

[Home](#)[Products and Services](#)[Training](#)[Support](#)[Partners](#)[Company Info](#)[Advanced Search](#)[Site Map](#)You are Here: [Global Services](#) > [relnotes](#) > 6H308-24 & 6H308-48 version 4.08.12

6H308-24 & 6H308-48
Firmware Version 4.08.12
January 2001

- ✂ [INTRODUCTION](#)
- ✂ [Firmware Specification](#)
- ✂ [HARDware compatibility](#)
- ✂ [Network Management Software Support](#)
- ✂ [SUPPORTED FUNCTIONALITY](#)
- ✂ [Installation and Configuration Notes](#)
- ✂ [Firmware Changes and Enhancements](#)
- ✂ [Known Restrictions and Limitations](#)
- ✂ [Compliance support](#)
- ✂ [Cabletron Private Enterprise MIB Support](#)
- ✂ [Cabletron Private Enterprise trap Support](#)
- ✂ [INTRODUCTION](#)
- ✂ [Firmware Specification](#)
- ✂ [HARDware compatibility](#)
- ✂ [BootPROM compatibility](#)
- ✂ [Network Management Software Support](#)
- ✂ [SUPPORTED FUNCTIONALITY](#)
- ✂ [Installation and Configuration Notes](#)
- ✂ [Firmware Changes and Enhancements](#)
- ✂ [Supported Functionality section.](#)
- ✂ [Known Restrictions and Limitations](#)
- ✂ [Compliance support](#)
- ✂ [Cabletron Private Enterprise MIB Support](#)
- ✂ [Cabletron Private Enterprise trap Support](#)

INTRODUCTION:

This document provides specific information for version 4.08.12 of firmware for the 6H308-24 & 6H308-48. This release supports 802.1D/Q and SecureFast 2.00.27 switching. The following hardware platforms are supported by this firmware release:

2E253-49R	2H252-25R	2H253-25R	2H258-17R
6E233-49	6H202-24	6H203-24	2H252-25RDC
6H253-13	6H258-17	6H259-17	6H252-17
6H262-18	6H302-48	6H303-48	6G306-06
	6H308-24	6H308-48	

It is recommended that one thoroughly review this release note prior to the installation or upgrade of this product.

Firmware Specification:

Status	Version No.	Type	Release Date
Current Version	4.08.12	Customer	January 2001
Previous Version	4.07.09	Customer	October 2000
Previous Version	4.06.05	Customer	July 2000
Previous Version	4.04.19	Customer	May 2000
Previous Version	4.04.14	Customer	April 2000
Previous Version	4.02.10	Customer	February 2000

Previous Version	4.01.08	Customer	January 2000
Previous Version	4.00.08	Customer	August 1999
Previous Version	3.11.04	Customer	July 1999
Previous Version	3.10.07	Customer	June 1999
Previous Version	3.05.07	Customer	March 1999
Previous Version	2.00.17	Customer	November 1998
Previous Version	1.03.10	Customer	September 1998
Previous Version	1.01.10	Customer	July 1998

 HARDWARE compatibility:

This version of firmware is supported on all Hardware revisions.

Boot PROM compatibility

This version of firmware supports all Boot PROM versions.

 Network Management Software Support:

NMS Platform	Version No.	Module No.
SPECTRUM	5.0	Rev. 1
SPMA (Spectrum Portable Management Application)	3.2	Rev. 1
SPEL (Spectrum Element Manager)	2.02.00	N/A
SmartSwitch ATM Administrator	2.3.16	N/A
NetSight Element Manager	2.2.1	N/A

If you install this image, you may not have control of all of 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 management platform for details.

 SUPPORTED FUNCTIONALITY:

Features Summary

3rd Generation Module Proxy Function for Matrix E7

Multi-Layer Frame Classification

Switch Configuration Upload/Download

IGMP Version 1 and 2 Support

802.1D Traffic Class Expediting (previously known as 802.1p Traffic Management/ Dynamic Multicast Filtering)

802.1Q VLAN tagging and identification

GARP protocol support including GMRP and GVRP

SNMP/IP access control lists

Support for the following HSIMs:

HSIM-FE6 HSIM-F6

HSIM-G01 HSIM-G09

HSIM-W6 HSIM-W84

HSIM-A6DP HSIM-SSA710

HSIM-SSR600

HSIM-W85 HSIM-W87

Support for the following VHSIMS:

VHSIM-G6 VHSIM-A6DP

VHSIM-G02 VHSIM2-A6DP

IP Fragmentation

Local Management via TELNET, HTTP, RS232

Port Mirroring

RMON

Runtime Address Discovery

Runtime Image Download

SmartTrunk Link Aggregation

Rate Limiting

WebView Management

 Installation and Configuration Notes:

In general, the product will be shipped to you pre-configured with this version of firmware. If you would like to upgrade an existing product, please follow the TFTP download instructions that are included with your Users Guide. TFTP download instructions are also available on the Enterasys Support web site at:

<http://www.enterasys.com/support/techtips/tk0020-9.html>

If you are downloading this firmware to a module or modules operating in Distributed Chassis Management Mode within a SmartSwitch 6000*, and you wish to use the chassis IP address, you must be connected to one of the front panel ports of the module you are downloading. In order to download to the chassis IP address the community name string must include the slot number. For example, to download a module in slot one, a community name of public.1 would be used (assuming the super-user community name was public). Non-Runtime downloads to the chassis IP address are not allowed. Please consult your Users Guide for more details on performing firmware downloads.

*Distributed Chassis Manager is not supported on the Matrix E7

Modules running this firmware version default to Distributed Chassis Management mode when installed into a SmartSwitch 6000**. An individual IP address does not need to be assigned to each module if a chassis IP address has been assigned. The management mode field in the General Configuration Screen of Local Management allows a user to select whether a module will operate in Distributed or Standalone mode. Changing this value will cause the module to reboot. **Distributed Chassis Manager is not supported in SecureFast mode.

Distributed Chassis Management (DCM) is not supported in the Matrix E7 chassis.

Module Generations and Chassis Definitions

Part #	Description
6C107	7-slot Matrix E7 chassis
6C105	5-slot SmartSwitch 6000 chassis
6X1XX	1st generation module, SmartSwitch 6000 family
6X2XX	2nd generation module, SmartSwitch 6000 family
6X3XX	3rd generation module, SmartSwitch 6000/Matrix E7 family
2X2XX	2nd generation standalone switch, SmartSwitch 2000 family

note: where X is a wildcard

The second and third generation modules default to an 802.1Q operational mode. In order for the modules to maintain backward compatibility with first generation SmartSwitch 6000 modules, which default to an 802.1D operational mode, the backplane ports do not insert an 802.1Q Frame Tag when forwarding traffic to other modules. If 802.1Q VLANs are configured on modules, then the backplane ports will need to be configured by the user as 802.1Q Trunk ports for all second generation modules. The first generation modules automatically configure their backplane ports as 802.1Q when changed from the default operational mode to the 802.1Q VLAN mode.

SmartTrunk

This version of firmware supports SmartTrunk, Cabletrons Link Aggregation feature. Below is a matrix that details the other SmartSwitch platforms and the firmware versions required to operate properly with this version of firmware.

Product Line	Firmware Version Required
First Generation SmartSwitch 2000/6000	4.05.09 or higher
SmartSwitch 9000 9X5XX series	1.01.10 or higher
SmartSwitch 9000 9X4XX series	1.11.08 or higher
SmartSwitch Router (SSR-2000, 8000, and 8600)	2.0 or higher

SmartTrunk with Non-Enterasys products

SmartTrunk has been successfully tested with some other vendors Link Aggregation methods. It is usually necessary to disable the SmartTrunk link protocol for proper operation. This is configurable in the SmartTrunk Local Management screen. When the SmartTrunk link protocol is disabled, some of the automated configuration protection is lost. Users must be very careful to observe the SmartTrunk configuration rules detailed in the SmartTrunk Users Guide. Failure to follow these configuration rules could result in unstable network operation. This guide is available at:
<http://www.enterasys.com/support/manuals> .

SecureFast

If you are mixing 2nd and 3rd generation modules in a Matrix E7 while running in SecureFast mode, no modules can be inserted in slots 6 and 7 of the Matrix E7 chassis. If you are running all 3rd generation modules in the Matrix E7 there are no slot restrictions.

Cabletron Discovery Protocol

By default the SmartSwitch products periodically transmit multicast CDP frames, these frames are used by network management products, and other switches to discover the topology of the network. CDP frames are sent by the switches using a multicast destination address of 01-80-C2-00-00-21. This feature can be disabled using the CDP command in Local Management.

System Interface Numbering

The System interfaces are numbered by MIB II as shown:

- 1 - Front Panel switch port 1 interface
- 2 - Front Panel switch port 2 interface
- 3 - Front Panel switch port 3 interface
- . .
- . .
- n - Last Front Panel switch port n interface
- n+1 - FTM1 slot/port 1 switch interface
- n+2 - FTM1 slot/port 2 switch interface
- . .
- . .
- n+6 - FTM1 slot/port 6 switch interface
- n+7 - Host in-band management interface

```
n+8 - SmartTrunk group 1
. .
. .
n+13 - SmartTrunk group 6
```

Source Address Table

The default size of Source Address Table (SAT) has been changed from 16,000 entries to 8,000 entries. The SAT size can be changed back to 16000 entries; however, the number of classification entries allowed is reduced. With the SAT size set to 8000 entries, up to 500 VLAN classifications can be created, and up to 500 priority classifications can be created. With the SAT size set to 16000 entries up to 200 VLAN classifications can be created, and up to 200 priority classifications can be configured. For more detailed information about SAT size and classifications refer to the Users Guide <http://www.enterasys.com/support/manuals/2-6.html>

Matrix E7 Proxy Function

The 3rd generation modules (6X3XX) provide a function called "proxy". This function provides 1st and 2nd generation modules (which have only 4 backplane ports that connect to slots 1-5) with the ability to communicate with 3rd generation modules installed in slots 6 and 7, by way of an intermediate (proxy) 3rd generation module installed in one of the first 5 slots.

A packet entering a port on a 1st or 2nd generation module in one of the first 5 slots is transmitted to the FTML backplane connection to the 3rd generation "proxy" module in one of the first five slots. This proxy module can then forward that packets out its appropriate backplane port to either slots 6 or 7.

The lowest numbered slot containing a 3rd generation module will always be the "proxy" module. The function moves dynamically if a new lower slot 3rd generation module is installed or removed.

Since the "proxy" function does consume switching resources, switching performance will be reduced on the module. For this reason, it is not recommended that an uplink module (like the 6G306-06) perform the "proxy" function in order to avoid performance degradation of the uplink. This is avoided by simply installing uplink modules in the higher numbered slots and using other 3rd generation modules in a lower numbered slot to perform the function.

ATM Specific

In order for the HSIM-A6DP, VHSIM-A6DP or VHSIM2-A6DP to operate with this version of firmware, the optional memory upgrade kit must be installed in the SmartSwitch 2000/6000. The Enterasys part number for this kit is

SS-16M-DRAM-UGK.

This version of firmware supports the ATM Forums standards for LAN Emulation version 1.0 and 2.0 as well as UNI 3.0, 3.1, and 4.0 Signaling, and Traffic Management 4.0 (TM). The VHSIM-A6DP, VHSIM2-A6DP, and HSIM-A6DP module are an 802.3 LAN Emulation (LEC) device and will operate in an 802.3 Emulated LAN using an MTU of 1516 bytes or 1580 bytes.

The VHSIM-A6DP, VHSIM2-A6DP and HSIM-A6DP support 32 LECs when operating in 802.1Q mode, and 16 LECs when operating in SecureFast mode.

For SVC operation, LAN Emulation Services must be installed on the ATM Network. At a minimum, there must be an 802.3 LAN Emulation Server (LES) and a Broadcast and Unknown Server (BUS). Additionally, there should be a LAN Emulation Configuration Server (LECS) on the ATM network. If there is not an LECS on the ATM network, the ATM address of the LES must be manually entered on the VHSIM-A6DP, VHSIM2-A6DP or HSIM-A6DP through Local Management.

On power-up, the VHSIM-A6DP, VHSIM2-A6DP/HSIM-A6DP will automatically join a default ELAN provided by ATM Services. The VHSIM-A6DP, VHSIM2-A6DP and HSIM-A6DP will contact the LECS using the Configured ATM address of the LECS. If an LECS ATM Address has not been configured (via Local Management), the VHSIM-A6DP, VHSIM2-A6DP, AND HSIM-A6DP will use the ATM Address provided by

Interim Local Management Interface (ILMI) to set up a connection to the LAN Emulation Configuration Server (LECS). If there is no response via ILMI, the VHSIM-A6DP, VHSIM2-A6DP, AND HSIM-A6DP will set up a connection to the well-known ATM address of the LECS. If a connection is not successful to the well-known ATM address, the VHSIM-A6DP, VHSIM2-A6DP, AND HSIM-A6DP will use the well-known VC of VPI=0, VCI=17 to connect to the LECS.

This firmware image supports Permanent Virtual Circuits via Local Management or the IETF ATOM MIB (RFC-2515). When PVCs are used, LAN traffic is encapsulated over ATM using RFC 1483 VC-based Multiplexing or LLC Encapsulation for Bridged Protocols.

HSIM-A6DP Specific

This version of firmware supports up to 2040 SVCs on the HSIM-A6DP using LANE 1.0/2.0 and UNI 3.0/3.1/4.0 signaling. In 802.1Q mode 64 PVCs are allowed. In SecureFast mode 32 PVCs are allowed. Each PVC is treated as a bridge port.

This firmware image provides support for Bandwidth Allocation over PVC's. When creating PVC's via Local Management, the user can specify a peak bandwidth allowed (in Mbps) to be transmitted on each PVC. This feature is available for PVC's only. SVCs are not allowed when using this feature.

This version of firmware provides support for the HSIM-A6DP Traffic Management Daughter-card (HSIM-TM-UGK and HSIM-A6DP-TM). When the TM card is installed, the HSIM-A6DP supports the following ATM Forum Traffic Management 4.0 traffic classes:

- o UBR
- o ABR (Explicit Rate and Relative Rate)
- o nrt-VBR
- o CBR

This version of firmware provides support for up to 64 Traffic Descriptors, which can be assigned to any of the aforementioned classes. A single Traffic Descriptor can be used for each PVC on the HSIM-A6DP. Traffic Descriptors (up to eight) are assigned to a LEC or virtual port based on 802.1p priority and transmit queue mapping. Traffic Management will carry priority across the ATM fabric for the given QoS level.

VHSIM-A6DP and VHSIM2-A6DP Specific

This version of firmware supports up to 8192 SVCs on the VHSIM-A6DP using LANE 1.0/2.0 and UNI 3.0/3.1/4.0 signaling. The number of PVCs allowed in 802.1Q mode is 64, in SecureFast mode 32 PVCs are allowed. Each PVC is treated as a bridge port.

This version of firmware provides support for the following ATM Forum Traffic Management 4.0 traffic classes:

- o UBR
- o nrt-VBR
- o CBR

This version of firmware provides support for up to 64 Traffic Descriptors, which can be assigned to any of the aforementioned classes. A single Traffic Descriptor can be used for each PVC on the VHSIM-A6DP. Traffic Descriptors (up to eight) are assigned to a LEC or virtual port based on 802.1p priority and transmit queue mapping. Traffic Management will carry priority across the ATM fabric for the given QoS level. Please refer to the users guide for proper utilization of this feature.

HSIM and VHSIMs

For the following HSIM/VHSIM modules to operate, the firmware revisions listed below must be installed and running on the SmartSwitch 2000/6000 prior to installing the HSIM/VHSIM. If this firmware is not running prior to installing the HSIM/VHSIM, the SmartSwitch may not boot-up:

VHSIM-G6 1.03.07 or higher

HSIM-A6DP 1.03.11 or higher

VHSIM-A6DP 3.05.07 or higher

VHSIM2-A6DP 4.03.06 or 4.06.05 and higher

Firmware Changes and Enhancements:

Changes made in 4.08.12

Added support for the 6H308-24 & 6H308-48

Added support for SNMP/IP access control lists

Added embedded Radius Client. This provides additional security for the switchs management entity. For details on configuring this feature refer to the latest User Guide located at <http://www.enterasys.com/support/manuals/>

Fixed a problem in which a soft-reset of an ELS box connected Via fiber Ethernet link to a second generation SmartSwitch module causes the SmartSwitch to reset.

Non Ethernet II frames are now translated correctly in configurations with the HSIM-F6.

Fixed a Node Alias Table initialization problem suspected of causing occasional resets.

Fixed a problem in which the SFS version number was not being displayed on the LM password screen for "Enterasys" modules (as opposed to "Cabletron" modules) or for any module installed in a Matrix E7 chassis.

Fixed a problem in which an ATM VHSIM with APIM-67 does not forward traffic at the proper rate according to the configuration.

Fixed a problem in which priority information is not maintained in a frame's Q tag when port redirection and priority classifications are configured together.

Fixed a problem in which the ifSpeed object would return an incorrect value for 10/100 ports operating at 10 Mb/sec on 3rd generation blades.

Fixed a reset problem seen when "status" changes (i.e., link up/down, new speed, new duplex) occur on the 10/100 ports of a 3rd generation blade under severe traffic load.

Modified "pvc_mesh" CLI command operation to properly interact with IGMP and flood reports out all ATM PVCs. Modified IGMP to provide support for dynamically removing ports from its flood mask (in the event that a PVC goes away) and for reporting the flood reports mask. Fixes connectivity problems with pvc_mesh and IGMP.

Fixed a port LED display problem seen on VHSIM-G6.

Removed check for number of Fast Ethernet ports when determining whether to allow any Gigabit Ethernet ports to be dual activated. Now, any Gig ports may always be dual-activated

Changes made in 4.07.09

SecureFast version 2.00.22 has been added.

SecureFast is now operational in the Matrix E7 chassis

The ability to re-write the TOS field of an IP packet has been added to the frame classification feature.

TCP and or UDP port (socket) ranges can now be specified when configuring Layer 4 VLAN and Priority Classification rules. This feature is especially useful when classifying traffic patterns that does not use a constant socket number.

An SNMP trap is now generated if the switch detects the presence of a loopback condition when the loopback_detect feature is enabled. The

enterprise trap numbers used are 418 (port up) and 419 (port down). The link LED is also toggled off during the loopback condition and back on when the loopback condition goes away.

The ability to enable and disable Telnet and WebView access to the switch (via the General Configuration LM screen) has been added.

Two new commands have been added to the Network Tools Local Management screen, `timed_reset` and `timed_soft_reset`. These features allow the switch to be automatically reset at a specified time interval.

Fixed a problem (introduced in 4.06.05) in which the default Trunk Protocol (PLAP) state of newly-created SmartTrunk would appear to be disabled (as indicated by the trunk configuration Local Management screen), but would actually be enabled.

The interface speed for FTM backplane ports is now correctly reported as 2.1Gb/s

The `ctIfconnectiontype` oid from Cabletron MIBII Extensions now returns the correct value for the 6G306-06.

Fixed an issue with the VHSIM2-A6DP that would cause the host platform to reset under certain conditions.

Fixed a problem where multicast priority on ATM PVCs was being handled incorrectly.

Fixed a problem in which WebView would not function correctly on modules operating in "standalone" mode in a 6C105 chassis.

Fixed a reset problem seen when polling `ctUPS` MIB objects on behalf of APC 700/1400 UPS models.

Using the `non_bridge_if_num` command the device can now be configured to report non-bridge ports, such as the SNMP host port with a MIB II `ifIndex` value of either 0 or 9999. The default value is 0. Some third party network management applications do not operate properly if a value of 0 is reported.

Automatic End System/Alias Discovery
 - Switches maintain a directory of locally attached users. Supported protocols learned in this release are:

ENET

Inet ARP

Inet IP

Inet UDP

Inet OSPF

IP RIP

IP BGP

IGRP

ATalk RTMP

ATalk ARP

NetBIOS IPX

IPX RIP

IPX SAP

IPX Type 20

802.1q VLAN Tag

Fixed a problem of Spanning Tree not working properly on ATM SVC ports.

Added support for ATM VHSIM2 in SFS mode.

Modified SFS dispatcher to log a debug message upon receipt of a frame whose RX descriptor contains an invalid physical port number; the frame is then discarded. Previously, a reset of the device would be initiated.

Fixed a problem of interface statistics decrementing on gigabit Ethernet ports.

Fixed a problem on ATM PVC redundancy failover not working when operating in SFS mode.

Fixed a memory leak in RMON that would cause occasional "Out of Heap" resets.

Known Restrictions and Limitations:

Rate Limiting is not supported over ATM interfaces or WAN VC interfaces (i.e., HSIM-W87/W85) in this release of firmware.

NVRAM must be cleared on a 6H302-48 module after upgrading from version 4.02.10 to this release because of interface numbering changes between firmware revisions.

1st and 2nd generation modules (6X1XX & 6X2XX series) are restricted to installation into slots 1-5 of the Matrix E7. In addition, if backplane communication from these modules in slots 1-5 to 3rd generation modules within slots 6 and 7 is desired, a 3rd generation module (6X3XX) is required to be installed within the first 5 slots to perform the proxy function. See Matrix E7 Proxy Function on page 5 of these notes.

The Matrix E7 chassis does not support Distributed Chassis Management.

If you are mixing 2nd and 3rd generation modules in a Matrix E7 while running in SecureFast mode, no modules can be inserted in slots 6 and 7 of the Matrix E7 chassis. If you are running all 3rd generation modules in the Matrix E7 there are no slot restrictions.

Port Redirect cannot be used to redirect frames between modules that use the proxy function within the Matrix E7 for their connectivity. Specifically, frames cannot be redirect from 1st or 2nd generation modules in slots 1-5 to a 3rd generation module in slots 6-7, and vice-versa. Note that frames may be redirected between any two 3rd generation modules in the chassis.

Using local management to configure 100 Mb/s Full Duplex via "Save to all ports" command on a module with a VHSIM-G6 or VHSIM-G02 installed, may cause constant spanning tree topology changes to occur. This will only occur on devices with a VHSIM-G6 or VHSIM-G02 installed. This issue can be avoided by configuring the ports individually.

BOOTP firmware download will not function over an 802.1Q trunk port.

The Port Path Cost spanning tree parameter for ports configured as "SmartTrunk" ports is not saved in non-volatile memory

The Max Age parameter in spanning tree is not saved in non-volatile memory.

Because of hardware limitations, undersize Ethernet frames received on front-panel ports of 3rd generation modules will never be passed into the switch. Consequently, such frames can never be captured by RMON or redirected as error frames.

During a Spanning Tree topology change, multiple Topology Change Notifications (TCNs) are generated which can lead to increased convergence times.

Hot swapping a 3rd generation module into a lower numbered slot, than the current Matrix E7 proxy module, will cause the proxy function to move to the

newly inserted module. This will cause a temporary loss of connectivity for 1st and 2nd generation modules communicating to slots 6 and 7. This does not happen with removal of the current proxy module causing the proxy function to move to the next 3rd generation module in a higher slot number.

There can be power redundancy issues when installing 6H302-48 or 6H303-48 modules in a SmartSwitch 6000 chassis. The 6C205-1 power supply provides 60 Amps and the 6C205-3 provides 90 Amps. The recommended maximum using the 6C205-1 is two modules (6H30x-48) and four when using the 6C205-3. Please contact Technical Support for assistance in determining which configurations have power redundancy and those that do not. The redundancy LEDs on the power supplies indicate the status of power redundancy. When both LEDs are green, the chassis has redundant power. The Matrix E7 does not have these power redundancy issues.

User configured VrrpPort settings in Local Management will be lost after downloading this image and resetting the switch. The settings must be manually reentered.

By default only one port is active on the VHSIM-G6 and 6H262-18. To activate the second port you must use the "gigabit_port_mode" option in the Network Tools Local Management screen. When the second port is activated, the switch will reset, and the contents of NVRAM with the exception of IP Address and Subnet Mask will be lost.

If the device is running boot code version 2.01.00 or lower and power-up sequence is interrupted, or if optional hardware is installed or removed, the device may run an extended diagnostics sequence. During the extended diagnostics, the CPU LED color will be solid amber. This sequence may take between five to forty minutes to complete depending on the hardware model.

Opening multiple active telnet sessions to a single module to access local management may cause the module to reset. This will be corrected in a future firmware release.

The SmartSwitch host (SNMP management) port must not be configured as an 802.1Q Trunk Port. If configured as an 802.1Q Trunk Port the device may reset. If the user wishes to make the host port accessible from some or all 802.1Q VLANs, then the 802.1D Trunk Port option must be used.

If an optional HSIM or VHSIM is installed or removed, the device will restore all configuration settings to the factory defaults. The only customer configuration settings that will be saved are the IP Address, Subnet Mask and the Operational Mode.

When upgrading from firmware version 2.00.17 to this version of firmware, the speed and duplex settings may not properly restored from NVRAM. In the event this occurs, you will need to reconfigure the speed and duplex settings if something other than the default setting is required. This issue only affects devices running firmware version 2.00.17.

If HSIM-F6 ports are configured to operate in full duplex mode, the Local Management Ring Map screen may not operate correctly.

The HSIM-F6 should not be configured as an 802.1Q "trunk" port.

The HSIM-F6 does not support 802.1Q tagged frames which exceed the maximum Ethernet frame length (1522 bytes). Frames larger than 1522 bytes will be discarded by the switch.

The HSIM-F6 was developed in the early stages of the 802.1Q standards process. The Ethernet & FDDI translation and frame tagging rules were not fully defined at development time. Changes made during the standardization process, result in the HSIM-F6 using a proprietary method of frame tagging on FDDI. The implementation used on the HSIM-F6 is fully functional between other HSIM-F6 devices.

The User configurable parameters available in the IGMP Local Management Screen should be set to match the parameters on the device acting as the IGMP Querier. The SmartSwitch 2000/6000 will not act as the IGMP Querier in this version of firmware.

Runtime firmware downloads will not work if the TFTP server is using one of the following RFC1918 reserved addresses:10.0.0.0, 172.16.0.0, 192.168.0.0

Distributed Chassis Manager is not supported in SecureFast mode.

SmartTrunking is not supported over ATM in this release of firmware.

When using 802.1Q VLANs with multiple LECs, Spanning Tree Algorithm must be disabled on ATM. This can be done via the Local Management "Network Tools" option using the atm_stp_state command.

MPOA is not supported in this version of firmware.

ATM Lane Services are not supported on the VHSIM-A6DP in this release of firmware.

The SmartSwitch cannot perform Layer 3/4 Classification on frames received by the HSIM-A6DP or VHSIM-A6DP. This means that inbound ATM frames cannot be classified into an 802.1Q VLAN based on the received frames protocol type. Also, inbound ATM frames cannot be assigned an 802.1p Priority value based on a frames protocol type, Layer 3 or Layer 4 information.

Each ATM LEC assigned to an 802.1Q VLAN must be assigned a unique Filtering Database Identifier (FID)

GVRP must be disabled on the following types of interfaces: 802.1Q Trunk Ports, 802.1d Trunk Ports, SmartTrunk Ports and all ATM ports.

802.1Q VLAN tagging is not supported on the VHSIM-A6DP.

All modules installed in a Matrix E7 must have consistent spanning tree settings i.e. all modules configured for spanning tree enabled or all modules configured for spanning tree disabled. This restriction does not apply to individual ports just the entire module. Also the same Spanning Tree algorithm (IEEE or DEC) must be used for all modules installed in a Matrix E7.

Backplane ports must not be administratively disabled on modules installed in a Matrix E7, when there is a device providing a Proxy bridge service in the chassis. If backplane ports are disabled unreliable Spanning Tree operation could occur. For details on Proxy Bridge refer to the "Matrix E7 Proxy Function" in the Installation and Configuration Notes section of this document.

A bootp firmware image download should not be performed to a device that is providing a proxy bridge service. If a download is performed on the proxy device, the proxy will stop bridging, and the other modules in the chassis will not assume the proxy function. An online (runtime) download can be performed on any module including the proxy without affecting proxy function.

When downloading a configuration file to a switch, the file used must have been originally uploaded from the same SmartSwitch type. (i.e. A file from a 6H202-24 can only be downloaded to a 6H202-24.)

Cannot fully populate the DeltaAlias Table in SecureFast after clearing the table.

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

Compliance support:

Compliance Level	Compliant
Year 2000	YES

Known Anomalies: None.

IETF Standards MIB Support:

RFC No.	Title
---------	-------

RFC 1190	Path MTU Discovery
RFC 1213	MIB II
RFC 1354	FIB
RFC 1493	Bridge MIB
RFC 1573	Evolution of MIBII Interfaces
RFC 1595	SONET MIB
RFC 2515	AToM MIB
RFC 1757	RMON MIB

 Cabletron Private Enterprise MIB Support:

Title and Version No.		
Actions-mib version 1.03.01	ctmib2-ext-mib version 1.05.01	Sys-res-mib version 1.00.02
Container mib version 1.02.00	Ctdownload version 1.06.01	Ctpic-mib version 1.02.01
Ctbridge-mib version 1.06.01	Fastethernet-mib version 1.02.01	Trap-mib version 1.01.03
Ctenvir-mib version 1.05.01	Ctbroadcast-mib version 1.00.01	Ctsmt-mib version 1.03.02
ctip-mib version 1.02.01	Cttranslat-mib version 1.01.06	Ctatm-mib version 1.03.00
Community-mib version 1.02.02	Ctrouter-mib version 1.01.00	CtEthernetParameters ver 1.00.00
Ctvlan-ext-mib version 1.03.01	ctups-mib version 1.03.00	CtPriority-mib version 1.00.00
CtIFRemap2-mib ver 2.00.02	CtTxQarb-mib ver 1.00.01	Ct-priority-classify-mib ver 1.00.00
Ct-vlan-classify-mib version 1.00.00	Ctwebview-mib Version 1.00.00	CtRatePolicing-mib version 1.00.00
ctron-cdp-mib version 1.00.03	Ctron-oids-mib version 1.19.14	ctron-alias-mib version 1.00.00
ct-policy-mib version 1.00.00		

ENTERASYS Private Enterprise MIB Support:

Title and Version No.

Enterasys-oids-mib version 1.00.00

Cabletron and Enterasys Private Enterprise MIBs are available in ASN.1 format from the Enterasys web site at:

<http://www.enterasys.com/support/mibs> . Indexed MIB documentation is also available.

SNMP Trap Support:

RFC No.	Title
1157	WARM START LINK UP LINK DOWN
1493	AUTHENTICATION FAILURE NEW ROOT TOPOLOGY CHANGE

 Cabletron Private Enterprise trap Support:

Title		
0x1A2	interfacePortLinkUp	0x40D fddiMACDuplicateMACAddress
0x1A3	interfacePortLinkDown	0x44E aPCLineFailRecovery
0x3EB	ctBroadcastThresholdReached	0x450 aPCLowBatteryRecovery
0x4B0	contLogicalChangesTrap	0x451 aPCAAbnormalCondition
0x4B1	contPhysicalChangesTrap	0x708 wgPsInstalled
0x44F	aPCLowBattery	0x709 wgPsRemoved
0x542	aPCAAbnormConditionRecovery	0x70A wgPsNormal
0x453	aPCShuttingDown	0x70B wgPsFail
0x44D	aPCLineFail	0x70C wgPsRedundant
0x400	fddiPortConnectStateChange	0x70D wgPsNotRedundant
0x403	fddiPortAction	0x70E wgBoardInserted
0x404	fddiPortLerAlarm	0x70F wgBoardRemoved
0x408	fddiMACRMTState	0x2EE0 activePortInATMRedundancyFailed
0x409	fddiMACCurrentPath	0x2EE1 aTMRedundantPortActivated
0x40A	fddi SMTCFstate	0x2EE2 aTMRedundantPortTestFailed
0x40B	fddiRingTopology	0x2EE3 aTMRedundPrimaryPortSkipped
0x40C	fddiMACFrameErrorRatio	0x2F12 atmLecStatus

SecureFast™ Virtual Networking Technology
 Firmware Version 2.00.27
 December 2000

 INTRODUCTION:

This document provides specific information pertaining to firmware version 2.00.27 of Cabletrons SecureFast™ Virtual Networking Technology. For information regarding support of this technology on individual Cabletron products, please refer to the release notes for those products.

It is recommended that one thoroughly review this release note prior to the installation or upgrade of this product.

 Firmware Specification:

Status	Version No.	Type	Release Date
Current Version	2.00.27	Customer	December 2000
Previous Version	2.00.22	Customer	October 2000
Previous Version	2.00.12	Customer	February 2000
Previous Version	1.09.12	Customer	November 1999
Previous Version	1.09.10	Customer	October 1999
Previous Version	1.09.04	Customer	August 1999
Previous Version	1.08.26	Customer	May 1999
Previous Version	1.09.02	Customer	March 1999
Previous Version	1.08.20	Customer	February 1999
Previous Version	1.08.15	Customer	October 1998
Previous Version	1.07.23	Customer	September 1998
Previous Version	1.07.14	Customer	April 1998
Previous Version	1.07.10	Customer	February 1998
Previous Version	1.06.13	Customer	December 1997
Previous Version	1.05.03	Customer	April 1997

 HARDware compatibility:

Please reference individual product release notes for supported hardware revisions or to verify if there is a hardware revision limitation.

 BootPROM compatibility:

Please reference individual product release notes for supported BootPROM versions or to verify the qualifying revisions.

 Network Management Software Support:

NMS Platform	Vers ion No.	Module No.
SPECTRUM VLAN Manager	V2.0	Rev. 0 (or greater)
SPECTRUM 5.0 Rev. 1 with Integrated VLAN Manager	5.0	Rev.1

If you install this image, you may not have control of all of 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 management platform for details.

 SUPPORTED FUNCTIONALITY:

The following lists and highlights the features from previous versions of SecureFast that are supported by this release.

Features	Support
Automatic Membership Registration (AMR) allows SPECTRUM VLAN Manager to recognize non-unicast packets based upon Layer 3 information and automatically assign the source user to protocol-specific VLANs.	Yes
IP Multicast allows for efficient and controllable use of IP Multicast services within SmartSwitch based networks. IGMP versions 1 and 2.	Yes
Active Mesh Topologies supported	Yes
Inter-VLAN communication with or without a traditional router	Yes
Complete user mobility A user that is statically assigned to a VLAN will retain the mapping during the process of a move from one port to another or from one SecureFast switch to another.	Yes
Complete multi-protocol connection-oriented switching services with enhanced resolution services for IP (i.e., IP ARP, AppleTalk AARP)	Yes
Broadcast containment via tag-based flooding of non-resolvable broadcasts. (i.e., IPX SAP, IP RIP)	Yes
Automatic End System/Alias Discovery - SecureFast™ switches maintain a directory of locally attached users. The directory information	Yes
Includes users Network layer address(es), MAC address, switch and switch port location as well	
as protocol stack(s) used by each user.	
Automatic Switch Adjacency Discovery - SecureFast™ switches automatically learn the topology of the switch fabric.	Yes
Automatic Call Re-routing upon link or switch failure as well as user moves	
Best Path Determination based on summarization of link metrics using Virtual Link State	
Protocol (VLSP)	
Connection Load balancing over equal cost paths	Yes

Quick Topology Convergence - Electrical loss detection for Ethernet and Fast Ethernet network links - SMT loss detection for FDDI network links Yes

VLAN Policy (Open/Secure) - Open Inter-VLAN communications via Call Processing and address/VLAN resolution - Secure Inter-VLAN communications via traditional router only Yes

Features Support

Unidirectional or Bi-directional calls may be tapped to any switch within the domain by selecting the network analyzer probe's MAC address or switch port through the SPECTRUM VLAN Manager. Yes

Multiple VLAN membership per user and per port Yes

VLAN membership inheritance per port for users without any previous VLAN association

VLAN locking per port - When a VLAN is locked to a port, the default VLAN for the port becomes the only VLAN mapping that can be associated with the port. This VLAN mapping overrides any and all VLAN mappings associated with the user when the user plugs into this port type.

Enable/Disable Tag-based Flooding per VLAN Yes

Router Port Configuration - Disable Learning of Alias Addresses (MAC and Network Layer Addresses) Per Port Yes

Provisioning for source to broadcast/multicast/unicast connections Yes

Broadcast/Unicast Suppression for Connection Attempts made to Unknown/Invalid Destination addresses Yes

SecureFast-specific CLI Commands via local management or telnet session. Yes

Ability to specify remote IP networks and the optimal router through which the IP networks are reachable. Yes

Ability to specify internal IP networks and a SecureFast Default Gateway per domain. Yes

Layer 3 address mobility: As Layer 3 addresses are reused by different users, the SecureFast switches will automatically purge the old Layer 3-to-MAC address binding from the SecureFast directory and dynamically learn the new Layer 3-to-MAC address binding for the new user. This feature is particularly useful in a DHCP environment where IP address reuse may occur frequently. Yes

SecureFast switches protect themselves against a misconfiguration where two ports on the same switch are looped. Two ports on the same switch that are connected via a jumper cable will automatically transition to "standby" ports, eliminating all traffic from both ports. Yes

Support for three or more (up to 14) SecureFast switches to be attached to a shared media segment. Yes

Support for allowing a network administrator to specify remote IP networks as well as the optimal router through which these IP networks are reachable. This information will be utilized by the SecureFast switches during their resolve process and will ultimately minimize the amount of flood traffic. Additionally, a network administrator will be allowed to specify internal IP networks (IP networks within a SecureFast domain) and a SecureFast Default Gateway. This information will also be utilized by the SecureFast switches during the resolve process. Connection attempts for remote hosts (hosts attached to a remote IP network behind a traditional router) will result in a connection being established to the specified SecureFast Default Gateway. Yes

SecureFast-specific CLI commands, accessible via the Network Tools utility of a SecureFast switch, is supported. The SecureFast-specific CLI commands allow a network administrator to view configuration parameters, view SecureFast directory information, tap a call, enter a resolve entry into the directory and search for a connection either remotely via a telnet session or locally via a console port. Details on SecureFast-specific CLI commands are documented within the Local Management Network Tools guide included on the SPECTRUM VLAN Manager 1.9 Rev0 CD.	Yes
Advanced VLAN Policy - Allows VLAN policy relationships to be controlled on a per VLAN pair basis. Provides granular control over a single VLAN by offering the ability to set inter-VLAN connection and flooding policies per VLAN, as opposed to globally.	Yes
IP Multicast Services - Allows for IP Multicast support through IGMP "snooping" or detection of IGMP v1 and v2 messages. Through IGMP snooping, a source rooted multicast distribution tree is dynamically created, ensuring that multicast transmissions are distributed to only those hosts requesting such services.	Yes
Restrict Port - Allows for the configuration of ports on a switch to be reserved for specified MAC addresses. A user attempting to connect to a "restricted port" with a MAC address not configured for this port will be denied access to the network.	Yes
Restrict User - Allows for the configuration of binding the alias (Layer 3 address, host name, VLAN membership) and MAC address of a user to a specific port. Protects critical resources such as routers and servers by limiting another user or device from utilizing known Layer 3 addresses of routers or servers. Also restricts Layer 2/3 mobility for users within a VLAN to predefined switch ports on a network.	Yes
Dynamic Uplink Switch Support - Allows for scalability beyond current domain size limitations by having VLSP automatically disabled on select periphery switches, yet still be involved in other SmartNetwork Services by assigning a flood path into the VLSP core.	Yes
DHCP Relay Agent - Allows for the definition of a logical address scope on a single DHCP server per switch and the ability to associate VLANs to this scope. Allows VLANs to dynamically represent logical subnets and allows administrators to point new users entering a network to a certain scope. This is accomplished by defining the DHCP attribute per VLAN and making that VLAN the default or inherited VLAN for a certain port.	Yes
AppleTalk Island Support- This feature scopes AppleTalk floods to an Island regardless of geographic location.	Yes
AppleTalk islands are available when AppleTalk AMR is enabled. The default VLAN for islands is the AppleTalk AMR VLAN 0x00000201	
Dynamic Apple Talk Router Address Discovery - Allows for the dynamic learning and discovery of router AppleTalk addresses via inspection of RTMP packets.	Yes
Redundant Access - Allows for redundant access into a VLSP domain for shared network segments. Allows critical devices located on shared network segments to have multiple connections from the shared network segment to a SecureFast fabric without transitioning the shared network segment into a network link.	Yes
CLI Network Tools enhancements - Allows for the ability to enable or disable Protocol Control and Destination/Source Blocker functionality through traditional and easy to use CLI commands.	Yes
Allows the user to clear the Source Block Table (sblock).	
Allows the user to clear the Flood Suppress Table (floodsuppress).	

Allows the user to Set priority of Bridge ST via the CLI.

Improved readability of VLAN name in "vlan list".

Changed portcli to correctly report the physical port number and changed "slot" to report the neighbor INB slot rather than the slot this board is in.

IP Address Learning - Allows for the suppression of Alias learning for IP users configured with a remote network (external IP) address. Also allows networks to be set as invalid to prohibit users with incorrect IP addresses to communicate. Yes

Flood Optimization - Adds an attribute called VLAN ID to VLAN properties to improve overall flood performance. Allows for the assignment of an ID to a VLAN, defining whether a VLAN flood will need to be processed by the switch host or not. VLAN floods are only processed by those hosts requesting such services. Option to flood AppleTalk ARP probe across VLANs. Default mode is to scope to VLAN. Yes

Protocol Control - Allows the network administrator to suppress certain protocols and/or associated frame types from being utilized on a per switch or switch domain basis. The following protocols may be suppressed: Yes

IPX

AppleTalk

NetBEUI

NetBIOS - IP

NetBIOS IPX

BPDU

All applicable frame types

Statistics are kept by recording the number of packets accepted and dropped for each protocol. This includes IPX Raw, IPX on SNAP, IPX on 802.2, IPX on Ethernet II, IP-NetBIOS, IPX-NetBIOS, NetBEUI on 802.2, AppleTalk DDP on SNAP, AppleTalk AARP on SNAP, Spanning Tree, OSPF Broadcast, OSPF Multicast, IGRP, RIP, RIP II, and AppleTalk RTMP. Additionally, a viewable table is created that identifies users that attempt to transmit prohibited protocols, allowing network administrators to isolate and identify users who break these rules. Yes

Automatic Membership Registration (AMR) - Allows for dynamic configuration of VLANs based upon the following protocol criteria: Yes

IP network

IPX by network number and frame type

NetBEUI

NetBIOS-IP

NetBIOS-IPX by network number

DECnet

AppleTalk

Banyan VINES

BPDU

Provisioned Redundant Access - Provisioned RA is a specialized Yes

version of RA. Rather than depending on T1 (front panel) RA hellos to learn of other RA ports on the shared segment, the neighbor information will be provisioned (that is manually entered via MIBs or VLAN Manager).

SNMP Traps - Allows for SNMP traps to be sent by the switch upon certain events: Yes

- o SNMP Trap #1400 - New User -> A new Mac Address has been detected on this switch.

- o SNMP Trap #1401 - Directory Violation -> Currently includes 7 different violation types. Each has its own description, cause, and action. All these may mean that an invalid user is trying to gain access into the network

- o SNMP Trap #1402 - Source Blocked -> A user (MAC address) has surpassed the Source Blocker thresholds and is being blocked from sending traffic on the network.

- o SNMP Trap #1403 - Flood Suppressed -> A destination address has been unresolvable past the thresholds, and is now being flood suppressed (only 1 per x seconds will pass).

- o SNMP Trap #1404 - Port to Standby -> A port has entered standby. There are different standby states: STAND_BY (Port has been manually set to standby or End Station Port (ESP) has detected a neighbor), STAND_BY_FCL (One way neighbor, incompatible neighbor, incompatible hello, etc. One way neighbors are normally during switch resets for a second or two), STAND_BY_LOOPED (Self-originated Hello seen on port), STAND_BY_RA (Redundant Access Port is standby/backup).

- o SNMP Trap #1405 - Port from Standby -> A port, which was in a standby state, is not anymore. Most common scenario is either a one-way neighbor situation is gone, or a standby RA port is now primary.

- o SNMP Trap# 1406 - AgeCnt -> The Connection table on the switch was near capacity (95%) and ran an Age Pass to remove connections. Note - 1 Age Pass equals removing 100 connections.

- o SNMP Trap# 1407 - ChgCnt -> The Flood/Control Channel has changed on the switch. May be due to adding/removing switches or links. Trap not sent for first 5 minutes of Switch uptime.

- o SNMP Trap# 1408 - Found Neighbor -> A new SecureFast Neighbor was found/added.

- o SNMP Trap# 1409 - Lost Neighbor -> A SecureFast Neighbor was lost/removed.

- o SNMP Trap# 1410 - VLAN Mg. Agent IP ->Vlan System Table detected a change in the Vlan Manager IP. The Vlan Manager inserts its IP Address into the Vlan System Table. A trap is sent if the IP changes to/from a non-zero IP Address.

- o SNMP Trap #1411 No Source VLANs -> Takes place when the source user is not mapped to a VLAN

- o SNMP Trap #1412 No Destination VLANs -> Takes place when the destination user is not mapped to a VLAN

- o SNMP Trap #1413 No Source VLANs Enabled -> Takes place when the source VLAN is disabled

- o SNMP Trap #1414 No Destination VLANs Enabled -> Takes place when the destination VLAN is disabled

- o SNMP Trap #1415 No Common Secure VLAN -> Takes place when the source and destination vlans don't share a common secure VLAN. This will result in the packet being flooded - need a Router to route between Secure VLANs.

o SNMP Trap #1416 Inc VLAN User Count -> Takes place when a user is mapped to a VLAN. This is User configurable - the default is to send this trap for every VLAN. Through the VlanTrapAPI, this can be configured to send the trap for ONE specific VLAN.

o SNMP Trap #1417 Dec VLAN User Count -> Takes place when a user is unmapped from a VLAN. This is User configurable - the default is to send this trap for every VLAN. Through the VlanTrapAPI, this can be configured to send the trap for ONE specific VLAN.

o SNMP Trap #1418 VRRP Primary Resign -> Takes place when a VRRP Router sends out a VRRP Hello Packet with Priority of 0 (Resign).

o SNMP Trap #1419 VRRP Primary Aged -> Takes place when the Primary VRRP Route is 'Aged' via the ingress SecureFast Switch.

o SNMP Trap #1420 VRRP Secondary Up -> Takes place when Backup VRRP Router becomes Primary.

o SNMP Trap #1421 HSRP Primary Resign -> Takes place when a HSRP Router sends out a HSRP Hello Packet with Priority of 0 (Resign).

SNMP Trap #1422 HSRP Secondary Up -> Takes place when Backup HSRP Router becomes Primary.

DHCP Enhancements- The option is now available to send DHCP replies directly to the intended DHCP client. The entire SecureFast fabric must have this feature enabled in order to function properly. This applies to BOOTP packets as well as DHCP packets. Yes

DHCP Islands- DHCP Islands are now supported. This new feature supports the ability to administer geographically based DHCP islands. This applies to BOOTP packets as well as DHCP packets. The entire SecureFast fabric must have this feature enabled in order to function properly. The DHCP Islands feature is similar to the DHCP Server feature and it provides geographical granularity. Yes

DHCP Server VLAN - Limits the flood scope of DHCP request packets to only DHCP Servers Yes

statically mapped to that VLAN.

HSRP Support - With IP Multicast enabled, establishes a multicast connection between the routers so that the inter-router messages no longer rely on the flood channel. Yes

Reap connections if MAC address is in Source Blocker. Yes

Vflood Support The default behavior is for SecureFast to learn the AppleTalk alias address on an AppleTalk ARP probe. This discovery step can increase the time it takes for the switch to flood the probe. Vflood support, when enabled throughout the switch fabric, reduces the time to flood by not discovering the source on the probe request. The probe is then flooded across all VLANs. In addition, when the switch receives the probe reply from the destination, SecureFast will not perform an inter-switch resolve or connection setup but rather flood the reply to all ports. Yes

VRRP Call Processor Support- Creates a connection between the routers running VRRP so that the VRRP packets are not dependent on the flood channel. Yes

VRRP Support - With IP Multicast enabled, cleans up the Directory after a router has transitioned between primary and standby. Yes

Improved IP Multicast setup time- Creates connections along Control Channel until the traffic exceeds a certain threshold (query-high limit, query-low limit). At that point, connections are made based on VLSP. This feature is disabled by default. The Entire SecureFast fabric needs to enable this feature in order for it to function properly. Yes

IP Multicast Policy Checking- There are now two modes of policy checking for IP Multicast: Yes

1. no check (default) 2. check only the default VLAN of the port

SecureFast Violation Table Addition- If a protocol is disabled from running in the SecureFast domain and user A attempts to talk using that protocol, the switch will intercept and drop the packet. It will then add an entry for user A and what protocol was used (but no Layer 3 information) to the Violation Table. The switch maintains counters for each packet that is dropped per protocol. Yes

Persistent IP- A delay timer has been added to save the modules IP address when NVRAM is cleared. Yes

Tier 2 Support - Allows extending the Uplink Switching model one more tier. In order to utilize this functionality, the switch needs to be provisioned as Tier 2. Yes

Installation and Configuration Notes:

In general, SecureFast Virtual Networking Technology is a set of features that are available within a version of firmware. If you would like to upgrade an existing product to support the above features which are available to the product, please follow the TFTP download instructions that are included with your firmware image upgrade kit. TFTP download instructions are also available on the Cabletron Systems Support Web Site at:

<http://www.cabletron.com/support/techtips/tk0020-9.html>

Firmware Changes and Enhancements:

The following is a list of major firmware changes and enhancements for the 2.00.27 release. These are in addition to the existing features listed in the

Supported Functionality section.:

A check has been added to insure that the VST/VLSP component can not be enabled when the uplink facility is enabled.

A problem was introduced in SecureFast 2.0.12; if you administratively provisioned a multicast connection through the McCnxTblApi, it could not be deleted unless you cleared NVRAM.

Made a fix when adding an IP Multicast Filter connection.

Update so connections with the Source MAC Address equal to the Destination MAC Address get a CRITICAL message instead of an ERROR message in the SecureFast Event Log.

Don't automatically Age the OSPF Directory/Connections based on the "OSPF DeadTimer". OSPF will now add BOTH 224.0.0.5 and 224.0.0.6 IP Multicast Groups when it sees a packet from either 224.0.0.5 or 224.0.0.6.

Refresh the IpMcRib Entries Age every time we get a router packet. SecureFast was aging out OSPF/HSRP/VRRP (non-provisioned) entries (default age time is 50 minutes).

Fixed SFPSTrap - potential for not saving persistent information (TRAP Disables).

Broken in SecureFast 2.00.22. SecureFast would only subnet validate an IP Multicast packet. A fix has been made to also subnet validate an IP packet.

VRRP Timer fix to prevent premature aging. Set the default VRRP Ageout Timer

to (3*SendFrequency)+1 rather than just (3*SendFrequency). This is the default setting. The VRRP Ageout Timer can be manually configured.

When SecureFast receives an AppleTalk ARP with the destination address in the startup range 65280.0 through 65534.255, it will no longer try to resolve (SecureFast does not learn these AppleTalk addresses). SecureFast will go directly to the flood step.

EIGRP fix to prevent the EIGRP IP multicast group (224.0.0.10) from being learned as "local" on ports when IP multicast is enabled.

Known Restrictions and Limitations:

Network Configuration

SecureFast switches attached to shared network links are supported in limited configurations of up to (14) SecureFast switches.

To optimize performance and to utilize full duplex capabilities, all SecureFast™ VLAN switches should be interconnected via point-to-point inter-switch links.

It is recommended not to exceed 7 switch hops from the root switch.

Users may not be attached to inter-switch links without using the Redundant Access feature.

A total of 1024 VLANs may be supported by a SecureFast™ VLAN domain. A SecureFast VLAN domain terminates at a traditional router boundary.

A single user (MAC address) may be a member of up to 8 static VLANs, not including AMR VLANs.

This release supports the following module configurations per SmartSwitch 9000 chassis:

All SmartSwitch 9000 INB-based modules configured for SecureFast mode.

One SmartSwitch 9000 INB-based module configured for SecureFast mode and multiple SmartSwitch 9000 FNB-based modules configured for Traditional Bridging mode.

All SmartSwitch 9000 INB-based modules configured for SecureFast mode and multiple FNB-based modules configured for Traditional Bridging mode.

The following SmartSwitch 9000 chassis configurations are not supported in this release: Two or more SmartSwitch 9000 INB-based modules configured for SecureFast™ mode and one or more SmartSwitch 9000 INB-based modules configured for Traditional Bridging mode. One SmartSwitch 9000 INB-based module configured for SecureFast mode and multiple SmartSwitch 9000 INB-based modules configured for Traditional Bridging mode. The work around for this configuration is to disable the INB on each traditional mode module and connect one front panel port of the traditional mode module to one of the front panel ports of the SecureFast mode module.

A device with multiple interfaces that share a common MAC address is supported in this release when the Duplicate MAC address support is enabled.

It is recommended that shared segments of less than 500 users be attached to a single switch port. This limitation is dependent on the specific traffic patterns and protocols used by end stations on the shared segment. Optimal performance is achieved by attaching a single user to a single switch port.

Layer 3 Mobility

Any multi-interface device routing IPX (type 20) packets may inadvertently be associated with incorrect NetBIOS names. Solution: Toggle ports to which the multi-interface device is attached, to router ports. Layer 3 learning is disabled by default for Router ports.

Automatic Membership Registration (AMR)

To reach quiet users, if the user is performing an ARP request and it is unresolvable, the switch will flood out one packet for every 20 unresolved hits. These go to all port(s) that do not have any IP subnet AMR VLANs mapped to it. Note that this behavior is for ARP requests only. If this is occurring and the user is being Flood Suppressed, the normal leaking of once every minute will be increased to once every two minutes.

IP-Subnet & NetBIOS Automatic Membership Registration (AMR): When IP Subnet and NetBIOS AMR configuration parameters are enabled, IP NetBIOS flood packets are delivered to all ports where IP-NetBIOS users have been discovered regardless of subnet.

If both the IP-subnet and the NetBIOS rules are used simultaneously, flooded IP-NetBIOS packets may not be flooded to the desired flood group. Since the IP-NetBIOS packet is processed by the call processing stack, it reaches the IP call processor before reaching the IP-NetBIOS call processor. Because the IP-NetBIOS call processor is last, it gets the final say as to which VLAN the packet will be flooded. The "IP-NetBIOS" VLAN is simply a VLAN and not a template for VLANs based on IP-Subnet, and therefore, when an IP-NetBIOS packet is flooded, it may violate the IP-Subnet rule. The following is a sample scenario: Both the IP-Subnet and NetBIOS rules are enabled. The IP-Subnet AMR mask is 255.255.255.0 Node A (134.141.xx.3) transmits an IP-NetBIOS packet and is joined to both the 134.141.xx.0 and IP-NetBIOS VLANs. Node B (134.141.xy.1) transmits an IP-NetBIOS packet that needs to be flooded. The IP-NetBIOS call processor determines that the packet should be flooded to the IP-NetBIOS VLAN, which means that the packet would also flood to the port where node A resides (a different subnet).

IPX RIP/SAP Automatic Membership Registration (AMR) - IPX-NetBIOS_00000000 VLAN When Windows for Workgroups 3.11 (WFW311) clients use IPX-NetBIOS for sharing resources, they do not follow the standard IPX procedure for determining their IPX network number. Instead of issuing a SAP GNS (Get Next Server) request, the client simply uses a network number of 0x00000000 until it hears an IPX RIP advertisement. It will then look at the source network number used in the RIP packet and start using it as its own network number. If the client never receives an IPX RIP advertisement, it will continue to use the zero network number forever. If the IPX RIP/SAP AMR rule is being used along with the NetBIOS AMR rule, the IPX RIPs from servers/routers will be auto-segmented and therefore not sent to the WFW311 client. The WFW311 client will continue to use the zero IPX network number for all it's IPX-NetBIOS packet transmissions. This will essentially cut off the WFW311 client from all other IPX-NetBIOS clients that are correctly using non-zero network numbers.

To remedy this, IPX RIP packets are flooded to the "IPX-NetBIOS_00000000" VLAN as well as the "IPX RIP/SAP__" VLAN, allowing the WFW311 client to receive a RIP response advertisement. When the WFW311 client receives the RIP advertisement, it changes it's IPX network number to a valid, non-zero number and starts sending packets using the new number. When the SecureFast switch sees that the WFW311 client is no longer using a zero network number, it removes the node from the "IPX-NetBIOS_00000000" VLAN and adds it to the "IPX-NetBIOS_,netNumber>" VLAN. Furthermore, the IPX RIP advertisements are no longer sent to the WFW311 client because it is no longer a member of the "IPX-NetBIOS_00000000" VLAN.

IP Multicast

The Time To Live (TTL) parameter within the Multicast feature can no longer be changed. This reduces unnecessary configuration and allows the switch fabric to rely on any attached multicast routers to handle TTL.

Potential switch reset problem if you provision an IP Multicast connection through the IP Multicast Connection API before IP Multicast is ever enabled on the switch.

Call Tap

A call cannot be tapped to more than one out port simultaneously.

A call cannot be tapped to a port that is already enabled as an out port for that call.

A tapped call will be lost if either the source or destination user of that call is moved.

Solution: To re-establish the tap, the tap for the original connection must be released and re-tapped for the new connection.

When a call is tapped and the connection is released or ages out, the tap will no longer function and may appear as a tapped call in the Connection Table.

Solution: To re-establish the tap, the tap for the original connection must be released and the new connection must be tapped.

Restrict User

If new user thrashing is occurring, or a user is continually being learned on different switches, the user cannot be restricted to a switch. An example of this is when two or more users are configured with the same IP address.

Redundant Access

Redundant Access supports two switches on a Redundant Access link. (A Redundant Access switch port can only have one Redundant Access neighbor) The Primary and Standby Redundant Access ports of a single Redundant Access connection cannot be located on the same switch. The maximum number of Redundant Access links between any two switches is 20. For example:

A 9E423-36 (switch 1) can have up to 20 ports configured for Redundant Access with switch 2. The remaining 16 ports could then be configured for Redundant Access with switch 3.

[Image]

Note: If a switch detects a 21st Redundant Access adjacency to any given switch, it will refuse the configuration and keep that port in RA_HALTED mode (which is essentially STANDBY). If a switch resets with over 20 Redundant Access ports configured, you may get less than the 20 Redundant Access links active once the switch returns.

FDDI configurations are also supported using Redundant Access. On the 6000 and 2000 platforms, Redundant Access can be configured on an actual FDDI Ring. On the 9000 platform, support is for a single FDDI user only. The way FDDI Redundant Access works is that the FDDI can be dual homed into the SecureFast network, but will be single attached due to the Primary and Standby Redundant Access ports.

The Redundant Access standby port will not detect a pulled link and a reinsertion of another link until it becomes primary.

Advanced VLAN Policy

The maximum number of VLANs supported by the Advanced VLAN Policy application is 128. This allows for each of the 128 VLANs to be paired against any or all of the remaining 127 VLANs for flooding and unicast connection policies to be defined. If there are more than 128 VLANs defined within the SPECTRUM VLAN Manager, Advanced VLAN Policy cannot be enabled.

Advanced VLAN Policy is a tool for static VLANs only and will not work with inherited AMR VLANs. SPECTRUM VLAN Manager will allow for Advanced VLAN Policy to be enabled when AMR VLANs are enabled, however, AMR flooding policy will defeat the purpose of the Advanced VLAN Policy. It is recommended to use either Automatic Membership Registration or Advanced VLAN Policy one at a time.

When Advanced VLAN Policy is enabled, the default VLAN policy (Open or Secure) is used to create the initial unicast connection relationships between all VLANs. Since there was no notion of inter-VLAN flooding in the legacy model, all VLANs are configured to not flood to each other by default in order to act consistently with the legacy model when no changes have been made. The default flood mode can be changed, but only through SPECTRUM VLAN Manager. If the default flood mode is enabled for a VLAN, other VLANs will be able to flood to that VLAN by default in Advanced VLAN Policy. Note, however the corollary is not true. Operationally, if the Advanced VLAN Policy table has not been changed after enabling it, flooding and unicast connection policies will behave as if they were in their legacy (default) mode. (i.e. Users in Open VLANs will be able to connect to users of other Open VLANs, while users in Secure VLANs will not be able to connect to any VLANs. All

flooding will remain intra-VLAN.) Outside of Advanced VLAN Policy, VLANs can still be toggled to be Open or Secure. If the Advanced VLAN Policy table remains unchanged, toggling the default mode of a VLAN to be Open or Secure will result in a change of the flooding and unicast connection policies within Advanced VLAN Policy to operate exactly as if the VLANs were in their new legacy (default) mode. If changes were made to any of the flooding or unicast connection policies within Advanced VLAN Policy between any two VLANs, changing the default policy mode will not affect the Advanced VLAN Policy table. The settings within Advanced VLAN Policy will remain as they were changed, however new VLANs to be added to the Advanced VLAN Policy table will default to operate with the existing VLANs according to the existing VLANs legacy (default) mode.

Note: The difficult part about the above scenario is what happens when a VLAN is toggled when Advanced Policy is running. Note that the logic checks to see if any changes were made respecting that particular VLAN. If no changes have been made in Advanced VLAN Policy regarding that particular VLAN, then the connect policy against all other VLANs will be re-calculated according to its new default value. If it has been modified in any way through Advanced VLAN Policy, then the previously defined settings will be left alone. The new default connect policy will be used to determine its initial value against any newly created VLANs. This check is done on an individual VLAN basis and not table-wide. That is, changes may have been made to the table for other VLANs, but if the behavior of a particular VLAN was not changed at all, then its settings will be modified in the Advanced VLAN Policy table.

If Advanced VLAN Policy is disabled, then VLAN behavior will resume according to the legacy rules.

Disabling flooding (flooding-off) for a VLAN prevents floods from being transmitted on that VLAN, but does not necessarily prevent floods from being received on that VLAN. A VLAN may still receive floods if it has been defined by Advanced VLAN Policy for this VLAN to see the broadcast floods from other VLANs where flooding is still enabled. For example: If flooding is disabled for the blue VLAN, floods will no longer be generated by users within the VLAN. However, if rules within Advanced VLAN Policy were changed to allow other VLANs to flood to the blue VLAN, the blue VLAN will still see these broadcasts. Disabling a VLAN through SPECTRUM VLAN Manager still disables a VLAN with Advanced VLAN Policy.

If switches reset, Advanced VLAN Policy will remain enabled, however VLANs will operate in default mode until the SPECTRUM VLAN Manager populates the Advanced VLAN Policy table. If NVRAM is toggled and reset on a switch, the VLANs will operate in default mode.

It is recommended to have users maintain single VLAN membership as opposed to multiple VLAN memberships when using Advanced VLAN Policy. With Advanced VLAN Policy, only a single membership is actually necessary because individual connection policies can be made on a VLAN pair basis. It is mechanically possible to maintain multiple VLAN memberships within Advanced VLAN Policy, however, there is a possibility for strange and unpredictable flood behavior.

DHCP

When configuring DHCP relay agents, it is a requirement that each VLAN have a unique relay agent value. The

relay agent values can represent the same subnet, but still need to be unique. For example:

Given subnet 134.141.xx.0

	VLAN Name	Relay Agent
Configuration DONT	Blue Red	134.141.xx.1 134.141.xx.1
Configuration DO	Blue Red	134.141.xx.1 134.141.xx.2

Note: The relay agent IP addresses represent the same subnet but they are unique.

DHCP requests will be flooded to the source VLAN from which the request originated. Using Advanced VLAN Policy, it is possible to create a DHCP

server target VLAN. This will allow the user to open flooding and connection policies between the DHCP Server VLAN, and all other VLANs which the user wishes the DHCP Server to service. This avoids the requirement of needing to place the DHCP Server as a member of all VLANs which it needs to serve.

Duplicate MAC Address Support

On non-Token Ring platforms, if the link goes down on the port with the active duplicate MAC user, the switch will not attempt to reconnect to the other duplicate MAC user(s). Workaround is to delete the connections to the inactive duplicate MAC user.

VRRP

VRRP packets will be flooded between routers via the Control channel. A problem can occur if the Control channel is broken. During this time, both routers will think they are master. Once the Control channel recovers, the routers will arbitrate so that only one is master but the virtual MAC will be in both places.

IP multicast has to be enabled in order for the switches to clean up the Directory after the primary router has transitioned to backup.

IP multicast has to be enabled in order for the switches to create a multicast connection between the routers and therefore, not rely on the flood channel for inter-router communication.

HSRP

IP multicast has to be enabled in order for the switches to create a multicast connection between the routers and therefore, not rely on the flood channel for inter-router communication.

If the router interfaces of a single group are located on different switches but same chassis, it is recommended to decrease the switch hello timers. This only needs to be done if the HSRP timeout value is less than 20 sec. This is needed since there is no electrical loss on the INB, therefore it will take 20 seconds for that neighbor to be reaped from the flood path of its INB neighbors.

The inability to ping the virtual IP is possible in certain instances. If a link/switch is lost that was in the multicast connection path and the new fail over link does not have the virtual IP in its remote cache this problem can occur. A resolve timeout will occur because the switch will try to resolve this IP, hence we have to wait for the flood path to be reestablished.

The configuration of a SecureFast Default Gateway in a network where Cisco HSRP (Hot Standby Routing Protocol) is deployed may cause connectivity problems when fail over occurs. Workaround: 1) Do not configure a SecureFast Default Gateway in networks where HSRP enabled routers are deployed. 2) It is recommended to place the connections for HSRP enabled routers (of the same HSRP group) within the same chassis.

Dynamic and Provisioned Uplink

With the addition of Dynamic Uplink support in the 1.8 release, there are now three different types of switches: VLSP Fabric switches, Dynamic Uplink switches, and Provisioned Uplink switches. The default mode for a switch is VLSP fabric switch. It is recommended to migrate to Dynamic Uplink switching, if the previous mode of Uplink switching (Provisioned Uplink Switching) is still being used. The older Provisioned Uplink feature will be deprecated in favor of Dynamic Uplink switching with all future releases. The way these switches operate with each other varies under different circumstances. The following rule applies to configuring switches with the three basic types of topology awareness: dynamic uplink, provisioned uplink, and fabric (VLSP); If Provisioned Uplink switching is the current mode of operation for a switch, the switch needs to be unprovisioned before being transitioned into a Dynamic Uplink or VLSP Fabric switch. However, a VLSP Fabric switch will automatically transition to a Dynamic Uplink switch upon configuration.

When reconfiguring uplink switches within a chassis to be VLSP Fabric switches, or when there is a chassis that has been configured for uplink and it is desired that the chassis be part of the VLSP fabric again, it is necessary to make sure all switches within the chassis are reconfigured before one switch can finish the reboot process. This is due to the various repercussions that occur when switches of mixed topology awareness are in the

same chassis.

It is recommended to rerun the "Uplinker" tool whenever switches are added or changes are made to an Uplink chassis. If it is unknown which mode the switch being placed into the chassis is currently in, it is possible to verify switch mode by checking the bindery status within Local Management. If the service.facility.uplink component is enabled, the switch is a Dynamic Uplink Switch. If the service.facility.fabric component is enabled, the switch is a VLSP Fabric switch. If both the service.facility.uplink and service.facility.fabric components are disabled, the switch is a Provisioned Uplink switch.

Additionally, connecting the chassis together via front port or Inter-Switch Link connections has different results when using the different types of switches. The following graph represents the various outcomes when connecting switches of various topology awareness to each other:

[Image]

Please note that clearing NVRAM and resetting the switches will enable the switch as a VLSP Fabric switch. This is the default mode for all switches, and will ensure that all previous uplink configurations are erased from memory.

Router Port

Layer 3 learning can be toggled on router ports only, instead of all ports. Previously, the layer three learning state on a port could be changed independently from the router port state. With this "fix", the only time there can be a layer three learning value other than "other", is when the port is a router port. When a port is a router port, there are two values for the layer three learning: enabled and disabled. The default value is disabled. This means that when the user migrates from 1.7 to 1.8, the non-router ports will show "other" for layer 3 learning. (Note: Learning will take place.) For the router ports, layer three learning will be disabled. This means that the switch will behave the same before and after a migration. (In regards to layer three learning.)

Connections

Provisioned source user-to-broadcast/multicast/unicast circuits configured through SPECTRUM VLAN Manager are not restored upon a switch reset. Solution: Reconfigure your provisioned circuits following a switch reset.

Persistence

There are a number of attributes that remain persistent in memory on SPECTRUM VLAN Manager or in the switch NVRAM, upon a switch reset. The following table below lists the various attributes which may or may not be persistent in 1.07.x, 1.08.x, and 1.09.x networks. The column title indicates the feature, switch firmware version, VLANServer application version and comments. For example, the column for SFS 1.7.x identifies all SmartSwitches operating with SFS version 1.7 firmware. The cells below the column will explicitly state the persistence of objects in memory or the maximum number of entries. Furthermore, VLAN membership (user) on item 6 will not be retained in the switch on a reset, however the application does store it in the database and will restore the original configuration after successfully contacting the switch.

Cabletron VLAN Switch and Application Matrix

Feature	SFS 1.7.x	SFS 1.8.x	VLANServer 1.8, 1.9 SFS 1.9.x and 2.0 SFS 2.0.x *	Comments
1 Default VLAN on port	yes	yes	yes	1
2 Policy for VLAN	yes	yes	yes	1
3 VLAN flooding enable/disable	yes	yes	yes	1
4 VLAN enable/disable	yes	yes	yes	1
5 VLAN name and attributes	no	no	yes	1
6 VLAN membership (user)	no	no	yes	1
7 Provisioned connection Protocol Policy Rules	no	no	no	
8 flood on/off	no	yes	yes	2

9 discover layer 3 address on/off	no	yes	no	
10 SAP (Service Access Point)	up to 64	up to 64	no	4
11 AMR rules	up to 32	up to 32	yes	5
12 AMR subnet prefix and mask	up to 32	up to 32	yes	5
Multicast				
13 IGMP query interval	yes	yes	yes	
14 Provisioned IP multicast receiver entries	up to 512	up to 512	no	5
15 Provisioned IP multicast connection	up to 256	up to 256	no	5
16 IPMC Sender policy	N/A	yes	yes	
17 IP multicast router interface	yes	yes	no	
18 IP Multicast Policy Modes	no	yes* (2.0 only)	no	
19 IP Multicast Traffic Control/Suppression Router Configuration	no	yes* (2.0 only)	no	
Router Configuration				
20 Router port	yes	yes	yes	
21 Subnet resolve entries	up to 128	up to 256	yes	5
22 Default gateway configuration	yes	yes	no	
23 Subnet resolve subnet mask	yes	yes	yes	
24 Subnet validation mode	N/A	yes	yes	6
Topology port parameters				
25 Lock port as access/network	yes	yes	no	7
26 static path uplink and downlink config.	yes	yes	no	8
27 VLAN hello TX and RX frequency	no	yes	no	9
Services				
28 IP	yes	yes	no	10
29 AppleTalk	yes	yes	no	10
30 Novell	yes	yes	no	10
31 NetBIOS over IPX	yes	yes	no	10
32 NetBIOS over IP	yes	yes	no	10
33 NetBEUI	yes	yes	no	10
34 User mobility	yes	yes	no	10
35 Source blocker	yes	yes	no	10
36 Flood Suppress	yes	yes	no	10, 13
37 Multicast	yes	yes	no	10
38 Redundant Access	yes	yes	no	10
39 DHCP relay address	no	yes	no	10, 11
Other				
40 Advanced VLAN Policy	N/A	no	yes	
41 Call tap	no	no	no	
42 Port redirect	N/A	no	no	
43 Connection age parameters	no	no	no	12
44 Flood suppress parameter	no	yes	no	13
45 Source block parameter	no	yes	no	
46 Static directory entries	no	no	no	14
47 Restricted port (see comments)	N/A	yes	no	15
48 Restricted node	N/A	yes	no	15
49 MAC / port unblockable	no	yes	no	15
50 Redundant access	yes	yes	no	

Comments:

1. If switch NVRAM is cleared, application will restore setting.
2. Flooding of protocol such as IP, IPX, AppleTalk, RTMP, and OSPF.
3. Switch proxy ARP reply enable / disable.
4. Frame type control is also available such as IPX 802.3 and IPX 802.2
5. Memory constraints on the switch cap the number of entries.
6. Options are learn IP address, discard packet, or off (do not learn IP).
7. ATM SVC logical ports will not be persistent because after reset, the logical port cannot be guaranteed.
8. Static uplink not FCS feature.

9. Available in the VLAN topology agent port table.
10. Bindery view exposed in UI.
11. Attributes of VLAN are persistent only when VLAN is mapped to a port.
12. Threshold for aging connections (represented in % of connections).
13. Earlier versions referenced it as "destination blocker".
14. The directory API provides a mechanism to add MAC, IP, etc.
15. Restricted port, restricted node, and MAC port unblockable entries are persistent in the switch and share a common pool of 128 entries (Restrict on alias uses two entries, alias and MAC)

Topology

VLSP, by default, will apply a 100 Mbps metric to all Fast Ethernet ports regardless of speed or duplex settings.

In topologies where three or more SecureFast switches are attached via a shared media segment, redundant physical connections or paths in parallel with the shared media segment may not exist. Please reference drawing below:

Shared Neighbor Support on Network Links

[Image]

Sample Topologies

[Image]

[Image]

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

Compliance support:

Compliance Level	Compliant
Year 2000	Yes*

Known Anomalies: None.

*NOTE: Please reference individual product release notes for complete information on Year 2000 Compliance.

IETF Standards MIB Support:

Please reference individual product release notes for a complete and updated list of RFCs and IETF standard MIBs supported by this release.

Cabletron Private Enterprise MIB Support:

Title	Version
ct-sfps-base	Rev. 0.0.18
ct-sfps-bindery	Rev. 0.0.4
ct-sfps-call	Rev. 0.0.16
ct-sfps-common	Rev. 0.0.13
ct-sfps-conn	Rev. 0.0.2
ct-sfps-connection	Rev. 0.0.11
ct-sfps-diagstats	Rev. 0.0.3
ct-sfps-directory	Rev. 0.0.11
ct-sfps-esys	Rev. 0.0.17
ct-sfps-eventlog	Rev. 0.0.4
ct-sfps-flood	Rev. 0.0.3
ct-sfps-inc	Rev. 0.0.14

ct-sfps-mcast	Rev. 0.0.5
ct-sfps-path	Rev. 0.0.8
ct-sfps-pktmgr	Rev. 0.0.15
ct-sfps-policy	Rev. 0.0.3
ct-sfps-port	Rev. 0.0.11
ct-sfps-resolve	Rev. 0.0.8
ct-sfps-sflsp	Rev. 0.0.6
ct-sfps-size	Rev. 0.0.3
ct-sfps-softlink	Rev. 0.0.3
ct-sfps-tap	Rev. 0.0.4
Ct-sfps-topology	Rev. 0.0.9
ct-sfps-vlan	Rev. 0.0.9
ct-sfps-vstp	Rev. 0.0.5

Cabletron Private Enterprise MIBs are available in ASN.1 format from the Cabletron Systems Support Web Site at: <http://www.cabletron.com/support/mibs>
Indexed MIB documentation is also available.

SNMP Trap Support:

Please reference individual product release notes for a complete and updated list of Traps/RFCs supported.

Cabletron Private Enterprise trap Support:

Please reference individual product release notes for a complete and updated list of Traps supported.

GLOBAL SUPPORT:

By Phone: (603) 332-9400

By Email: support@cabletron.com

By Web: <http://www.cabletron.com/support>

By Fax: (603) 337-3075

By Mail: Cabletron Systems P.O. Box 5005 Rochester, NH 03867-5005

For information regarding the latest firmware available, recent release note revisions, or if you require additional assistance, please visit the Cabletron Systems Support Web Site.

Contact Us | Locations | Terms and Conditions | Feedback | Privacy ©2003, Enterasys Networks, Inc. All rights reserved.