled. If the port does not achieve link, a reset of the system may workaround the issue for that boot.	8.21.01
Host Mobility Problems Corrected in 8.32.01.0024	Introduced in Version:
Host-mobility is now supported in a segmented VRRP network design where the routers' VRRP interfaces are not connected.	8.31.01
IGMP Problems Corrected in 8.32.01.0024	Introduced in Version:
If a configuration is enabled for IGMP on a VLAN that becomes an SPVID, you can not delete the config.	8.31.01
When using a BaseVid without spvid allocation due to an insufficient spvid pool or a lack of boundary egress, IGMP may not forward traffic.	8.31.01
IPStack Problems Corrected in 8.32.01.0024	Introduced in Version:
Within a network environment where DHCP clients are active, over time, could see an exhaustion of resources that prevent IP host communication and loss of device management.	7.91.01

7100-Series Customer Release Notes

Licensing Problems Corrected in 8.32.01.0024	Introduced in Version:
On the 7100G platform, the supported irl-reference range is 0 through 7. Attempts to configure 8 through 15 results in error(s) similar to: "Unsupported irl reference (8) on 7100G. Only 0 - 7 supported."	8.21.01
The Help text for the 'set/clear license { I3-7100k, I3-7100g }' commands on the 7100-Series shows <CR> as an alternative to the "chassis" qualifier. The Help text for entering a slot or chassis number has changed from <value> to <slot_number> or <chassis_number>.	7.91.01
Link Aggregation Problems Corrected in 8.32.01.0024	Introduced in Version:
Occasionally when a port is added to a lag, the port's discard state will follow the port's configuration rather than the lag's configuration.	7.91.01
IGMP packets received on a LAG port will be reflected out other LAG ports elsewhere in the stack.	8.31.02
Management Problems Corrected in 8.32.01.0024	Introduced in Version:
When syslog servers are configured, if any of the following cli commands are issued: <ul style="list-style-type: none"> • show support • show config • show config logging <p>The switch will lose (leak) 144 bytes of memory. If commands are issued frequently enough the switch will reset, logging a message similar to:</p> <p>Message 3/30 EDR Record 07.62.05.0001H 07/27/2014 19:55:11 Severity/Facility: FATAL/KERNEL Task: tCLI0 Injection Point: memPartLib.c:2498 Address: 0x00000000</p> <p>memPartAlloc: block too big 84624 bytes (0x10 aligned) in partition 0x2234548</p>	7.91.01
MultiAuth Problems Corrected in 8.32.01.0024	Introduced in Version:
Number of mac authentication sessions being used might be larger than the hardware can support.	8.21.01
Multicast Problems Corrected in 8.32.01.0024	Introduced in Version:
If IGMP or MLD snooping is enabled on a VLAN that belongs to a VRF, and SPB-Vlan is enabled, but not for that VLAN, then hardware may be mis-programmed causing multicast traffic to be flooded.	8.31.02
First multicast packets may not be forwarded if Join(S,G) comes before data.	8.31.01

7100-Series Customer Release Notes

Multicast Problems Corrected in 8.32.01.0024	Introduced in Version:
Disabling IGMP/MLD on a VLAN or disabling a VLAN itself may result in multicast flows remaining programmed in hardware that do not recover after re-enabling.	8.21.01
Multicast flows ingressing LAG ports may be removed and re-added, causing a brief disruption in traffic, as well as causing "show ip mcache" to show lower age of flow.	8.20.02
Multicast cache entries show up in the router even without a multicast routing protocol enabled on an interface.	8.31.01
When running in a SPBV topology, IP multicast that is received on a SPVid may be flooded when the corresponding BaseVlan belongs to a VRF.	8.31.01
Multicast frames that are buffered and forwarded do not have TTL decremented.	8.31.01

MVRP Problems Corrected in 8.32.01.0024	Introduced in Version:
The CPU utilization may spike up to 99% indefinitely due to MVRP. The system may crash or require manual intervention to force a reset.	8.31.01
VLANs that are either forbidden or mapped to the SPBV MST at bootup will not allow dynamic registration via MVRP or GVRP after the VLAN forbidden egress status or MST mapping is cleared.	8.31.01

PIM-DM Problems Corrected in 8.32.01.0024	Introduced in Version:
"show ip mcache" shows a corrupted/incorrect Source/Destination IP in the display output.	8.31.01

Platform Problems Corrected in 8.32.01.0024	Introduced in Version:
When uploading a new image to chassis, errors noticed when distributing image to compatible slots in the chassis.	7.91.01
"Unable to delete a file/image from the users directory if it has the same name as the current running image. You will get the following error return. (su)->delete slot1/myImage The active image cannot be removed. Failed to remove /slot1/myImage"	7.91.01
Message similar to the following might be seen when bonding is disabled:<163>Feb 12 10:42:00 100.10.10.22 dot3Mgt[4.tEmanate10]dot3MgtDist::ifJackEntryGet():sendMessage(ackReq)!=kDs_good;sendMask=0x10000	7.91.01
When the FDB table has 40K+ entries and they are deleted by disabling a port, a message of the type "Unit 0: CLEAR_RESTORE: L2_ENTRY_2[2018] blk: ism0 index: 5536 : [0][60000000]" may be seen resulting in a slot reset.	7.91.03
Bootloader version 02.03.02 is included with this release. The bootloader's flash memory driver no longer refreshes flash pages found to have ECC-corrected errors in order to eliminate a window where the flash page could be corrupted.	7.91.01
On 7100K, 40G ports 1 and 3 would potentially not get link up with QSFP from transceiver 40GB-SR4-QSFP.	8.31.01
After bootup the system can become unmanageable from the network.	8.31.02

7100-Series Customer Release Notes

PoE Problems Corrected in 8.32.01.0024	Introduced in Version:
Manually enabling PoE on one port at a time by way of the CLI could occasionally cause the system to incorrectly believe that the PoE controller had encountered a fatal error and cause loss of PoE functionality on all ports until the unit was rebooted.	8.21.01
Policy Problems Corrected in 8.32.01.0024	Introduced in Version:
Under higher frame rates, some source MAC addresses may not be authenticated, and may instead receive the default port policy.	8.22.01
RMON Problems Corrected in 8.32.01.0024	Introduced in Version:
"show rmon stats" report might fail to include a bond port. This problem is intermittent (all of the bond ports might show up on some reboots), and the omitted bond port could change from reboot to reboot.	7.91.01
Shortest Path Bridging Problems Corrected in 8.32.01.0024	Introduced in Version:
Traffic traversing an SPBV network does not egress out access ports. Filter database entries indicate traffic is not received on the correct internal ports. If the filter database is cleared, traffic correctly egresses out the access ports.	8.31.01
Occasionally when a port's operational state is changed, layer 2 static multicast traffic over Shortest Path Bridging is lost on that port.	8.31.01
Port falls out of the SPB region when a spanning tree MSTI is created.	8.31.01
Ports may become blocked when adding a BVLAN or SPVID and then immediately removing it. Spanning tree reinitializes the port topology information calculated by ISIS-SPB, but the information is not refreshed because the topology calculated by ISIS-SPB has not actually changed.	8.31.01
When Shortest Path Bridging is globally disabled, Layer 2 multicast traffic will not be forwarded across a Virtual Switch Bond, when using a configured Shortest Path Bridging BaseVLAN.	8.31.02
When an SPB regional port becomes a boundary port and then reenters the region, ISIS-SPB and Spanning Tree may become out of sync with respect to the value the port is using for agreement digest. The value transmitted in an SPT BPDU may differ from the value transmitted in the SPB-Digest sub-TLV of the SPB Hello PDU. This may result in traffic loss due to agreement not being reached between the connected ports.	8.31.03
CIST root port may become stuck in the listening state when disabling and reenabling the global SPB status for all the nodes in an SPB region.	8.31.03
If Shortest Path Bridging is enabled, or enabled then disabled, a "show mac addr ..." command could take minutes or tens of minutes to complete. All matching Filter Database entries should still be returned.	8.31.01
SPB configurations using manual SPVID allocation mode without manually configured SPVIDs can lead to high CPU utilization and network instability.	8.31.01
SPB devices may not agree topology agreement digest after changing master role.	8.31.01
Traffic may not recover after disable/re-enable SPB.	8.31.01

7100-Series Customer Release Notes

Shortest Path Bridging Problems Corrected in 8.32.01.0024	Introduced in Version:
In a Shortest Path Bridging domain, when a device becomes the new regional root, designated ports on this new regional root go into listening state. Consequently, CIST traffic using this path is blocked. The issue is resolved by forcing a BPDU to be sent by the root port on the peer device.	8.31.01
In a Shortest Path Bridging-VLAN domain, when a device becomes the new regional root, customer traffic that ingresses the network on a base VID does not reach the intended destination endpoint(s). The associated SPVID lacks egress on some bridges throughout the SPBV network, and there is no clear indication of why this is so. The issue is resolved by forcing a BPDU to be sent by the root port on the peer device.	8.31.01
In a Shortest Path Bridging VLAN (SPBV) domain, some multicast traffic, including statically programmed L2 multicast entries, loops in the network.	8.31.01
In a Shortest Path Bridging VLAN (SPBV) domain, ports are incorrectly set to backup role and a state of blocking. The only ports affected are internal to the region and the consequence is limited network connectivity. Toggling the SPB configuration on the port may fix the problem, but not always.	8.31.01
For Software Bonded flows, from SPB ports, the first 4 bytes of the Software Bond Header is not getting removed properly, causing loss of L2 multicast traffic.	8.31.01
The agreement protocol for Spanning Tree internal to the SPB region requires an exchange of BPDUs greater in number than what is required for rapid failover in RSTP or MSTP. Spanning Tree software rate limiters may cause a BPDU drop during this exchange causing the protocol to be interrupted for a HELLO period, two seconds by default, until the next periodic transmit of a BPDU. This will delay convergence when SPB has the digest convention configured for loopFreeBoth.	8.31.01
In a Shortest Path Bridging VLAN domain, traffic loops are seen on directly adjacent 7100-series devices. Packet captures show that SPVID-tagged traffic egresses on ports that are not actually part of the VLAN egress membership. The problem is not seen if 7100-series devices are not connected to each other directly.	8.31.01
Port state may be listening for SPB internal port due to neighbor transmitting BPDUs with the agreeDigestValid flag persistently false.	8.31.01
Updating the L2 FDB may cause the device to reset and add a log entry into the message log reading "Assertion failed".	8.31.03

Spanning Tree Problems Corrected in 8.32.01.0024	Introduced in Version:
A root or alternate port may get stuck in a state where it will not respond to a proposal BPDU with an agreement BPDU. This will cause port forwarding for the connected designated port to use timers rather than the rapid forwarding mechanism. Additionally, if the designated port is configured for lp (Loop Protect), it will detect a loop protect event and remain in the listening state.	7.91.01
The Multisource function detects multiple BPDU sources received on a point-to-point link and sets the point-to-point operational status to false. The point-to-point operational status is an input into the rapid transition to forwarding capability for rapid spanning tree. It is also a factor in the Loop Protection mechanism and in Shortest Path Bridging. A port that receives BPDUs from multiple sources where those sources are exclusively different ports on the same transmitting bridge will not be triggered for multisource and will	8.31.01

7100-Series Customer Release Notes

Spanning Tree Problems Corrected in 8.32.01.0024	Introduced in Version:
remain operationally point-to-point.	
FDB entry not removed for IST port in an SPB region during a topology change. This can cause traffic assigned to VLANS mapped to SID 0 to be directed out the wrong port until the FDB entry times out.	8.31.01
A port on the root bridge may select a backup role instead of a designated role if it receives a BPDU from another bridge where the role in the flags field indicates a designated role, the root identifier is the ID of the receiving bridge and the transmitting port ID is lower than the receiving port ID.	7.91.01
A temporary loop may be created when the root bridge relinquishes its root status and the direction of root in the network reverses, i.e. designated ports become root/alternate ports and root/alternate ports become designated.	7.91.01
Configuring dot1dStpPortEnable for a port, either with SNMP or with the spantree portenable CLI command, to the disabled value, will prevent the port from attaching to a LAG. This changes prior behavior.	8.31.03

Stacking Problems Corrected in 8.32.01.0024	Introduced in Version:
Stacking port in Spanning Tree may block when it should not have been.	7.91.01

Transceiver Problems Corrected in 8.32.01.0024	Introduced in Version:
Some QSFP+ transceivers will come up as "Unauthenticated".	8.31.02

Transmit Queue Monitor Problems Corrected in 8.32.01.0024	Introduced in Version:
<p>During initialization, a 7100G may unexpectedly reset producing an error log message similar to the following:</p> <p>Message nn/mmm Syslog Message vv.vv.vv.vvvv dd/mm/yyyy hh:mm:ss</p> <p>Full Version: <version></p> <p><0>bcmStrat[5.tusrApplnit]bcm_cosq_gport_attach(0 (50), 240C8309, 34017B32, 9) failed: Operation timed out (0x00ddae08 0x00ed1564 0x00ec72a0 0x0043ccdc 0x02714760 0x02716824 0x02718aa0 0x02722c68 0x00b5b2ec 0x02fdd0c0 0x02720af4 0x0270e9d4 0x027214dc 0x009328b8 0x017970e4 0x00000000)</p>	8.21.01

VLAN Problems Corrected in 8.32.01.0024	Introduced in Version:
VLAN egress registered dynamically by MVRP may bounce when the system is in a steady state.	7.91.01

Problems Corrected in 8.31.03.0003

Shortest Path Bridging Problems Corrected in 8.31.03.0003	Introduced in Version:
In a large Shortest Path Bridging network, running the command "show spb path" will cause the Shortest Path network traffic to stop forwarding.	8.31.02

7100-Series Customer Release Notes

Shortest Path Bridging Problems Corrected in 8.31.03.0003	Introduced in Version:
SPB Port configuration will be lost if hello parameters are configured and lower port numbers do not have hello parameters configured.	8.31.02

Spanning Tree Problems Corrected in 8.31.03.0003	Introduced in Version:
A temporary loop may be created when the root bridge relinquishes its root status and the direction of root in the network reverses, i.e. designated ports become root/alternate ports and root/alternate ports become designated.	7.91.03

Link Aggregation Problems Corrected in 8.31.03.0003	Introduced in Version:
IGMP packets received on a LAG port will be reflected out other LAG ports elsewhere in the stack.	8.31.01

Hardware Problems Corrected in 8.31.03.0003	Introduced in Version:
A 1-Gigabit port with a copper SFP inserted on a 71G21K2L2-24P24 or 71G11K2L2-48 system may not achieve link when the system is booted even though both the local and remote port are enabled.	8.21.01

Problems Corrected in 8.31.02.0014

Layer 2 Problems Corrected in 8.31.02.0014	Introduced in Version:
A performance reduction causes the throughput of new traffic processing to be reduced with default configuration.	8.31.01
Changing Bridge Mode while FDB is full & unicast traffic causes core & reset.	7.91.03
MAC gets stuck in FDB and will not age out.	8.22.01
VLANs that are either forbidden or mapped to the SPBV MST at bootup will not allow dynamic registration via MVRP or GVRP after the VLAN forbidden egress status or MST mapping is cleared.	8.31.01
A root or alternate port may get stuck in a state where it will not respond to a proposal BPDU with an agreement BPDU. This will cause port forwarding for the connected designated port to use timers rather than the rapid forwarding mechanism. Additionally, if the designated port is configured for LP (Loop Protect), it will detect a loop protect event and remain in the listening state.	7.91.01
Multisource will fail to trigger for multi BPDUs sent from same switch.	8.31.01
FDB entry not cleared on topology change resulting in temporary traffic loss.	8.31.01
A port on the root bridge may select a backup role instead of a designated role if it receives a BPDU from another bridge. Where the role in the flags field indicates a designated role, the root identifier is the ID of the receiving bridge and the transmitting port ID is lower than the receiving port ID.	7.91.01

Layer 2 Multicast Problems Corrected in 8.31.02.0014	Introduced in Version:
IGMP: mgmdStdMIB errors when running traffic.	7.91.01
IGMPv3 GS query message resets 'Other Querier Present Timer'.	7.91.01

7100-Series Customer Release Notes

Layer 3 Problems Corrected in 8.31.02.0014	Introduced in Version:
BFD neighbor state does not return to full state after master failure.	8.31.01
Module might reset with message indicating DSI Exception in Thread Name: tTrackBfdS.	8.31.01
The size of the IPv6 frame for the ICMPv6 redirect error message does not conform to the maximum of the IPv6 minimum MTU size of 1280 bytes.	7.91.01
OSPF/PIM - OSPFv3 neighbors bounce when mcast traffic started.	8.31.01
If an OSPFv2 virtual link is configured with an invalid timer value of 0, the router will crash with the following syslog message, "sms_get timeout: oid=3e000001, tRtrPtcls state: running, last wakeup: 1 tics, IPS in use cnt: 1968, Bytes: 6527728".	7.91.01

Layer 3 Multicast Problems Corrected in 8.31.02.0014	Introduced in Version:
Multicast flows are not correctly forwarded after disabling then re-enabling PIM on an interface.	8.21.01
First multicast packets may not be forwarded if Join(S,G) comes before data.	8.31.01
Disabling IGMP/MLD on a VLAN or disabling a VLAN itself may result it multicast flows remaining programmed in hardware that do not recover after re-enabling.	8.21.01
When running in a SPBV topology, IP multicast that is received on a SPVid may be flooded when the corresponding BaseVLAN belongs to a VRF.	8.31.01
IP Multicast flows may revert to a "register state" after PIM events such as neighbor loss, RP loss, etc.	8.31.01
User is unable to disable or delete an IGMP configuration for a VLAN if the Vid becomes configured as an Spvid.	8.31.01
CLI Syslog may indicate that a failed IGMP configuration succeeded, when it did not.	7.91.01
If adding an SPB base vid, before enabling IGMP, IGMP may not recognize the base vid, resulting in traffic issues.	8.31.01
A User is able to enable IGMP query on an SPBV Spvid.	8.31.01
Multicast cache entries show up in the router even without a multicast routing protocol enabled on an interface.	8.31.01
Multicast frames that are buffered and forwarded do not have TTL decremented.	8.31.01

Data Center Bridging Problems Corrected in 8.31.02.0014	Introduced in Version:
A message similar to "DCB[1]HW: getCnCpHwIndexDD:CPHwIndex mapping was unable to map port 258, priority 2 to a txQueue" may be seen at startup on a 7100G with a CNPV enabled.	8.31.01
DCB application priority will restore only one port's configuration.	7.91.01

Host Services Problems Corrected in 8.31.02.0014	Introduced in Version:
Slot reset with message similar to "this server has been invalidated".	7.91.01
Module might reset with messages similar to: "DSI exception" and "Thread Name: tDSrecv4".	7.91.01

7100-Series Customer Release Notes

Host Services Problems Corrected in 8.31.02.0014	Introduced in Version:
The Help text for the 'set/clear license { I3-7100k, I3-7100g }' commands (on the 7100-Series) and the 'set/clear license vsb' commands (on the S-Series) shows <CR> as an alternative to the "chassis" qualifier. The Help text for entering a slot or chassis number has changed from <value> to <slot_number> or <chassis_number>.	7.91.01
Disabling auto-negotiation on 40G port is not persistent and it generates message similar to: "fg.1.1 does not support specified feature".	7.91.01
The "no ip forward-protocol udp" commands do not return to the configuration after reboot.	8.01.01
When writing to a file on a remote blade, if the connection becomes unresponsive, the local blade could reset. An example would be running the following command from the master slot to a slot across a bond link: ""show config all outfile slot13/showCfgAll.out"" The log should have something similar to the following: Message 83/263 Exception PPC750 Info 08.30.01.0036 08/13/2014 08:54:27 Exc Vector: DSI exception (0x00000300) Thread Name: tCLIO"	7.91.01
"show rmon stats" report might fail to include a bond port. This problem is intermittent (all of the bond ports might show up on some reboots), and the omitted bond port could change from reboot to reboot.	7.91.01
Slot reset with message similar to "nvFilePtrMgr::fopen_ab(4,0,0,50, 4) fopen(/flash1/nonvol/0/b0000000.032,ab+)".	7.91.01
Underlying transport errors will cause the messages "TIPC discarding incoming Ethernet message with destination <mac_address>" to be displayed resulting in internal network buffer loss and a segmentation of a slot in a chassis to stand alone mode.	8.31.01
Layer 1 Problems Corrected in 8.31.02.0014	Introduced in Version:
A CPU under heavy load may prevent transmission of OAMPDUs which can lead to a discovery timeout on an OAM peer.	8.31.01
Disabling auto-negotiation on 1G transceiver, inserted in 10G SFP+ port, is not persistent.	7.91.01
Inserting a "Seimon 40G 0.5m copper QSFP cable" into a 40G port will result in board resetting.	8.31.01
When displaying debug CLI base information for some copper SFP cable assemblies, the output may incorrectly display the interface type as "40G Act Cbl" instead of "1000BASE-CX".	8.22.02
LLDP packets may be dropped when the port llb mode is set to flow control.	8.31.01
Policy Problems Corrected in 8.31.02.0014	Introduced in Version:
Under higher frame rates, some source MAC addresses may not be authenticated, and may instead receive the default port policy.	8.22.01

7100-Series Customer Release Notes

Platform Problems Corrected in 8.31.02.0014	Introduced in Version:
A stack that includes one or more 7100Gs may segment with messages similar to "FtmLi[2.bcmATP-RX]heartbeat rx on slot 2: from slot (3) != origin slot (1)." The segmentation may be triggered by a reset of a module, system or by enabling or inserting a link on a stacking port. The segmentation and messages will persist until the system is reset.	8.21.01
System instability might be experienced with messages similar to "Interhost Unit 1 no rx space in Net Pool".	7.91.01
Bootloader version 02.03.02 is included with this release. The bootloader's flash memory driver no longer refreshes flash pages found to have ECC-corrected errors in order to eliminate a window where the flash page could be corrupted.	7.91.01

Shortest Path Bridging Problems Corrected in 8.31.02.0014	Introduced in Version:
If individual blades are reset on stacked systems, Traffic transmitted via SPB with VRRP mac addresses may fail to route properly after egressing the SPB domain.	8.22.03
When a bridge has the SPB global status set from disabled to enabled, ports on the bridge and ports on attached bridges may have spanning tree port states stuck in listening. This can occur for the default spanning tree as well as shortest path trees. Toggling the port's administrative status will clear the condition.	8.31.01
Traffic within a SPBV topology does not recover when pulling and reinserting links. When traffic is inspected, packets are traversing the network without an 802.1Q VLAN tag, required to reach the next hop within the domain.	8.31.01
In a Shortest Path Bridging VLAN (SPBV) domain, some multicast traffic, including statically programmed L2 multicast entries, loops in the network. The problem lies only in 7100-Series devices that may have ingress filtering disabled. The S-Series, K-Series, and 7100-Series, which flood certain multicast traffic on all internal ports, rely on the peer device to drop traffic with ingress filtering. If the peer 7100-Series device is affected by the problem, the device will not drop the traffic and loops form.	8.31.01
SPB convergence times may take longer than expected when region topology changes.	8.31.01
In a multi-slot bonded chassis, LAG port egress may not be set properly for an SPVID on a non-switch master blade. There is a small timing window where the distributed spanning tree port state information is missed.	8.31.01
Insertion or removal of a module in a bonded system can cause poor network convergence times as well as a temporary loss of traffic.	8.31.01
SPB devices may not agree topology agreement digest after changing master role.	8.31.01
When running spanning tree in SPB mode, traffic is lost when connected ports have differing configuration for SPB port status. One side sees the port as internal to the region while the other sees it as external. This results in a disputed BPDU status causing the port to remain in the listening state.	8.31.01
Traffic may not recover after disable/re-enable SPB.	8.31.01
An new root port for an SPT may forward before the old root port on a remote blade disables forwarding opening a transient loop.	8.31.01
When there is a change in the topology of the SPB region, ports might get stuck in the listening state.	8.31.01
Port may not become internal to the region even though ISIS adjacency is indicated.	8.31.01

7100-Series Customer Release Notes

Shortest Path Bridging Problems Corrected in 8.31.02.0014	Introduced in Version:
In a Shortest Path Bridging domain, when a device becomes the new regional root, designated ports on this new regional root go into listening state. Consequently, CIST traffic using this path is blocked. The issue is resolved by forcing a BPDU to be sent by the root port on the peer device.	8.31.01
In a Shortest Path Bridging VLAN (SPBV) domain, ports are incorrectly set to backup role and a state of blocking. The only ports affected are internal to the region and the consequence is limited network connectivity. Toggling the SPB configuration on the port may fix the problem, but not always.	8.31.01
For Software Bonded flows, from SPB ports, the first 4 bytes of the Software Bond Header is not getting removed properly, causing loss of L2 multicast traffic.	8.31.01
MVRP may propagate SPBV Base-VID registrations on ports within the SPBV domain.	8.31.01
System crashes when reboot one blade in a multi-blade system with message similar to: "<161>Oct 30 08:40:27 0.0.0.0 System[7]Chassis coherency timeout exceeded, resetting. delta:222000 curr:335186 nts:113186 nto:30000 hw:0x37000000 lnk:0x37000000 nv:0x37000000 img:0x37000000 max:0x37000000 (0x00e8535c 0x0071b18c 0x01ad4564 0xe8535353)".	8.31.01
In a Shortest Path Bridging VLAN domain, traffic loops are seen on directly adjacent 7100-series devices. Packet captures show that SPVID-tagged traffic egresses on ports that are not actually part of the VLAN egress membership. The problem is not seen if 7100-series devices are not connected to each other directly.	8.31.01
Port state may be listening for SPB internal port due to neighbor transmitting BPDUs with the agreeDigestValid flag persistently false.	8.31.01

Feature Enhancements in 8.31.01.0006

Feature Enhancements in 8.31.01.0006
<p>SPBV - – IEEE 802.1aq Shortest Path Bridging (SPB) provides data traffic a shortest cost path between any pair of switches in the SPB network. SPB features dynamic route calculation in a loop-free Layer-2 network and fast convergence time using IS-IS. The 7100-Series supports Shortest Path Bridging VLAN (SPBV).</p> <p>VRF - & scale info - Support for multiple VRFs has been added to the 7100-Series with this release. VRF provides a method of partitioning your network into different routed domains. A VRF is a segregated domain for the routed forwarding of packets. An interface configured to a particular VRF is considered a member of that VRF. VRFs can either be static or dynamic.</p> <p>Static VRFs employ only static or policy based routing.</p> <p>Dynamic VRFs employ dynamic routing protocols such as: OSPF, BGP, RIP, PIM, DVMRP, VRRP</p> <p>The default VRF is known as the Global Router and only interfaces assigned to the Global Router may be used to manage the device.</p> <p>VRF Route Leaking - Static Routing has been modified to allow routes to leak from a VRF to the Global Router and vice-a-versa.</p> <p>VRF Aware Policy Based Routing - Policy Based Routing has been modified to allow inter-vrf routing based on Route-Maps.</p> <p>VRF-Aware DHCP Relay - DHCP Relay has been modified to allow DHCP requests to be relayed either within a VRF or between a VRF and the Global Router.</p> <p>(VRF requires an advanced routing license)</p>

7100-Series Customer Release Notes

Feature Enhancements in 8.31.01.0006
IS-IS Graceful Restart - Graceful Re-Start for the IS-IS protocol has been added. Graceful Re-Start provides for an IS-IS router to continue to forward existing traffic and remain on the forwarding path during a restart of the IS-IS software process.
Remote Port Mirroring - The mirror source port is the source of the mirrored packets found on the local router of interest. The mirror encapsulates the L2 traffic seen by the mirrored source port and delivers it to the tunnel destination address.
Extended Transceiver Information Display - Extended Information display for supported transceivers is provided. In addition to serial number and model details, digital diagnostic information is displayed such as Temperature, Voltage, Transmit Current, Receive Power, Alarm State as well as High/Low thresholds.
Network Load Balanced Servers - Network load balancer or similar proprietary server NIC load balancing technologies, comprised of multiple physical machines responding to a single "virtual" IP address, expect the switch to flood its traffic to all ports on the destination VLAN using a static unicast or multicast MAC address.
100BASE-T Support on 71K91L4-24 and 71K91L4-48 10GBASE-T Ports – 100Mb speed option is now supported on 10GBASE-T ports.

Problems Corrected in 8.31.01.0006

802.1d Filter Database Problems Corrected in 8.31.01.0006	Introduced in Version:
MAC addresses that should age out from filter database will fail to do so. The frequency of this will increase with lower mac age times.	7.91.01
ACL Problems Corrected in 8.31.01.0006	Introduced in Version:
When a packet with a protocol other than IPv4 or IPv6 matches an L2 ACL, the L2 source and destination addresses will be displayed in place of the IPv4 and IPv6 addresses and the ethertype will be displayed as a hex value.	8.11.01
When an L2 ACL is applied to an interface, removed from an interface, or when an L2 ACL currently in use is modified, connections may not be removed. This can cause traffic to flow as it did before the change was made. Toggling the interface down then up will clear all connections and allow the L2 ACL to be correctly applied to traffic.	8.11.01
IPv6 Neighbor discovery messages may be dropped if IPv6 Ingress ACL's are applied.	8.21.01
Configuring unsupported access-group types to interfaces results in a confusing error message.	8.21.01
ARP Problems Corrected in 8.31.01.0006	Introduced in Version:
The router configured on a service provider switch may respond to ARPs received on a customer VLAN when the VLAN ID matches a router's interface VLAN ID. Conversely, the router configured on a customer switch may respond to ARPs received on a service provider VLAN when the VLAN ID matches a router's interface VLAN ID.	7.91.01
Using the command "clear arp <ipAddress>" may not function properly when clearing an ARP or ND entry in the stale state. If the host is still up a new ARP or ND entry will be added immediately after it is deleted.	7.91.01
BGP Problems Corrected in 8.31.01.0006	Introduced in Version:
BGP does not provide a CLI command to allow the user to specify a per peer local AS number.	7.91.01

7100-Series Customer Release Notes

BGP Problems Corrected in 8.31.01.0006	Introduced in Version:
If a BGP Update message is received with no NLRI path attribute the peering session is torn down.	7.91.01
CFM Problems Corrected in 8.31.01.0006	Introduced in Version:
CFM PDUs that contain the SenderID TLV will be improperly discarded as invalid frames.	8.21.01
Remote MEP states may be incorrect on CFM MEPs that have no VLAN configuration ("Port MEPs").	8.21.01
Data Center Bridging Problems Corrected in 8.31.01.0006	Introduced in Version:
The "show dcb cn ?" output shows "<cr>" as a valid option.	7.91.01
CN does not properly update the automatic alternate priority when a new CNPV is created with a value one less than an existing CNPV. The existing CNPV will continue to remap priorities to the new CNPV on ingress.	7.91.01
CPs on the same port will generate CNMs with the same CPID when multiple CPs exists on the 7100G-series.	8.21.01
The "set dcb cn congestion-point" configuration does not persist when multiple CNPVs are created in a bonded system.	7.91.01
The CLI set or clear "dcb cn congestion-point" command with a port-string of "*.*.*" will fail with an error similar to "Error: Failed to clear congestion point 5 for port tg.1.25". In a stacked system, all subsequent ports will not be set or cleared by the command.	8.21.01
The CN domain defense mode that is automatically configured by LLDP is not cleared when the LLDP neighbor ages out.	7.91.01
The CPID in the cp-mapping table may differ from the CPID in the CNM generated by the CP if the qp-index parameter is modified on the 7100G-series.	8.21.01
The ieee8021CnCpTransmittedFrames MIB object does not return the correct value for the number of frames transmitted on a CN queue. For congestion points corresponding to priorities 1, 2, or 3 the MIB object will return a value of 0. The ieee8021CnCpTransmittedFrames MIB object corresponds to the "Transmitted Frames" value in the "show dcb cn congestion-point" CLI.	8.20.02
The MIB supports setting the ieee8021CnCpQueueSizeSetPoint and ieee8021CnCpFeedbackWeight per ieee8021CnCpEntry, however the hardware does not support this parameter on a per CN queue basis. In the CLI, these objects are configured via the "set dcb cn congestion-point" command.	7.91.01
The min-sample setting for q-profile 0.1 does not persist.	7.91.01
The qp-index setting of the "set dcb cn congestion-point" CLI command does not appear in "show config" or "show config all".	7.91.01
Congestion point and queue profile settings do not display valid ranges for the min-sample and weight parameters in CLI help strings.	7.91.01
HostDos Problems Corrected in 8.31.01.0006	Introduced in Version:
Enabling the HostDoS portScan feature mistakenly filters inbound packets on port 22 when SSH is enabled. HostDoS should only filter these packets when SSH is disabled. This may render the switches SSH server inoperable, and the DoS attack detection logic may produce false positives. A workaround is to not enable HostDos portScan, or to enable it but with a	7.91.01

7100-Series Customer Release Notes

HostDos Problems Corrected in 8.31.01.0006	Introduced in Version:
relatively high portScan rate limit. Another workaround is to disable and then re-enable SSH (via a Telnet or console connection). However, the problem will return following a system reboot.	
IGMP Problems Corrected in 8.31.01.0006	Introduced in Version:
IGMP may lose track of where a flow entered the system. It may cause flow Interruption due to bad internal hardware programming.	7.91.01
It is possible for IGMP to lose track of which port a flow comes in, and cause an IGMP verify failed, status:0x00020000 message.	7.91.01
When the command "set igmp flow-wait" has both oper-state and time set on the same line, only the oper-state is set.	8.11.01
IP Interface Manager Problems Corrected in 8.31.01.0006	Introduced in Version:
When removing a Layer-3 interface using the "no <interfaceName>" command you may receive a difficult to decipher error message if the interface does not exist.	7.91.01
IPSLA Problems Corrected in 8.31.01.0006	Introduced in Version:
The SLA scheduler sub-mode command 'reset' cannot be entered while the SLA entry is scheduled. In order to reset the attributes for the entry, the user must stop the SLA entry via the 'stop' command in the SLA scheduler sub-mode.	8.01.01
The user will see the following CLI error when attempting to configure an SLA entry that had been previously configured in another VRF: ' Error: Command failed - create IpSla Entry ' The user will either have to remove the SLA entry from VRF in which it is configured, or choose a different SLA entry to configure.	8.11.01
Host Services Problems Corrected in 8.31.01.0006	Introduced in Version:
ICMP echo requests to IP interface addresses exceeding 100 per second will not all be answered.	8.20.02
"Unexpected syslog messages may be displayed if an interface is removed after the underlying vlan is cleared. These syslog messages are benign. Example of syslog messages: rtrHwApi[2.tRtrHwApi]ERROR: failed to update iif at index 591. rtrHwApi[2.tRtrHwApi]bcm_vlan_control_vlan_get(0, 591,..) failed."	8.21.01
Message "masterTrapSem time out, dropping trap" may appear in message log indicating an SNMP trap being dropped.	
"Blade may reset with the following log message after a configuration change: <1>NonVol[5.tNVolCU]cleanup:Remove() of first file on store=0, fileIndex=0 majorId=162 failed retval=3".	8.20.02
"Debug syslog message generated when an attempt to create a layer 3 interface is made with an out of range value. PiMgr[1.tConsole]generatelfIndex():retval=0;owner(0);mediaType(7);mediaPos(4096)".	7.91.01
Changing the owner string within an rmon command will result in a small memory leak.	7.91.01
"Failed to set -101" error is seen during logging configuration.	7.91.01

7100-Series Customer Release Notes

Host Services Problems Corrected in 8.31.01.0006	Introduced in Version:
"show support" or "debug messageLog message" result in an exhaustion of memory and a "memPartAlloc: block too big" message stored in the log.	7.91.01
"show system utilization slot <slot>" allows invalid slot numbers such as 0.	7.91.01
"Module might reset with message similar to "<1>DistServ[4.tDsBrdOk]serverWatchDog.1(Config), client 63(PEME) in recv for 6007 tics (0x00d0f9e4 0x0067b420 0x006707ac 0x01683264 0x00000000)" while PoE Controller is being updated."	7.91.01

Layer 1 Phy Problems Corrected in 8.31.01.0006	Introduced in Version:
Bonded 40G port with CR4 QSFP can potentially get into a link down condition when otherwise its link would be up. This can happen at bootup or any other link bounce condition.	7.91.01
Admin disabled 7100G-series tg ports do not bring link down with forcelinkdown enabled.	7.91.01
If nodealias is disabled on a given port and the maxentries value is set to default, after upgrading to firmware version 8.11.01 or newer will cause the maxentries value to be set to the previous default value.	8.11.01
POE may log a message similar to "bcPoE[4.tDSrecv5]bc_poeShutDown: Unable to get poeUpdateSemId" when a POE system is rebooted.	8.22.01

Layer 2 Problems Corrected in 8.31.01.0006	Introduced in Version:
CNM messages generated on a 7100G-series will be dropped if the reverse path is across a bond link.	8.22.02
Setting the mac age time to 10 seconds may cause the tNtpTmr task to use high amounts of CPU processing time.	8.21.01
"clear dcb cn priority <pri> lldp" will trigger a reset.	7.91.01
When GVRP adds a port to a VLAN that is not statically created, traffic will be dropped when not received on the same slot as the port added through GVRP.	7.91.01

L2 Multicast Problems Corrected in 8.31.01.0006	Introduced in Version:
It is possible for 7100-series modules to reset with the following message Machine Check exception Thread Name: tlgmplnp, at boot time, and may also get stuck in a constant reboot loop.	8.11.01
When setting IGMP setting for unknown input action to flood 7100-series does not flood the first packet.	7.91.01
IGMP may not properly send IGMP queries out interfaces on 7100 series product.	8.21.01

MVRP Problems Corrected in 8.31.01.0006	Introduced in Version:
Dynamic VLANs that were registered by MVRP may still show up in "show vlan" when there are no longer any egress ports. This can happen if the egress was registered on a module port that has since joined a lag.	7.91.01
Dynamic VLANs registered by MVRP fail resulting in no egress.	7.91.01
The ""show vlan"" command may show that egress on a port unexpectedly continues to be	7.91.01

7100-Series Customer Release Notes

MVRP Problems Corrected in 8.31.01.0006	Introduced in Version:
seen on a VLAN that once was dynamically registered by MVRP if the VLAN is configured statically on that port and then subsequently removed.	
Spanning Tree Problems Corrected in 8.31.01.0006	Introduced in Version:
Output from the command <code>show spantree blockedports</code> shows a port state of <code>Invalid</code> instead of <code>Disabled</code> . This error occurs when the port has the <code>dot1dStpPortEnable</code> value set to <code>disabled</code> and the port operstatus is up.	7.91.01
BPDUs are not processed when marked for discard by Policy. The port role and state will be designated forwarding. When the port is an inter-switch link and the attached port is designated forwarding, a loop will form if there is redundancy.	7.91.01
The <code>set spantree backuproot</code> command completes successfully but will not modify the value.	7.91.01
Layer 3 Problems Corrected in 8.31.01.0006	Introduced in Version:
The <code>age</code> column for the command <code>show ipv6 neighbors</code> displays the last time the ND entry was updated instead of the entry's age.	7.91.01
The <code>description cli</code> command is unavailable on a tunnel interface.	7.91.01
The following syslog message can be seen on 7100 series switches after a system reset has been issued. <code>trHwApi[1.tRtrHwApi]lock timeout warning. waited 10 seconds for the lock</code> . This message can be ignored as long as it occurs when the system is being reset.	8.21.01
When using VRRP fabric route mode, if a packet is sent to a host that is connected to the router that is in fabric-route mode (through the master router), the ARP response for that host will not make it back to the master router. This is because the ARP response will be consumed by the router in fabric route-mode.	7.91.01
Host routes for loopback interface addresses may not be distributed to all blades on a system reset causing connectivity issues to those addresses.	7.91.01
Port Jumbo MTU settings allowed for values below 1519.	8.01.01
Host routes advertised from the host-mobility routers are installed in other host-mobility peers that direct frames to the core instead of the directly connected networks.	8.21.01
IPv6 Neighbor Discovery (ND) Problems Corrected in 8.31.01.0006	Introduced in Version:
ARP/ND entries may expire early if the host does not respond to periodic ARP/ND refresh attempts.	8.21.01
It is possible to configure a Static ND entry which uses the same IP address as an interface address or VRRP address if the static ND entry is created before the other address.	8.21.01
The configuration commands <code>arp</code> and <code>ipv6 neighbor</code> allow invalid VLAN interfaces such as <code>vlan.0.4095</code> .	8.21.01
OSPF Problems Corrected in 8.31.01.0006	Introduced in Version:
If a config file saved prior to version 7.60 contains an OSPF passive interface, it will cause the box to hang if a configure is executed on an upgrade. The config file can be edited to format <code>vlan.0.#</code> instead of <code>vlan #</code> to allow upgrade.	8.22.02

7100-Series Customer Release Notes

OSPF Problems Corrected in 8.31.01.0006	Introduced in Version:
The "debug ip ospf packet" display for virtual interfaces reads "Interface not found for ifIndex 0".	8.21.01
When changing an OSPF network's area id then failing over, the original area ID is running seen in "show ip ospf interface", though the config reflects the new area ID.	8.21.01
With the removal of passive-interface default, the no passive-interface commands are removed, but they return on reboot of the router. They have no adverse effect.	8.21.01
If OSPF is configured to use a non-existent track object for cost, it does not calculate the cost based on the configured reference bandwidth but leaves it at default.	8.21.01

RIP and RIPng Problems Corrected in 8.31.01.0006	Introduced in Version:
If RIP is configured with passive interfaces and RIPng is configured, the passive-interfaces will function correctly but be displayed under RIPng.	8.21.01
When a RIPng interface is configured to be passive, the passive setting takes effect, but it is not displayed in show running.	8.21.01

VRRP Problems Corrected in 8.31.01.0006	Introduced in Version:
A VRRP router that owns the IP address may relinquish mastership if a packet is received from another VRRP router also claiming to the VRRP owner.	8.21.01
When a VRRP VRID is the master the "show ip vrrp" command will show the default "Master Advertisement Interval" when the correct value should match "Advertisement Interval" of the VRID (since it is the master).	8.21.01
When creating more than the maximum number of allowed VRRP critical IP addresses the error returned indicates that the IP address is bad when it should indicate that the maximum number of critical IP addresses already exists.	8.21.01
When removing a VRRP VRID from configuration the VIP may not be available to use on subsequent VRIDs if the command for the VIP address is negated just before the VRID is disabled.	8.21.01

COS Problems Corrected in 8.31.01.0006	Introduced in Version:
Setting cos IRL reference to a value greater than 15 causes the device to continuously reset.	7.91.01
If an invalid configuration is detected on upgrade the following syslog will display:	
<pre> SYSLOGX(kDbg_UPN,LOG_WARNING, ""CosTable unable to restore IRL "" ""reference %d mapping to resource %d "" ""for group %d.%d. Mapping is fixed for "" ""this product"" ,i,nvValue.ref[i],nvValue.group, nvValue.type); </pre>	
A change to the port configuration will prevent these messages from displaying after future reboots.	

7100-Series Customer Release Notes

Policy Problems Corrected in 8.31.01.0006	Introduced in Version:
Policy mac address rules may not be immediately applied to flows on Tunneled Bridge Ports.	8.21.01

KNOWN RESTRICTIONS AND LIMITATION:

10GBASE-T ports on 71K91L4-24 and 71K91L4-48 support 1Gb/10Gb speeds only. With 8.41.01, 100Mb port speed is not supported on 10GBASE-T ports.

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

The 7100-Series does not support half-duplex port configuration at any speed.

MACsec Limitations:

100Mb/1Gb SFP ports and 40Gb QSFP+ ports are not MACsec capable.

The MGBIC-02 copper 1Gb SFP transceiver cannot be used with MACsec enabled in SFP+ ports.

L2 MAC address aging could take up to 2x the desired MAC age time.

For SPBv: When changing the ISIS areaID, spb should be disabled before the change, and re-enabled after the new areaID is configured.

During an power down Machine Check and/or NonVol SysLog Messages may occur: These messages do not indicate a serious condition and may be ignored:

Example Machine Check SysLog Message

Message 52/128 Exception PPC750 Info 07.90.04.0000 02/23/2013 03:19:04

Exc Vector: Machine Check exception (0x00000200)

Thread Name: tPhyIntr

Exc Addr: 0x00c15588

Thread Stack: 0x073c9000..0x073c6000

Stack Pointer: 0x073c8e30

Traceback Stack:

[0] 0x00c10fb8

[1] 0x00c11104

[2] 0x00f86b6c

....

Example NonVol SysLog Message

Message 67/143 Syslog Message 07.90.04.0000 02/23/2013 03:16:36

<0>NonVol[1.tusrAppIntr]nonvol_init_dd: The persistent store for 0 is in complete. This data has been erased and the board will reset. (0x00b5a874 0x0092f644 0x007c06b4 0x011f90ac 0x00000000)

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

RFC STANDARDS SUPPORT:

RFC No.	Title
RFC0147	Definition of a socket
RFC0768	UDP
RFC0781	Specification of (IP) timestamp option
RFC0783	TFTP
RFC0791	Internet Protocol

7100-Series Customer Release Notes

RFC No.	Title
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
RFC0959	File Transfer Protocol
RFC1027	Proxy ARP
RFC1027	Using ARP - transparent subnet gateways
RFC1034	Domain Names - Concepts and Facilities
RFC1035	Domain Names - Implementation and Specification
RFC1157	Simple Network Management Protocol
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
RFC1191	Path MTU discovery
RFC1195	Use of OSI IS-IS for Routing in TCP/IP
RFC1213	MIB-II
RFC1245	OSPF Protocol Analysis
RFC1246	Experience with the OSPF Protocol
RFC1265	BGP Protocol Analysis
RFC1266	Experience with the BGP Protocol
RFC1323	TCP Extensions for High Performance
RFC1349	Type of Service in the Internet Protocol Suite
RFC1350	TFTP
RFC1387	RIPv2 Protocol Analysis
RFC1388	RIPv2 Carrying Additional Information
RFC1389	RIPv2 MIB Extension
RFC1492	TACACS+
RFC1493	BRIDGE- MIB
RFC1517	Implementation of CIDR
RFC1518	CIDR Architecture
RFC1519	Classless Inter-Domain Routing (CIDR)
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

7100-Series Customer Release Notes

RFC No.	Title
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
RFC2012	TCP-MIB
RFC2013	UDP-MIB
RFC2018	TCP Selective Acknowledgment Options
RFC2030	SNTP
RFC2080	RIPng (IPv6 extensions)
RFC2082	RIP-II MD5 Authentication
RFC2096	IP Forwarding Table MIB
RFC2104	HMAC
RFC2113	IP Router Alert Option
RFC2117	PIM -SM Protocol Specification
RFC2131	Dynamic Host Configuration Protocol
RFC2132	DHCP Options and BOOTP Vendor Extensions
RFC2233	The Interfaces Group MIB using SMIPv2
RFC2236	Internet Group Management Protocol, Version 2
RFC2260	Support for Multi-homed Multi-prov
RFC2270	Dedicated AS for Sites Homed to one Provider
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
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

7100-Series Customer Release Notes

RFC No.	Title
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
RFC2548	Microsoft Vendor-specific RADIUS Attributes
RFC2553	BasicSocket Interface Extensions for IPv6
RFC2577	FTP Security Considerations
RFC2578	SNMPv2-SMI
RFC2579	SNMPv2-TC
RFC2581	TCP Congestion Control
RFC2597	Assured Forwarding PHB Group
RFC2613	SMON-MIB
RFC2618	RADIUS Client MIB
RFC2674	P/Q-BRIDGE- MIB
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
RFC2787	VRRP MIB
RFC2796	BGP Route Reflection
RFC2819	RMON MIB
RFC2827	Network Ingress Filtering
RFC2858	Multiprotocol Extensions for BGP-4
RFC2863	IF-MIB
RFC2864	IF-INVERTED-STACK-MIB
RFC2865	RADIUS Authentication
RFC2866	RADIUS Accounting
RFC2869	RADIUS Extensions
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
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
RFC3273	HC-RMON-MIB
RFC3291	INET-ADDRESS-MIB
RFC3315	DHCPv6

7100-Series Customer Release Notes

RFC No.	Title
RFC3345	BGP Persistent Route Oscillation
RFC3359	TLV Codepoints in IS-IS
RFC3373	Three-Way Handshake for IS-IS
RFC3376	Internet Group Management Protocol, Version 3
RFC3392	Capabilities Advertisement with BGP-4
RFC3411	SNMP Architecture for Management Frameworks
RFC3412	Message Processing and Dispatching for SNMP
RFC3412	SNMP-MPD-MIB
RFC3413	SNMP Applications
RFC3413	SNMP-NOTIFICATIONS-MIB
RFC3413	SNMP-PROXY-MIB
RFC3413	SNMP-TARGET-MIB
RFC3414	SNMP-USER-BASED-SM-MIB
RFC3415	SNMP-VIEW-BASED-ACM-MIB
RFC3417	SNMPv2-TM
RFC3418	SNMPv2 MIB
RFC3446	Anycast RP mechanism using PIM and MSDP
RFC3484	Default Address Selection for IPv6
RFC3493	Basic Socket Interface Extensions for IPv6
RFC3509	Alternative Implementations of OSPF ABRs
RFC3513	RFC 3513 IPv6 Addressing Architecture
RFC3542	Advanced Sockets API for IPv6
RFC3562	Key Mgt Considerations for TCP MD5 Signature Opt
RFC3576	Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)
RFC3579	RADIUS Support for Extensible Authentication Protocol (EAP)
RFC3584	SNMP-COMMUNITY-MIB
RFC3587	IPv6 Global Unicast Address Format
RFC3590	RFC 3590 MLD Multicast Listener Discovery
RFC3595	Textual Conventions for IPv6 Flow Label
RFC3596	DNS Extensions to Support IP Version 6
RFC3621	POWER-ETHERNET-MIB
RFC3623	Graceful OSPF Restart
RFC3635	ETHERLIKE-MIB
RFC3678	Socket Interface Ext for Mcast Source Filters
RFC3704	Network Ingress Filtering
RFC3769	Requirements for IPv6 Prefix Delegation
RFC3787	Recommendations for Interop IS-IS IP Networks
RFC3810	MLDv2 for IPv6
RFC3879	Deprecating Site Local Addresses
RFC3956	Embedding the RP Address in IPv6 MCAST Address
RFC3973	Protocol Independent Multicast - Dense Mode (PIM-DM)
RFC3986	URI Generic Syntax
RFC4007	IPv6 Scoped Address Architecture
RFC4022	MIB for the Transmission Control Protocol (TCP)
RFC4109	Algorithms for IKEv1
RFC4113	MIB for the User Datagram Protocol (UDP)
RFC4133	ENTITY MIB
RFC4167	Graceful OSPF Restart Implementation Report
RFC4188	Bridge MIB

7100-Series Customer Release Notes

RFC No.	Title
RFC4193	Unique Local IPv6 Unicast Addresses
RFC4213	Basic Transition Mechanisms for IPv6
RFC4222	Prioritized Treatment of OSPFv2 Packets
RFC4264	BGP Wedgies
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
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
RFC4384	BGP Communities for Data Collection
RFC4443	ICMPv6 for IPv6
RFC4444	MIB for IS-IS
RFC4451	BGP MULTI_EXIT_DISC (MED) Considerations
RFC4456	BGP Route Reflection
RFC4486	Subcodes for BGP Cease Notification Message
RFC4541	IGMP Snooping
RFC4541	MLD Snooping
RFC4552	Authentication/Confidentiality for OSPFv3
RFC4560	DISMAN-PING-MIB
RFC4560	DISMAN-TRACEROUTE-MIB
RFC4560	DISMAN-NSLOOKUP-MIB
RFC4577	OSPF as PE/CE Protocol for BGP L3 VPNs
RFC4601	PIM-SM
RFC4602	PIM-SM IETF Proposed Std Req Analysis
RFC4604	IGMPv3 & MLDv2 & Source-Specific Multicast
RFC4607	Source-Specific Multicast for IP
RFC4608	PIM--SSM in 232/8
RFC4610	Anycast-RP Using PIM
RFC4632	Classless Inter-Domain Routing (CIDR)
RFC4668	RADIUS Client MIB
RFC4670	RADIUS Accounting MIB
RFC4673	RADIUS Dynamic Authorization Server MIB
RFC4724	Graceful Restart Mechanism for BGP

7100-Series Customer Release Notes

RFC No.	Title
RFC4750	OSPFv2 MIB
RFC4760	Multiprotocol Extensions for BGP-4
RFC4835	Crypto Algorithm Requirements for ESP and AH
RFC4836	MAU-MIB
RFC4836	IANA-MAU-MIB
RFC4861	Neighbor Discovery for IPv6
RFC4862	IPv6 Stateless Address Auto-configuration
RFC4878	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
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
RFC5132	IP Multicast MIB
RFC5176	Dynamic Authorization Extension to RADIUS
RFC5186	IGMPv3/MLDv2/MCAST Routing Protocol Interaction
RFC5187	OSPFv3 Graceful Restart
RFC5240	PIM Bootstrap Router MIB
RFC5250	The OSPF Opaque LSA Option
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
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
RFC5643	OSPFv3 MIB
RFC5798	Virtual Router Redundancy Protocol (VRRP) V3
RFC6164	Using 127-Bit IPv6 Prefixes on Inter-Router Links
RFC6296	IPv6-to-IPv6 Network Prefix Translation
RFC6329	IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging
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)

7100-Series Customer Release Notes

RFC No.	Title
Drafts	draft-ietf-isis-ipv6-te (Partial Support)
Drafts	draft-ietf-ospf-ospfv3-mib
Drafts	draft-ietf-ospf-te-node-addr
Drafts	draft-ietf-idmr-dvmrp-v3-11
Drafts	draft-ietf-rrp-unified-spec-03.txt

EXTREME NETWORKS PRIVATE ENTERPRISE MIB SUPPORT:

Title	Title	Title
CISCO-CDP-MIB	ENTERASYS-IF-MIB-EXT-MIB	ENTERASYS-SPANNING-TREE-DIAGNOSTIC-MIB
CISCO-TC	ENTERASYS-JUMBO-ETHERNET-FRAME-MIB	ENTERASYS-SYSLOG-CLIENT-MIB
CT-BROADCAST-MIB	ENTERASYS-LICENSE-KEY-MIB	ENTERASYS-TACACS-CLIENT-MIB
CTIF-EXT-MIB	ENTERASYS-LICENSE-KEY-OIDS-MIB	ENTERASYS-TRANSMIT-QUEUE-MONITOR-MIB
CTRON-ALIAS-MIB	ENTERASYS-LINK-FLAP-MIB	ENTERASYS-UPN-TC-MIB
CTRON-BRIDGE-MIB	ENTERASYS-MAC-AUTHENTICATION-MIB	ENTERASYS-VLAN-AUTHORIZATION-MIB
CTRON-CDP-MIB	ENTERASYS-MAC-LOCKING-MIB	ENTERASYS-VLAN-INTERFACE-MIB
CTRON-CHASSIS-MIB	ENTERASYS-MAU-MIB-EXT-MIB	IANA-ADDRESS-FAMILY-NUMBERS-MIB
CTRON-ENVIROMENTAL-MIB	ENTERASYS-MGMT-AUTH-NOTIFICATION-MIB	IEEE8021-CN-MIB
CTRON-MIB-NAMES	ENTERASYS-MGMT-MIB	IEEE8021-PAE-MIB
CTRON-OIDS	ENTERASYS-MIB-NAMES DEFINITIONS	IEEE8021-PFC-MIB
CTRON-Q-BRIDGE-MIB-EXT	ENTERASYS-MSTP-MIB	IEEE8023-LAG-MIB
ENTERASYS-AAA-POLICY-MIB	ENTERASYS-MULTI-AUTH-MIB	LLDP-EXT-DOT1-MIB
ENTERASYS-CLASS-OF-SERVICE-MIB	ENTERASYS-MULTI-USER-8021X-MIB	LLDP-EXT-DOT3-MIB
ENTERASYS-CONFIGURATION-MANAGEMENT-MIB	ENTERASYS-OIDS-MIB DEFINITIONS	LLDP-EXT-MED-MIB
ENTERASYS-CONVERGENCE-END-POINT-MIB	ENTERASYS-PFC-MIB-EXT-MIB	LLDP-MIB
ENTERASYS-CN-MIB-EXT-MIB	ENTERASYS-POLICY-PROFILE-MIB	RSTP-MIB
ENTERASYS-DIAGNOSTIC-MESSAGE-MIB	ENTERASYS-PWA-MIB	U-BRIDGE-MIB
ENTERASYS-DNS-RESOLVER-MIB	ENTERASYS-RADIUS-ACCT-CLIENT-EXT-MIB	USM-TARGET-TAG-MIB
ENTERASYS-IEEE8023-LAG-MIB-EXT-MIB	ENTERASYS-RADIUS-AUTH-CLIENT-MIB	SNMP-RESEARCH-MIB
ENTERASYS-IETF-BRIDGE-MIB-EXT-MIB	ENTERASYS-RESOURCE-UTILIZATION-MIB	VSB-SHARED-SECRET-MIB
ENTERASYS-IETF-P-BRIDGE-MIB-EXT-MIB	ENTERASYS-SNTP-CLIENT-MIB	ENTERASYS-DOT3-LLDP-EXT-MIB
ENTERASYS-IEEE8021-CFM-EXT-MIB	ENTERASYS-IEEE8021-CFM-EXT-MIB	
ENTERASYS-OSPF-EXT-MIB	ENTERASYS-PIM-EXT-MIB	ENTERASYS-DVMRP-EXT-MIB
ENTERASYS-ETH-OAM-EXT-MIB	ENTERASYS-RIPv2-EXT-MIB	ENTERASYS-ENTITY-SENSOR-MIB-EXT-MIB
IEEE8021-SECY-MIB		

7100-Series Customer Release Notes

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

SNMP TRAP SUPPORT:

RFC No.	Title
RFC 1493	New Root Topology Change
RFC 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 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
Enterasys-link-flap-mib	etsysLinkFlapViolation
Enterasys-ietf-bridge-mib-ext-mib	etsysletfBridgeDot1qFdbNewAddrNotification etsysletfBridgeDot1dSpanGuardPortBlocked etsysletfBridgeDot1dBackupRootActivation etsysletfBridgeDot1qFdbMovedAddrNotification etsysletfBridgeDot1dCistLoopProtectEvent
Enterasys-notification-auth-mib	etsysMgmtAuthSuccessNotificiation etsysMgmtAuthFailNotificiation

7100-Series Customer Release Notes

RFC No.	Title
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	etsysCosIriExceededNotification
Enterasys-policy-profile-mib	etsysPolicyRulePortHitNotification
Enterasys-mstp-mib	etsysMstpLoopProtectEvent
Ctron-environment-mib	chEnvAmbientTemp chEnvAmbientStatus

RADIUS ATTRIBUTE SUPPORT:

This section describes the support of RADIUS attributes on the 7100-Series. 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

7100-Series Customer Release Notes

RADIUS ACCOUNTING ATTRIBUTES:

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

GLOBAL SUPPORT:

By Phone: 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/contact/

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.