

Managed Switches

Ethernet Managed Switching Hubs

CTRLink®

Software Manual for Web Browser

Version 5.0.x

Covering Product Series : EI-WPMR-M

TD070446-0MA



CONTEMPORARY CONTROLS®

Trademarks

Contemporary Controls and CTRLink are registered trademarks of Contemporary Control Systems, Inc. Other product names may be trademarks or registered trademarks of their respective companies.

Copyright

© Copyright 2006, by Contemporary Control Systems, Inc. All rights reserved. No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual, or otherwise, without the prior written permission of:

Contemporary Control Systems, Inc. 2431 Curtiss Street Downers Grove, Illinois 60515 USA	Tel: +1-630-963-7070 Fax: +1-630-963-0109 E-mail: info@ccontrols.com WWW: http://www.ccontrols.com
Contemporary Controls (Suzhou) Co., Ltd 11 Huoju Road Suzhou New District China 215009	Tel: +86-512-68095866 Fax: +86-512-68093765 E-mail: info@ccontrols.com.cn WWW: http://www.ccontrols.aisa
Contemporary Controls Ltd Sovereign Court Two, UWSP Sir William Lyons Road Coventry CV4 7EZ UK	Tel: +44 (0)24 7641 3786 Fax: +44 (0)24 7641 3923 E-mail: info@ccontrols.co.uk WWW: http://www.ccontrols.co.uk
Contemporary Controls GmbH Fuggerstraße 1 B 04158 Leipzig, Germany	Tel: +49 341 520359 0 Fax: +49 341 520359 16 E-mail: info@ccontrols.de WWW: http://www.ccontrols.de

Disclaimer

Contemporary Control Systems, Inc. reserves the right to make changes in the specifications of the product described within this manual at any time without notice and without obligation of Contemporary Control Systems, Inc. to notify any person of such revision or change.

**WARNING — This is a Class A product as defined in EN55022.
In a domestic environment this product may cause radio interference
in which case the user may be required to take adequate measures.**

1 Table of Contents

1	TABLE OF CONTENTS	3
2	HISTORY	5
3	INTRODUCTION.....	5
3.1	Software	7
3.2	Product Images.....	7
4	ADVANCED OPERATION	8
4.1	General Considerations	8
4.1.1	LEDs	8
4.1.2	Accessing the Web Server	8
4.1.2.1	Web Browser	8
4.1.2.2	Initial Access.....	8
4.1.3	On-Screen Help.....	11
4.1.4	Username and Password	12
4.1.5	Restoring Factory Default Settings	12
4.2	Home Page (Main Menu).....	13
4.2.1	Upload/Download Settings	15
4.2.2	Uptime.....	15
4.2.3	Switch Temperature	15
4.2.4	Port Configuration and Port Statistics.....	15
4.2.4.1	Port Configuration (Figure 11 upper panel).....	15
4.2.4.2	Port Frame Statistics (Figure 11 lower panel)	17
4.2.5	System Configuration	18
4.2.5.1	Configure IP Address.....	19
4.2.5.2	Configure Trunking (Copper Ports Only).....	20
4.2.5.3	Configure Port Mirroring	21
4.2.5.4	Virtual Local Area Networks (VLANs)	22
4.2.5.5	Configure Filtering and Forwarding Table	29
4.2.5.6	Configure Quality of Service (QoS)	32
4.2.5.7	Configure Fault Relay	38
4.2.5.8	Configure Redundancy	40
4.2.5.9	Configure Rate Control.....	48
4.2.5.10	Configure Port Security	50
4.2.5.11	Configure IGMP Snooping	51

4.2.6	SNMP Configuration	52
4.2.6.1	Configure System Information (Figure 39, upper panel).....	53
4.2.6.2	Configure SNMP Community (Figure 39, middle panel).....	53
4.2.6.3	Configure SNMP Trap Receivers (Figure 39, lower panel)	53
4.2.7	Performance Monitoring	54
4.2.7.1	Browse Address Table	54
4.2.7.2	Traps Log	55
4.2.7.3	Monitor STP Port Status	55
5	APPENDIX.....	57
5.1	SNMP.....	57
5.1.1	Managed Objects for TCP/IP Based Internet (MIB-II) — From RFC 1213	57
5.1.1.1	‘System’ group 1.3.6.1.2.1.1.....	57
5.1.1.2	‘Interfaces’ group 1.3.6.1.2.1.2.....	58
5.1.1.3	‘IP’ group 1.3.6.1.2.1.4.....	60
5.1.1.4	‘ICMP’ group 1.3.6.1.2.1.5.....	61
5.1.1.5	‘TCP’ group 1.3.6.1.2.1.6.....	61
5.1.1.6	‘UDP’ group 1.3.6.1.2.1.7.....	62
5.1.1.7	‘Transmission’ group 1.3.6.1.2.1.10.....	62
5.1.1.8	‘SNMP’ group 1.3.6.1.2.1.11.....	62
5.1.2	Managed Objects for Bridges — From RFC 1493	65
5.1.2.1	‘dot1dBase’ group 1.3.6.1.2.1.17.1.....	65
5.1.2.2	‘dot1dTp’ group 1.3.6.1.2.1.17.4.....	69
5.1.2.3	‘dot1dTpFdbTable’ 1.3.6.1.2.1.17.4.3.....	69
5.1.2.4	‘dot1dTpPortTable’ 1.3.6.1.2.1.17.4.4.....	69
5.1.3	Managed Objects for Ethernet-like Interface Types — From RFC 1643.....	70
5.1.3.1	Ethernet-like Statistics Group — ‘dot3StatsTable’ 1.3.6.1.2.1.10.7.2	70
5.1.4	Evolution of the Interface Group of MIB-II — From RFC 1573	72
5.1.5	Private Managed Objects	73
5.1.5.1	Relay Group – 1.3.6.1.4.1.17384.1.1.2.....	73
5.1.5.2	RapidRing Group – 1.3.6.1.4.1.17384.1.1.3.....	73
5.1.6	Message Format for SNMP Operations.....	74
5.1.6.1	Format of Command Messages	74
5.1.6.2	Traps for SNMPv1.....	75

2 History

9/27/2004	Initial Release
3/01/2005	Added IGMP Snooping, Rate Control & Port Security
9/01/2005	Added STP/RSTP
8/28/2009	Added Multi-RingID(Max support one hundred RindID)

3 Introduction

Managed switches in the CTRLink[®] family provide capabilities beyond those in both Plug and Play (PnP) and Configurable Switches. Besides conventional PnP features (auto-negotiation, 10/100 Mbps data rate, half- or full-duplex operation, flow control), a managed switch adds advanced features usually found only in high-end switches:

Rapid Spanning Tree Protocol (RSTP) provides a standardized network redundancy scheme with improved network recovery time over Spanning Tree Protocol (STP).

RapidRing[®] provides high speed network redundancy — allowing recovery from a link loss in under 300 ms.

VLAN allows the physical network to be configured as multiple virtual local area networks — limiting broadcast/multicast domains and improving performance.

Trunking allows ports to be associated in groups — each group functioning as a high-speed backbone to another managed switch.

QoS provides message priority with one of these priority schemes — port-based, MAC-based, 802.1p, DiffServ, or TOS.

Rate Control allows variable data rates by port for bandwidth allocation.

Port Security limits port traffic to only those devices with listed MAC addresses.

IGMP Snooping allows multicasts to be limited to only relevant ports.

Port Mirroring copies traffic from one or more ports to a monitoring port.

Programmable Fault Relay provides a dry contact to a supervisory system if the switch senses a condition such as the loss or addition of a link.

Non-blocking wire-speed operation provides a maximum data rate of 148,810 packets per second for 100 Mbps Ethernet on all ports at full duplex.

These features place managed switches from Contemporary Controls among the most powerful and versatile of Industrial Ethernet switches. Configuration is done through a web browser or a console port. Individual port parameters (data rate, duality, flow control) can be pre-set or auto-negotiated. Auto-MDIX can be disabled, if desired. Each port supports the PAUSE function for full-duplex links, and uses the backpressure scheme for half-duplex links.

Each switch is powered from wide-range, low-voltage AC or DC sources — and redundant power connections are available for backup considerations. Each comes with the attachments for either DIN-rail or panel mounting. The front panel features a power LED, a management status LED and bi-color LEDs for the link status, activity, and data rate of each port.

This software is used in EI-WPMR-M products.

3.1 Software

The provided CD-ROM contains:

- This Software Manual for Web Browser
- A Software Manual for Console Access
- An Installation Guide
- An Ethernet Glossary
- Additional information of interest

3.2 Product Images



EI-WPMR-M

4 Advanced Operation

4.1 General Considerations

Configuration is accomplished while the switch is connected to a computer running a web browser that accesses the switch's onboard web server.

4.1.1 LEDs

To aid in troubleshooting, several LEDs have been provided.

Each **port LED** glows solid if a link exists, flashes to show activity and shows data rate by color — green for 100 Mbps and yellow for 10 Mbps.

The **Power LED** glows solid green to indicate the presence of adequate power.

On *all* models the **Status LED** on the front panel of the switch glows solid green when operation is fault-free. For EI-WPMR-M models, this LED flashes if a fault occurs.

4.1.2 Accessing the Web Server

4.1.2.1 Web Browser

The switch contains an interactive web server, accessible from any Internet-compatible PC on the local network. It is compatible with recent versions of Internet Explorer (5.0 or later, suggested) or Netscape Navigator (7.1 or later, required). It is factory-programmed with a default IP address of 192.168.92.68 and a Class C subnet mask of 255.255.255.0. Changing the switch IP address is strongly encouraged.

4.1.2.2 Initial Access

The hardware arrangement for initial setting of the switch IP address by web browser appears in Figure 1. The PC should be temporarily disconnected from the Ethernet LAN in case the default address of the switch matches that of a device on the existing LAN. Therefore, the procedure for changing the IP address of the switch employs the creation of a temporary LAN composed of nothing but the switch, the PC used to configure it and a CAT5 cable that connects the PC to any switch port.

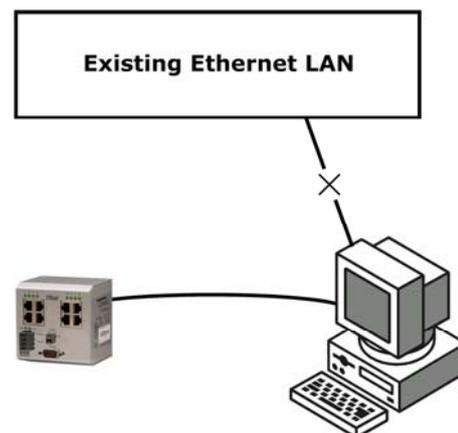


Figure 1 — Setup for Initial IP Address Configuration by Web Browser

For initial configuration, the PC chosen for the procedure should temporarily have its IP address modified as shown in Figure 2 — which employs a Windows 2000 example.

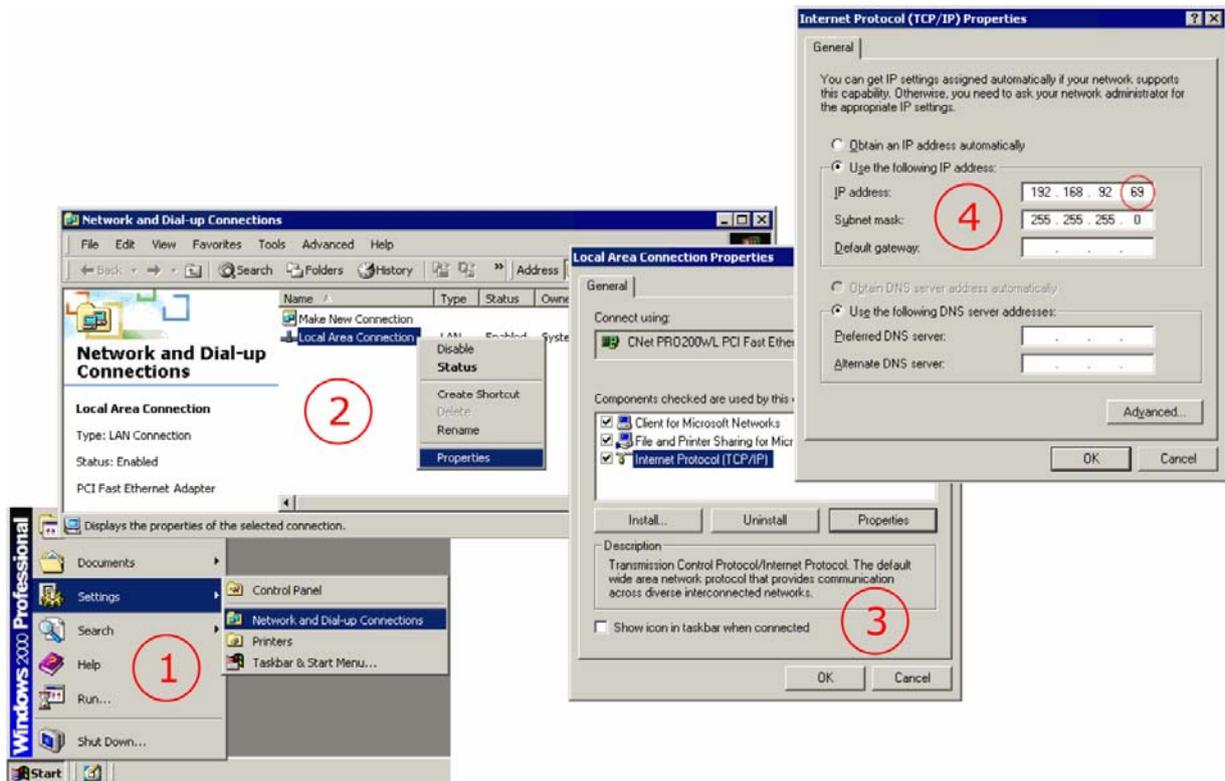


Figure 2 — Steps for Changing the IP Address of the PC Used for Setup

The example in Figure 2 suggests an IP address for the PC of 192.168.92.69, but the final quad of the address could be any value from 1 to 255 — except for 68 which is used by the switch. After the IP address of the PC has been set to the same LAN as the switch, a web browser can access the switch via its default IP address.

Upon accessing the switch, the screen of Figure 3 appears and prompts for **Username** and **Password** — both of which are blank by default. The first time through, leave both of these fields blank and click “Submit”.

Enter Username & Password:	
Username	<input type="text"/>
Password	<input type="text"/>
<input type="button" value="Submit"/>	

Figure 3 — Login Screen

Following the login screen, a Home Page similar to that in Figure 4 appears — **but some 10–12 seconds** for “downloading picture” will be needed before all imagery are displayed. If needed, the MAC Address of the switch is listed on this page.

EI-WPMR Information:	
Name	Contemporary Controls EI-WPMR Managed Switch. V5.0.6
Location	
Contact	
MAC Address	00-50-DB-00-22-64
Firmware Version	5.0.6
Uptime	1469 seconds
Switch Temperature	47°C

Note: The port and relay status and settings can be accessed by clicking on their images above.

Figure 4 — Web Server Home Page

Clicking on **System Configuration** in Figure 4 displays the menu options of Figure 5.

- [System Configuration](#)
- [Configure IP Address](#)
- [Configure Trunking](#)
- [Configure Port Mirroring](#)
- [Configure VLAN](#)
- [Configure Filtering and Forwarding Table](#)
- [Configure Quality of Service](#)
- [Configure Fault Relay](#)
- [Configure Redundancy](#)
- [Configure Rate Control](#)
- [Configure Port Security](#)
- [Configure IGMP Snooping](#)

Figure 5 — System Configuration Menu

Choose **Configure IP Address** to display the screen shown in Figure 6.

TD070446-0MA

By default, the IP Address configuration method is “Fixed”. Either leave it as is or choose the “DHCP” method of changing the IP address. If the “DHCP” option is selected, it will only take effect after the switch is connected to a network which contains a DHCP server. At that point, any of various network procedures must be employed to ascertain the IP Address assigned by the DHCP server. Otherwise, Console Port access via a serial cable can be used to learn the IP Address — as described in the *Software Manual for Console Access* located on the CD-ROM containing this manual.

Configure IP Address:	
Assign by	<input checked="" type="radio"/> Fixed <input type="radio"/> DHCP
IP Address	192.168.92.68
Subnet Mask	255.255.255.0
Default Gateway	10.0.0.1

Figure 6 — The Default IP Address

After the switch has been given its initial configuration, it will be ready for use in the full original LAN. The temporary LAN should then be dismantled and the PC re-configured to restore its original IP address.

4.1.3 On-Screen Help

There are many configuration screens. The upper-right portion of each screen contains a context-sensitive **Help** option. Clicking this option launches another browser window which contains helpful information about the current screen and which pertains to a certain product series. The example of Figure 7 pertains to the EI-WPMR-M Series.

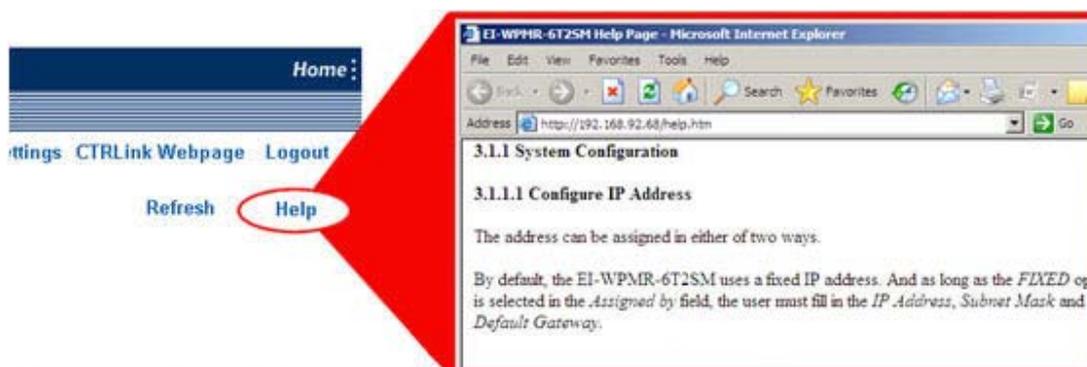


Figure 7 — Help Window

4.1.4 Username and Password

Each time the switch is accessed by a web browser, the Login screen of Figure 8 will appear.

Enter Username & Password:	
Username	<input type="text"/>
Password	<input type="password"/>

Figure 8 — Login Screen

NOTE: As a safeguard, *Username* and *Password* can **only** be modified by a terminal-emulation program communicating via a null-modem cable and the Console Port. If this is not done, web browser secure access cannot be achieved.

Both *Username* and *Password* are case sensitive and can be any combination of printable characters. The number of characters in each string can range from 0 (blank) through 10 — and the two strings can match each other, if desired.

Once modified via the Console Port, future access by web browser will be denied unless the correct information is entered at Login. Attempting to access the switch with an incorrect *Username* or *Password* via web browser will cause the switch to fail to download the Home Page and a blank Login screen will remain onscreen.

If either string is forgotten, they can be reset to their default (blank) values by shorting pins 2 and 3 of the Console Port (Figure 9) for **10 seconds** during boot. This will reset **only** the *Username* and *Password* (**and only for this current login session**) — remaining parameters will not be changed. After login, the *Username* and *Password* should be changed and saved to EEPROM to overwrite the old values. If the access strings are not changed and saved, the reset process must be repeated on power cycle.

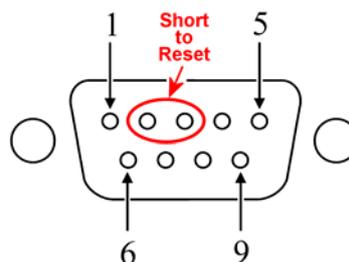


Figure 9 — Restoring the Default Access Strings

4.1.5 Restoring Factory Default Settings

For security purposes, this function can only be performed via Console Access. Consult the **Main Menu** options under *Software Manual for Console Access* for the procedure.

4.2 Home Page (Main Menu)

After login, the **Home Page** (Figure 10) can be accessed from any page via the “Home” link at the right edge of the banner. The Home Page identifies the switch *Name*, *Location*, *Contact* person, *MAC Address*, *Firmware Version*, *Uptime* and *Temperature*. Immediately beneath the banner, the following links are displayed:

System Configuration	(explained in Section 4.2.5)
SNMP Configuration	(explained in Section 4.2.5.8.3)
Performance Monitor	(explained in Section 4.2.7)
Save Settings	(explained below)
CTRLink Webpage (Internet access required)	(explained below)
Logout	(explained below)

Save Settings will store any currently modified settings to the non-volatile memory contained in the switch. For storage to a PC, see Section 4.2.1.

NOTE: If the “Save Settings” option is NOT selected after modifications have been made, any modified settings will be lost when a power cycle occurs.

CTRLink Webpage links to the CTRLink home page for further helpful information.

Logout restarts the session and prompts the user for a *Username* and *Password*.

CTRLink® EI-WPMR Series Home

System Configuration SNMP Configuration Performance Monitor Save Settings CTRLink Webpage Logout

CTRLink Compact Managed Switching Hub Refresh Help

EI-WPMR Information:

Name	Contemporary Controls EI-WPMR Managed Switch. V5.0.6
Location	
Contact	
MAC Address	00-50-DB-00-22-64
Firmware Version	5.0.6
Uptime	1469 seconds
Switch Temperature	47°C

The diagram shows the physical switch with ports 1-4 on the left and 5-8 on the right. It also shows power inputs for 10Vdc36 and 8Vac24, a console port, and status indicators for PWR, STAT, and Fault.

Note: The port and relay status and settings can be accessed by clicking on their images above.

Figure 10 — Main Menu

Also, individual **Port Configuration** and **Port Statistics** (Section 4.2.4) are invoked from the Home Page by clicking on the image of the port socket of interest. Before

clicking — as the cursor hovers over a picture element — a helpful Tool Tip appears to confirm the item being accessed.

4.2.1 Upload/Download Settings

A web browser can save the switch settings to the EEPROM onboard the switch. The settings can also be stored to a PC and retrieved from a PC, but this can only be done via the Console Port — not by a web browser. For a description of the process, refer to the *Software Manual for Console Access* on the CD-ROM.

4.2.2 Uptime

Uptime displays the number of seconds since the previous hard power recycle or soft restart. The **Refresh** option updates the value (the counter only resets due to a restart).

4.2.3 Switch Temperature

Switch Temperature displays the current internal temperature of the switch in degrees Celsius ($\pm 3^\circ$). Select the **Refresh** item to update the currently displayed value.

4.2.4 Port Configuration and Port Statistics

To configure a port or view its statistics, go to the switch Home Page and click on the image of the socket associated with the port of interest. Figure 11 shows the default settings for port 4, as an example, with controls in the upper panel and typical traffic statistics in the lower panel. The port number appears in the top center of the screen.

4.2.4.1 Port Configuration (Figure 11 upper panel)

With the upper-panel controls, the user can turn the port on or off (*Port State*), set its data rate and duplicity (*Mode*), manage its *Auto-MDIX* function and enable or disable its *Flow Control*. The type of cabling which the port supports is identified in the box to the right of the word *Media Type*.

NOTE: Some models include **fiber optic cable** ports. Such ports do not auto-negotiate. Their only *Mode* option is Full or Half Duplex because their data rate fixed at 100 Mbps.

Auto-Negotiation for Copper Ports

A single cable links two Ethernet devices. When these devices auto-negotiate, the data rate will be 100 Mbps only if both are capable of that speed. Likewise, full-duplex will only be selected if both can support it. If only one device supports auto-negotiation, then it will match the data rate and duplex mode of the non-auto-negotiating device. Sometimes it is advantageous to select a fixed data rate and duplicity setting on both Ethernet devices to eliminate the auto-negotiation process.

Auto-MDIX

When interconnecting two switches, crossover cables are traditionally used — but if one switch uses Auto-MDIX, the communication can be via either straight-through cable or crossover cable. This functionality does not require *both* switches to have Auto-MDIX.

Current State:	
Port State	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
Media Type	Copper
Mode	Auto Negotiate
Auto-MDIX	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
Flow Control	<input type="radio"/> Enable <input checked="" type="radio"/> Disable

Apply

Port Packet Statistics:	
Unicast Packets Received	32855
Unicast Packets Sent	25039
Multicast Packets Received	40993
Multicast Packets Sent	0
Broadcast Packets Received	130834
Broadcast Packets Sent	129
Dropped Packets	0
Oversize Packets	0
Undersize Packets	0
Fragments	0
Jabbers	0
Collisions	0
Deferred Transmissions	0

Figure 11 — Port Configuration and Port Statistics

4.2.4.2 Port Frame Statistics (Figure 11 lower panel)

These numbers will remain static until the **Refresh** option is used to provide an update. The displayed values are the total number of these events from when the switch was last powered-up, its IP address was redefined or its parameters were reset to their default values. Recycling power, redefining the IP address or resetting parameters to their default values will reset the **Port Frames Statistics** to zero.

Port Frame Statistics are available for the following counts:

<i>Unicast Frames Received</i>	This counts unicast frames received by the switch.
<i>Unicast Frames Sent</i>	This counts unicast frames transmitted by the switch.
<i>Multicast Frames Received</i>	This counts multicast frames received by the switch.
<i>Multicast Frames Sent</i>	This counts multicast frames transmitted by the switch.
<i>Broadcast Frames Received</i>	This counts broadcast frames received by the switch.
<i>Broadcast Frames Sent</i>	This counts broadcast frames transmitted by the switch.
<i>Dropped Frames</i>	This counts the frames dropped on transmission due to an excessive number of collisions.
<i>Oversize Frames</i>	This counts the frames dropped on reception because they exceeded 1518 bytes — often due to faulty drivers.
<i>Undersize Frames</i>	This counts “runt” frames dropped on reception because they were under 64 bytes with no valid CRC/FCS. Runt frames are usually created by collisions.
<i>Fragments</i>	This counts the number of fragmented frames.
<i>Jabbers</i>	These frames exceed 1518 bytes, have invalid CRC/FCS values and are due to constant transmissions from a network interface card which likely is faulty.
<i>Collisions</i>	This counts the collisions on a half-duplex segment — most often due to physical-layer issues.
<i>Deferred Transmissions</i>	This counts the frames delayed on the first transmission attempt because the media was busy.

4.2.5 System Configuration

Each of the **System Configuration** menu selections of Figure 12 will activate additional sub-menus from among the following list:

Configure IP Address	(explained in Section 4.2.5.1)
Configure Trunking	(explained in Section 4.2.5.2)
Configure Port Mirroring	(explained in Section 4.2.5.3)
Configure VLAN	(explained in Section 4.2.5.4)
Configure Filtering & Forwarding Table	(explained in Section 4.2.5.5)
Configure Quality of Service	(explained in Section 4.2.5.6)
Configure Fault Relay	(explained in Section 4.2.5.7)
Configure Redundancy	(explained in Section 4.2.5.8.3.4)
Configure Rate Control	(explained in Section 4.2.5.9)
Configure Port Security	(explained in Section 4.2.5.10)
Configure IGMP Snooping	(explained in Section 4.2.5.11 below)



Figure 12 — System Configuration Menu

4.2.5.1 Configure IP Address

Figure 13 displays the **Configure IP Address** menu with its default values. The address can be assigned in either of two ways.

Configure IP Address:	
Assign by	<input checked="" type="radio"/> Fixed <input type="radio"/> DHCP
IP Address	<input type="text" value="192.168.92.68"/>
Subnet Mask	<input type="text" value="255.255.255.0"/>
Default Gateway	<input type="text" value="192.168.92.1"/>

Figure 13 — Configure IP Address Menu

By default, the switch uses a fixed IP address. And as long as the *FIXED* button is selected in the *Assign by* field, the user must fill in the *IP Address*, *Subnet Mask* and *Default Gateway*. The switch can also have its *IP Address* assigned *automatically* by a Dynamic Host Configuration Protocol (DHCP) server if the *DHCP* button is selected. The response to clicking the *Apply* button is as follows.

Using Fixed configuration: When the *Apply* button is clicked, the switch restarts using the new IP configuration, the *Uptime* clock restarts and a redirect page displays:

The new IP address is [[hyperlink to new IP Address](#)]

Clicking on the hyperlink begins a *new session* as the Login screen appears.

Using DHCP: When the *Apply* button is clicked, the *Uptime* clock restarts and a redirect page displays the following:

IP Address assigned by DHCP.

At this point, further browser access to the switch will not be possible until some method is used to learn the new address assigned by the DHCP (a process of many seconds — perhaps half a minute, depending on network conditions).

If the switch fails to contact a DHCP server within two minutes, the message

“No Server has been found. Set IP manually”

will appear for 5 seconds on the **Console Port** — followed by the Login screen. In this case, the Figure 13 values for *IP Address*, *Subnet Mask* and *Default Gateway* will all be zeroes until new user values are entered.

Note for SNMP: The “Warm Start” trap is transmitted once the address assignment has been made — immediately for *Fixed* address assignment but after a delay of many seconds for *DHCP* address assignment.

4.2.5.2 Configure Trunking (Copper Ports Only)

Port Trunking allows two or more ports to be grouped with the resulting group behaving as a single logical link. A managed switch supports multiple trunks — each constructed of 2 or more fixed physical ports.

To keep frames in order, packets with the same source/destination MAC addressing are sent over the same trunk path — but the reverse path may follow a different link because a hash algorithm is used to balance the load between links in a trunk.

Adding more ports (links) to a trunk group will increase the communication bandwidth between two switches.

Port Trunking on managed switches from Contemporary Controls also provides redundancy with a fast recovery time (several milliseconds). If a link in the trunk group is lost, the remaining links immediately take over and maintain communication between the switches.

Figure 14 illustrates three two-link trunks (A, B and C) connecting two computers through four switches. This configuration could sustain a link loss in Trunk A, but within milliseconds a redundant data path would be reconstructed between the two computers. A similar recovery would manifest for a link lost in Trunk B or C. Indeed, even multiple link losses — one in each trunk group — would not disrupt communication between the two end stations except for the brief recovery time.

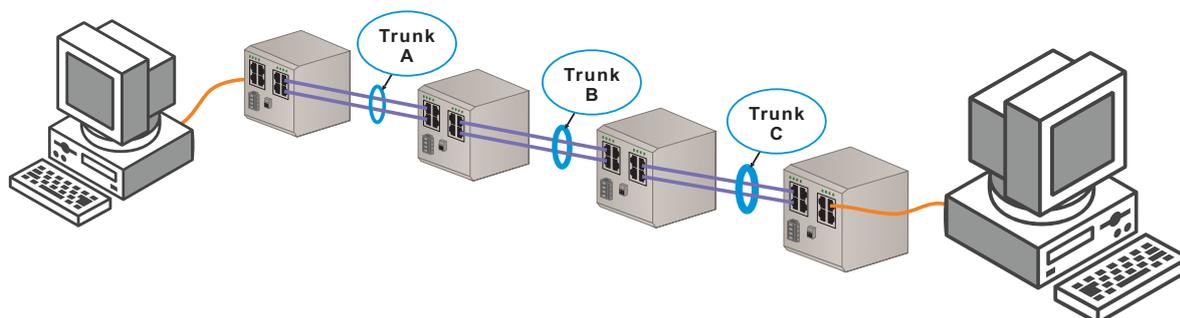


Figure 14 — Trunking

Figure 15 displays a sample **Configure Trunking** screen with two trunk groups defined. The sample shows that only *Group 1* is currently active. It would be permissible for both of these two sample groups to be enabled simultaneously — but only because they do not have any ports in common (overlapping ports).

Configure Trunking		
Group 1	Port: <input checked="" type="checkbox"/> 1 <input checked="" type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
Group 2	Port: <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input checked="" type="checkbox"/> 5 <input checked="" type="checkbox"/> 6 <input checked="" type="checkbox"/> 7 <input type="checkbox"/> 8	<input type="radio"/> Enable <input checked="" type="radio"/> Disable

Submit

Figure 15 — Configure Trunking

TD070446-OMA

4.2.5.3 Configure Port Mirroring

Port mirroring allows a port to *copy* or “mirror” traffic for one or many ports — useful when a diagnostic tool is used. Switches normally send frames only to ports involved in a conversation, so mirroring is required if a diagnostic tool must capture network traffic. Various settings control how data is mirrored, but apply *only* if (in Figure 16, upper panel) the *Status* option is enabled and a *Mirror Port* (to which the data is copied) is specified.

Ingress (Ingoing) Mirroring Rules (center panel) apply to data *received* by the *Source Ports* being mirrored; **Egress (Outgoing) Mirroring Rules** (lower panel) for data *sent* by *Source Ports* being mirrored. Both sets of rules allow a *MAC Address Filter* and a *MAC Address* to be defined. The *Divider* specifies the *fraction* of frames copied from each source port.

The **MAC Address** options are:

All — mirrors *all* traffic regardless of MAC address. This is the default.

Source — mirrors frames whose *source* addresses match the entered value.

Destination — mirrors frames whose *destination* addresses match the entered value.

Configure Port Mirroring	
Status	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
Mirror Port	Port: <input checked="" type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5 <input type="radio"/> 6 <input type="radio"/> 7 <input type="radio"/> 8
Ingress (Ingoing) Mirror Rules:	
Source Ports	<input type="checkbox"/> 1 <input checked="" type="checkbox"/> 2 <input checked="" type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8
Divider	1 (1-1023)
MAC Address Filter	<input type="radio"/> Capture ALL <input checked="" type="radio"/> Capture by Source <input type="radio"/> Capture by Destination
MAC Address	0050DBFF0000
Egress (Outgoing) Mirror Rules:	
Source Ports	<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input checked="" type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8
Divider	10 (1-1023)
MAC Address Filter	<input checked="" type="radio"/> Capture ALL <input type="radio"/> Capture by Source <input type="radio"/> Capture by Destination
MAC Address	

Figure 16 — Configure Port Mirroring

The Figure 16 example has ports 2 and 3 as *Ingress Source Ports* and port 4 as the sole *Egress Source Port* — all to be copied to port 1. Since the *Ingress Divider* is 1, each frame from ports 2 and 3 will be mirrored — but only if its source address matches 0050DBFF0000. The *Egress Divider* value is 10, so only 1 in 10 frames transmitted by port 4 will be mirrored — and since the *Egress MAC Address* setting is *All*, the source/destination address field in the transmitted frames will be ignored.

4.2.5.4 Virtual Local Area Networks (VLANs)

A VLAN (Virtual Local Area Network) is comprised of devices grouped on some basis other than geographic location (i.e., by work group, security level, user type, or application). The devices *logically* behave as if tied to the same wire although they may be physically located on very different LAN segments. VLANs are configured with software, which offers much greater flexibility than hardware configuration.

A chief advantage of VLANs is that they block *broadcasts* and *multicasts* from non-VLAN ports. Most switches tend to transmit *unicast* frames sent only to ports involved in a conversation (directed messages) and cannot accommodate broadcast or multicast frames. VLANs keep broadcasts and multicasts within a VLAN group.

Another advantage of VLANs is that despite being physically relocated, a device can remain in the same VLAN — with no hardware reconfiguration needed. The VLAN supervisor can change/add workstations and manage load-balancing (bandwidth) far more easily than with a LAN modified only by hardware. Management software maintains a virtual image of how the logical and physical networks compare.

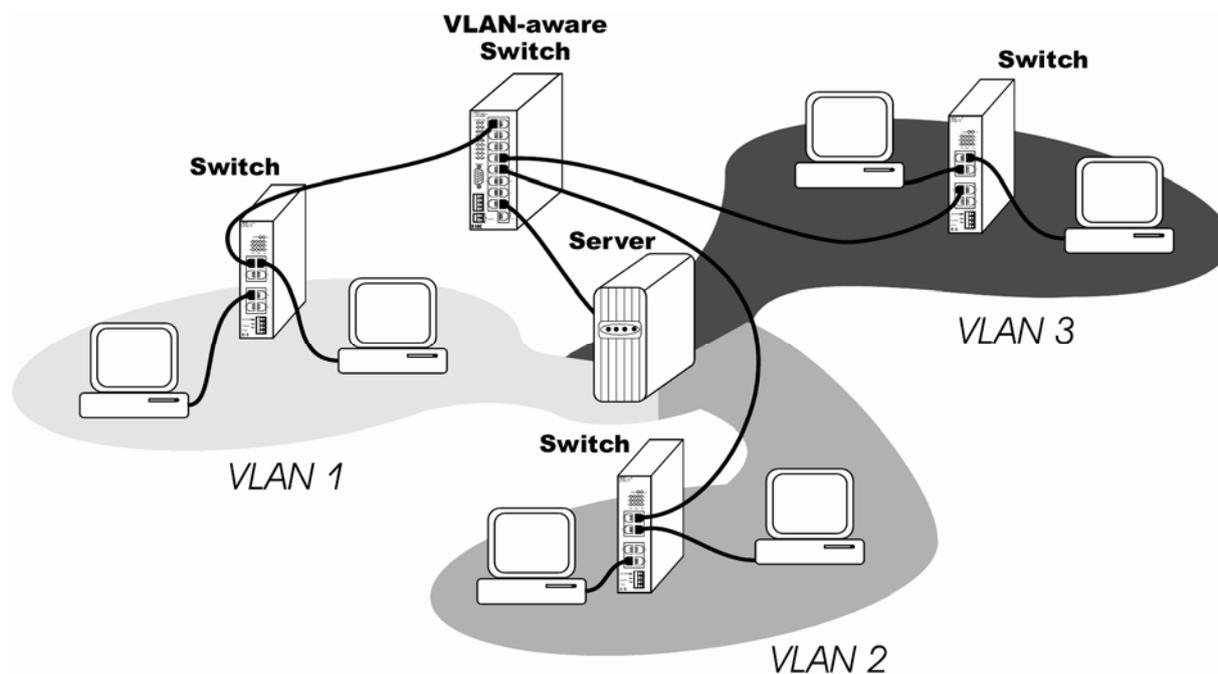


Figure 17 — VLANs

4.2.5.4.1 All Ports Should Be VLAN Ports

When VLANs are enabled on the switch, all ports should be assigned to one or more VLANs. Such ports are called *VLAN ports*. If a port is *not assigned* to a VLAN while VLANs are enabled, that port **cannot receive messages** from the switch. A frame received from a VLAN port will *only* be forwarded to those ports with which it shares a VLAN membership. If the destination belongs to another VLAN, the frame will be discarded. This topology allows networks to share a common server or router, but use different VLANs for security or performance reasons.

4.2.5.4.2 VLAN Tags and VLAN Identifiers (VIDs)

Each VLAN frame contains an 802.1Q VLAN tag having a *VID* (VLAN Identifier) indicating to which VLAN this message belongs. The switch can be configured to allow frames with specific VIDs to be received on specific ports within a VLAN.

VID values can range from 1 to 4094 — but only *within one contiguous block* of 512 values. The allowable VID blocks (ranges) are:

1–511	1024–1535	2048–2559	3072–3583
512–1023	1536–2047	2560–3071	3584–4094

Packets having VID values outside of the one defined block will be dropped.

4.2.5.4.3 Two Types of VLANs

The managed switch supports two types of VLANs, Port VLAN and 802.1Q VLAN. A **Port VLAN** is normally used to interconnect VLAN-unaware devices (such as desktop computers) which do not use VLAN tags. But **802.1Q VLANs** require 802.1Q tags in the frames passing through the switch.

4.2.5.4.4 VLANS and Trunking

A problem can occur if all of the ports in a trunk do not share the same VLAN. If two trunk ports were in different VLANs and one port suffered a link failure, frames would then pass through the other port. As a result, the frames would be discarded. If **more** than two ports support a trunk, the problem is not quite the same because when a path fails, the alternate path is *not user selectable* — thus, the alternate port might or might *not* be in the same VLAN unless **all** of the ports in the trunk were in the same VLAN.

Therefore, the rule when using trunking and VLANs is for *all* ports in the trunk to be in the *same* VLAN and have the *same* default VID number — and every port in potential use should have a VLAN defined for it.

4.2.5.4.5 Core Switches and Edge Switches

A **core switch** is connected only to devices that are VLAN aware — thus, all frames received by a core switch should *already contain* 802.1Q VLAN tags. An **edge switch** *adds* VLAN tags to frames sent by non-VLAN aware devices and it *removes* VLAN tags from frames destined for non-VLAN aware devices. Figure 18 illustrates how core and edge switches differ in their placements within a typical network.

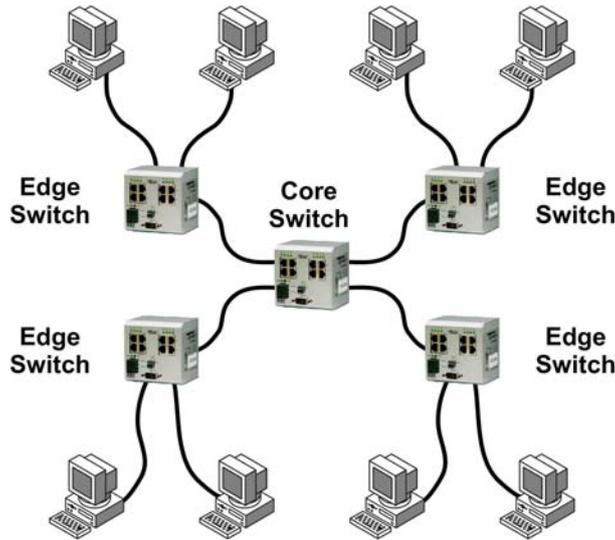


Figure 18 — Core Switches and Edge Switches

To function in VLAN mode, the switch *requires* VLAN tags. When it performs as a *Port VLAN* switch (connected **only** to non-VLAN aware devices that do not use VLAN tags), a *default tag* will be applied to the untagged frames entering the switch. When these frames leave the switch, the tags should be removed. The switch can act as either a core or an edge switch on a port-by-port basis — and this functionality allows the switch to isolate non-VLAN aware devices by tags and impart added security.

Figure 19 shows how an edge switch **adds** a VLAN tag to an untagged frame from PC1 so it can be used by the core switch and thus all of the VLAN, but the edge switch also **removes** the tag from the frame destined for the VLAN-unaware PC2.

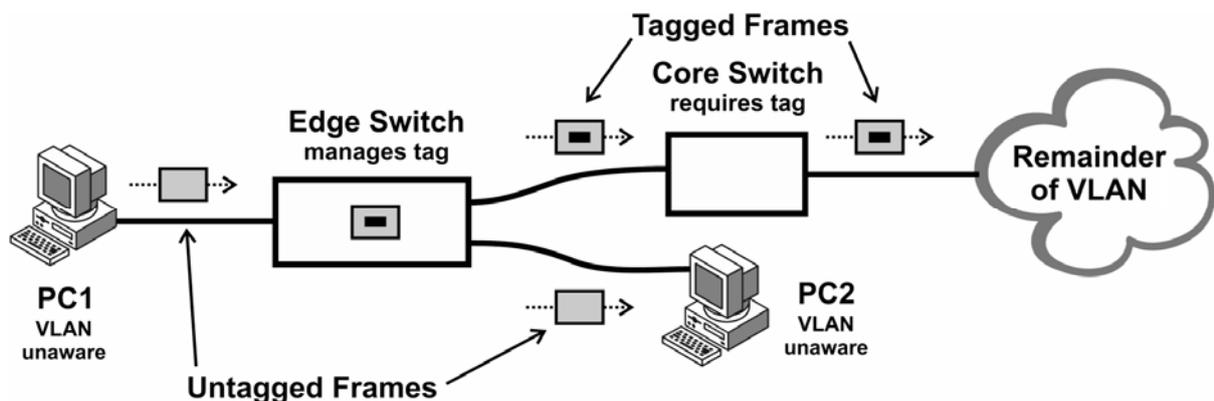


Figure 19 — Edge Switches Add or Drop VLAN Tags as Needed

4.2.5.4.6 Creating VLAN Groups and VIDs

When creating VLANs, several steps and two screens are required to complete the configuration. The first step is to define each VLAN *Group* using the **Configure VLAN Group and VID** screen shown Figure 20 (depicting an 8-port switch) — where the first 7 groups are defined. The last 2 groups are in the top panel of **VLAN Settings Page 2** — and accessible via the link at the bottom of this screen. Each group can have its *Status* “Enabled” or “Disabled” and is assigned a default *VID*. Define the *Members* of each *Group* (up to 24 ports on some switches) via a drop-down option for each port in the *Group*. Within a *Group*, each port can be either omitted or included in the group — with or without an egress *Tag Filter* being applied.

When a VLAN frame leaves the switch, take care regarding the VLAN tag it contains. In Port VLAN mode as shown in Figure 19, the switch will insert a VLAN tag into any frame arriving from a non-VLAN-aware device. When the frame exits the switch, the tag will still be present and could cause a problem if the receiving device is not VLAN-aware.

In an 802.1Q-compliant network, does the unit act as a core switch or an edge switch? If acting as a **core** switch, VLAN tags should be kept in the frame. For an **edge** switch, VLAN tags should be removed from those ports connected to non-VLAN-aware devices. (The switch can act as both a core and edge switch on a port-by-port basis.)

To remove the tag from a frame leaving the switch via a port tied to a non-VLAN-aware device such as PC2 in Figure 19, choose the “MF” drop-down option as has been done for most ports in Figure 20. (The only port **not** filtering egress tags is port 2 of group 1). A **core** switch should have **all selected ports** set to “MN” since each of its ports needs to pass frames with tags intact (unfiltered).

Configure VLAN Group and VID										
Group	VID	Members and Tag Filter								Status
		-- = Non Member, MN = Member without Filter, MF = Member with Filter								
1 Apply	1	1[--]	2[MN]	3[MF]	4[MF]	5[--]	6[--]	7[--]	8[--]	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
2 Apply	2	1[--]	2[--]	3[--]	4[MF]	5[--]	6[--]	7[MF]	8[MF]	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
3 Apply	3	1[MF]	2[--]	3[--]	4[--]	5[--]	6[MF]	7[MF]	8[--]	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
4 Apply	4	1[--]	2[--]	3[--]	4[--]	5[--]	6[--]	7[--]	8[--]	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
5 Apply	5	1[--]	2[--]	3[--]	4[--]	5[--]	6[--]	7[--]	8[--]	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
6 Apply	6	1[--]	2[--]	3[--]	4[--]	5[--]	6[--]	7[--]	8[--]	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
7 Apply	7	1[--]	2[--]	3[--]	4[--]	5[--]	6[--]	7[--]	8[--]	<input type="radio"/> Enable <input checked="" type="radio"/> Disable

[Continue to VLAN Settings Page 2](#)

Figure 20 — Configure VLAN Groups and VIDs

4.2.5.4.6.1 Management Port

Below each group's *Members* area in Figure 20 is the "Management Port" checkbox. If a VLAN is to have full management functionality for any of its ports, it must include the "Management Port" as one of its members. This port accomplishes the management logic solely within the Ethernet controller chip and outside that chip it has no physical presence whatsoever. It extends management functionality only to the ports with which it communicates. A VLAN that *excludes* the "Management Port" can bestow added security to its ports — but it cannot support IGMP Snooping, the web server or SNMP.

4.2.5.4.6.2 Example of VLAN Configuration

See Figure 20 where only the first three *Groups* have been defined as follows:

- *Group 1* consists of ports 2, 3, 4 and the "Management Port".
- *Group 2* consists of ports 4, 7 and 8.
- *Group 3* consists of ports 1, 6, 7 and the "Management Port".

For the *Groups* listed above (and assuming VLANs are enabled), note the following:

Groups 1 and 3 include the "Management Port" — therefore, ports 1, 2, 3, 4, 6 and 7 will have full management functionality when their *Groups* are "Enabled".

Group 3, however, will not currently function because it is "Disabled". Also, because ports 1 and 6 are included **only** in *Group 3*, these two ports can not communicate **whatsoever** as long as their *Group* is "Disabled".

Group 2 is "Enabled", but it does not include the "Management Port" — so ports 4, 7 and 8 seem to have no access to management functions. However, port 4 is an **overlapping port** and still has management due to its membership in *Group 1*. Port 8 will be **unmanaged** as long as it is defined **only** in a VLAN that excludes the "Management Port" — and that same VLAN is enabled. If *Group 2* were disabled, its members could be managed.

Finally, port 5 is not in any VLAN so it **cannot communicate** unless VLANs are disabled.

4.2.5.4.7 Configure VLAN Frame Drop Rules (Figure 21, second panel)

The switch supports the ability to drop non-802.1Q frames (frames without VLAN tags) on a port-by-port basis. This is a useful feature for core switches because untagged frames could be received due to the improper configuration of an edge switch.

This feature can also **add extra security** because a correct VID value does not guarantee a frame will travel through the switch — **the ingress port must also belong to the defined group** to pass the frame through the switch. When **Drop VID Violation Frame** is *Enabled* and **for only the selected ports**, this is what happens: When a frame arrives at the switch, its VID tag will be examined to confirm that the port through which the frame will enter the switch is part of the group using this tag. If the port does not belong to the group, the frame will be dropped.

4.2.5.4.8 Configure 802.1Q VLAN Tag (Figure 21, third panel)

For frames from a non-VLAN device to function in a VLAN, the default VID tag of the ingress port through which they pass **must match the VID of the group** to which the frames are destined. (The VLAN group must also be *enabled*.)

Example: Assume *Group 3* in Figure 20 has been enabled. Because the *Group* uses a VID value of “3”, a member port *will not pass frames* to fellow ports **unless** the port’s *Default Tag* has also been set to “3” as in Figure 21.

Configure VLAN Group and VID			
Group	VID	Members and Tag Filter <small>-- = Non Member, MN = Member without Filter, MF = Member with Filter</small>	Status
8 <input type="button" value="Apply"/>	8	1 2 3 4 5 6 7 8 <input type="text" value="--"/> <input type="text" value="--"/> <input type="checkbox"/> Management Port	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
9 <input type="button" value="Apply"/>	9	1 2 3 4 5 6 7 8 <input type="text" value="--"/> <input type="text" value="--"/> <input type="checkbox"/> Management Port	<input type="radio"/> Enable <input checked="" type="radio"/> Disable

Configure Non-802.1Q Frame drop Rules:	
Drop VID Violation Frame	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
Drop Non-802.1Q Frames by Ports	<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8

Configure 802.1Q VLAN Tag	
Port	Default Tag
1	15
2	15
3	15
4	15
5	15
6	15
7	15
8	15
M	15

VLAN Status:	
VLAN Status	<input type="radio"/> Enable <input checked="" type="radio"/> Disable

Note: Ensure the settings are correct before enabling or some ports may become inaccessible.

Figure 21 — VLANs Settings Page 2

4.2.5.4.9 VLAN Status (Figure 21, bottom panel)

The final step in VLAN configuration is to *activate global VLAN functionality* for the switch. The default setting is *Disabled*.

4.2.5.4.10 Port VLAN Security Example

Suppose that 3 computers on ports 1, 2 and 3 must talk to a printer on port 4, but not with each other — and none of these devices is VLAN-aware. One way to accomplish this is to define 4 VLANs, assign each included port the option “MF” (enabling its *Tag Filter*) and match the *Default Tags* to the *VIDs* as follows:

<u>Group</u>	<u>VID</u>	<u>Member Ports</u>	<u>Default Tags</u>
1	1	1 & 4	Port 1 = 1
2	2	2 & 4	Port 2 = 2
3	3	3 & 4	Port 3 = 3
4	4	1, 2, 3 & 4	Port 4 = 4

Ports 1, 2 and 3 do not talk to each other because each port uses a different tag. However, each port’s tag is shared with port 4. The printer on port 4 can still talk to ports 1, 2 and 3 using its own tag (4).

4.2.5.5 Configure Filtering and Forwarding Table

An Ethernet switch learns which devices are tied to which ports by monitoring the traffic carried through the switch. This information (MAC addresses and the ports to which they are attached) is stored in the **address table** which holds up to 4000 address/port associations. The switch will only transmit traffic destined for a *registered* address via the specific port associated with that address. The table behavior follows.

During frame reception, the frame's destination address is compared with table entries for a possible match. If a match is found, the port associated with the address is noted from the table and the frame is directed to that port only. If no match is found, the frame is flooded (transmitted) to all ports.

During the reception of a *unicast* frame, the frame's source address is also compared with table entries for a match. If a match is not found, the unregistered address (and the port by which it arrived) will be added to the table. However, the address will NOT be saved if the frame is from the Management Port or if the frame has an error or is illegal in length — or if the table has no room for the new entry.

If a port-device association is not refreshed within the address table's "MAC aging time", its information will be discarded. Figure 22 displays the typical MAC aging time of 300 seconds, but this can be set as high as 1048575 seconds (over 12 days).

Configure Filtering and Forwarding Table	
Global Settings:	
MAC Aging Time (seconds)	<input type="text" value="300"/> (Default = 300, Max = 1048575)
<input type="button" value="Apply"/>	
Configure Multicast Filtering Table	
Configure Static Forwarding Table	

Figure 22 — Configure Filtering and Forwarding Table

The screen of Figure 22 also provides options to display the two additional screens of the **Configure Filtering and Forwarding Table: Configure Multicast Filtering Table** (Section 4.2.5.5.1) and **Configure Static Forwarding Table** (Section 4.2.5.5.2).

4.2.5.5.1 Configure Multicast Filtering

A multicast message is one destined for two or more Ethernet devices. By default, the switch transmits such a message over all of its ports. However, the switch can filter up to ten multicast addresses so that messages sent to these addresses will only exit the switch via certain designated ports. The switch also supports IGMP Snooping which allows automatic filtering of multicast messages for devices that support multicasting as described in Section 4.2.5.11. As shown in the screen of Figure 23, the user must enter the multicast *MAC Address* to represent each multicast group, the *Ports* that will carry messages to that group and the *Priority* of those messages. By default, each address is assigned a *Priority* of “Low”.

Configure Filtering and Forwarding Table				
Multicast Filtering:				
Current Entries: # (10 Maximum)				
		MAC Address	Ports	Priority
Apply	Delete	159159159159	<input checked="" type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8	Normal
Apply	Delete	333333333333	<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input checked="" type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8	High
Apply	Delete	456123789456	<input checked="" type="checkbox"/> 1 <input type="checkbox"/> 2 <input checked="" type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8	Medium
Apply	Delete	7766a30dc22e	<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input checked="" type="checkbox"/> 4 <input checked="" type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input checked="" type="checkbox"/> 8	Low
Apply	Delete	0105e0000000	<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input checked="" type="checkbox"/> 7 <input type="checkbox"/> 8	Low
Apply	Delete		<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8	Low
Apply	Delete		<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8	Low
Apply	Delete		<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8	Low
Apply	Delete		<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8	Low
Apply	Delete		<input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7 <input type="checkbox"/> 8	Low

Figure 23 — Configure Multicast Filtering

When defining or editing a group, the “Apply” button must be clicked to register settings for that group — **before** editing another group or refreshing the screen — or the edits will be lost. When editing many addresses per session, many browser screens will accumulate in the browser’s history — thus, proceeding to another switch function is best done with the onscreen navigation links, instead of the browser’s “back” button.

For an address to be accepted as a valid **multicast** address, its second digit must be odd. If the second digit is even, the error screen of Figure 24 will result.

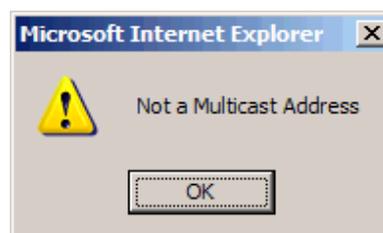


Figure 24 — Not a Multicast Address

4.2.5.5.2 Configure Static Forwarding

The forwarding (address) table (Figure 25) can hold not only learned addresses, but also up to 30 *static* unicast addresses (on two screens, each of 15 addresses). Static addresses perform as if learned, but are not subject to the aging process. *Priority* is applied to messages sent to each defined MAC address via its specified *Egress Port*.

When defining or editing a group, the “Apply” button must be clicked to register settings for that group — **before** editing another group or refreshing the screen — or the edits will be lost. When editing many addresses per session, many browser screens will accumulate in the browser’s history — hence, proceeding to another switch function is best done with the onscreen navigation links, instead of the browser’s “back” button.

Configure Filtering and Forwarding Table				
Static Forwarding: Current Entries: # 1-15 (30 Maximum)				
		Static MAC Address	Egress Port	Priority
Apply	Delete	123456789132	Port 1	High
Apply	Delete	222222222223	Port 6	Normal
Apply	Delete	0a1b12300232	Port 4	Medium
Apply	Delete	000000000000	Port 7	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low
Apply	Delete		Port 1	Low

[Go to Static Forwarding Page 2](#)

Figure 25 — Static Forwarding Table

4.2.5.6 Configure Quality of Service (QoS)

In addition to the MAC-based priority applied in multicast filtering and static forwarding, the switch can assign other types of priority to its traffic to achieve various levels in what is known as **Quality of Service (QoS)**. It can do this regardless of frame content (*Port QoS*) or by examining the content of every frame received by a port and assigning priority based on the port of origin. The default screen displayed in Figure 26, shows that *QoS Status* must be **enabled** before any (or all) of these methods can be applied. At the bottom of the screen there are links to three configuration subscreens — one for each type of QoS.

Configure Quality of Service (QoS):	
QoS Status	<input type="radio"/> Enable <input checked="" type="radio"/> Disable

[Configure Port QoS](#)
[Configure 802.1p Priority](#)
[Configure TOS/DiffServ Priority](#)

Figure 26 — Configure Quality of Service (QoS)

Although not recommended, all types of priority can be active simultaneously — giving rise to conflicts. The following hierarchy shows how these conflicts are resolved:

1. If *Port QoS* is enabled, apply its rules — otherwise,
2. If *TOS/DiffServ Priority* is enabled, apply its rules — otherwise,
3. If *802.1p Priority* is enabled, apply its rules — otherwise,
4. Apply the MAC-based priority used in multicast filtering and static forwarding.

4.2.5.6.1 Configure Port QoS

When QoS is enabled, *Flow Control* for each port can be enabled or disabled and two levels of *Port Priority* applied. This is the highest method of priority and it is known as **Port QoS**. For QoS to be most effective, it is recommended that each port have its flow control disabled (the default setting). When a port is operating in half-duplex mode, flow control is accomplished with backpressure. But in full-duplex mode, flow control is accomplished using the PAUSE protocol. Either method of flow control can affect the ability of a port to deliver messages and can cause some messages to be delayed.

In the example of Figure 27, *Port 2* has been configured for “High” priority while all other ports are set to “Normal”. But *Port 1* (set to “Normal” priority) is the only port on which Flow Control will function — because only it has had *Flow Control* enabled.

At the bottom of the screen there are two links to the other types of QoS.

Configure Port QoS:		
Port	Flow Control	Port Priority
1	<input checked="" type="radio"/> <i>Enable</i> <input type="radio"/> <i>Disable</i>	<input checked="" type="radio"/> <i>Normal</i> <input type="radio"/> <i>High</i>
2	<input type="radio"/> <i>Enable</i> <input checked="" type="radio"/> <i>Disable</i>	<input type="radio"/> <i>Normal</i> <input checked="" type="radio"/> <i>High</i>
3	<input type="radio"/> <i>Enable</i> <input checked="" type="radio"/> <i>Disable</i>	<input checked="" type="radio"/> <i>Normal</i> <input type="radio"/> <i>High</i>
4	<input type="radio"/> <i>Enable</i> <input checked="" type="radio"/> <i>Disable</i>	<input checked="" type="radio"/> <i>Normal</i> <input type="radio"/> <i>High</i>
5	<input type="radio"/> <i>Enable</i> <input checked="" type="radio"/> <i>Disable</i>	<input checked="" type="radio"/> <i>Normal</i> <input type="radio"/> <i>High</i>
6	<input type="radio"/> <i>Enable</i> <input checked="" type="radio"/> <i>Disable</i>	<input checked="" type="radio"/> <i>Normal</i> <input type="radio"/> <i>High</i>
7	<input type="radio"/> <i>Enable</i> <input checked="" type="radio"/> <i>Disable</i>	<input checked="" type="radio"/> <i>Normal</i> <input type="radio"/> <i>High</i>
8	<input type="radio"/> <i>Enable</i> <input checked="" type="radio"/> <i>Disable</i>	<input checked="" type="radio"/> <i>Normal</i> <input type="radio"/> <i>High</i>

Apply

[Configure 802.1pPriority](#) | [Configure TOS/DiffServ Priority](#)

Figure 27 — Configure Port QoS

4.2.5.6.2 Configure 802.1 Priority

The IEEE 802.1p extension of IEEE 802.1Q prioritizes traffic at the data-link/MAC layer through a 3-bit header field which was never articulated in the original VLAN standard. IEEE only suggests 802.1p definitions; it does not mandate them.

The upper panel of Figure 28 shows that 802.1p priority can be applied individually to ports. The lower panel shows that the switch provides 4 priority queues to which the 8 tags can be mapped in various schemes.

Configure 802.1p Priority:	
Port	Priority
1	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
2	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
3	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
4	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
5	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
6	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
7	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
8	<input checked="" type="radio"/> Enable <input type="radio"/> Disable

Apply

Map 802.1p Priority:	
Tag	Priority Queue
0	Low
1	Low
2	Normal
3	Normal
4	Medium
5	Medium
6	High
7	High

Apply

Configure Port QoS | Configure TOS/DiffServ Priority

Figure 28 — Configure 802.1p Priority

4.2.5.6.3 Configure TOS/DiffServ Priority

When configuring the TOS/DiffServ priority, a detailed screen of three tables appears. Each table — upper, middle and lower — is discussed and illustrated individually in turn. At the bottom of the screen there are two links to the other types of QoS.

Upper Table — Configure TOS/DiffServ Priority

The user must first decide which service to be used: TOS or DiffServ. This choice is made in the top-right portion of the *upper* table shown in the default screen of Figure 29.

The IP header contains an eight-bit field originally known as the **Type of Service (TOS)** field — but TOS priority had little acceptance. Subsequently, these eight bits were redefined to become the more popular **Differentiated Services (DiffServ)** field.

Regardless of the type of QoS chosen, each port (disabled by default) must be individually enabled to establish its QoS service.

After the choice of service has been made and the port behaviors determined for the upper table of Figure 29, these parameters must be registered by clicking the “Apply” button immediately below the table *before* editing the other tables or leaving this screen or refreshing this screen — otherwise, the upper-table settings will be lost.

Configure TOS/Diff Serv Priority:	
QoS Type	<input checked="" type="radio"/> TOS <input type="radio"/> Diff Serv
Port	Priority
1	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
2	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
3	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
4	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
5	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
6	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
7	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
8	<input type="radio"/> Enable <input checked="" type="radio"/> Disable

Figure 29 — Configure TOS/DiffServ Priority

Middle Table — Map TOS Precedence and Priority Queue

Although TOS priority is supported by few TCP/IP implementations, it is provided as a QoS option in managed switches from Contemporary Controls. One of the earliest methods of QoS for Internet Protocol, TOS uses the second octet (the TOS field) of the IP frame header — as described in RFC791 and RFC1349. The first three bits of this octet set the priority (*Precedence*). The next four bits (known as the *TOS* bits) define the tradeoffs among these four service objectives:

- minimize *monetary cost* (*M-Type Priority*) *
- maximize *reliability* (*R-Type Priority*)
- maximize *throughput* (*T-Type Priority*)
- minimize *delay* (*D-Type Priority*)

* This is sometimes referred to as “C-Type” priority or type of service.

The final bit of the TOS field was unused until DiffServ was defined.

Map TOS Precedence and Priority Queue:							
M-Type Priority		R-Type Priority		T-Type Priority		D-Type Priority	
0	Low	0	Low	0	Low	0	Low
1	Low	1	Low	1	Low	1	Low
2	Normal	2	Normal	2	Normal	2	Normal
3	Normal	3	Normal	3	Normal	3	Normal
4	Medium	4	Medium	4	Medium	4	Medium
5	Medium	5	Medium	5	Medium	5	Medium
6	High	6	High	6	High	6	High
7	High	7	High	7	High	7	High

Apply

Figure 30 — Map TOS Precedence and Priority Queue

Lower Table — Map DiffServ DSCP and Priority Queue

Sometimes called the “second generation of Internet QoS”, **Differentiated Services (DiffServ)** was first described by RFC2474 which redefines the TOS octet to allocate its first six bits as the **Differentiated Services Code Point (DSCP)**. The DSCP field selects the per-hop behavior (PHB) which defines how packets are queued at network nodes. RFC 2475 describes DiffServ methods for implementing scalable differentiated services on the Internet.

Map DiffServ DSCP and Priority Queue:							
DSCP	Priority	DSCP	Priority	DSCP	Priority	DSCP	Priority
0	Low	16	Normal	32	Medium	48	High
1	Low	17	Normal	33	Medium	49	High
2	Low	18	Normal	34	Medium	50	High
3	Low	19	Normal	35	Medium	51	High
4	Low	20	Normal	36	Medium	52	High
5	Low	21	Normal	37	Medium	53	High
6	Low	22	Normal	38	Medium	54	High
7	Low	23	Normal	39	Medium	55	High
8	Low	24	Normal	40	Medium	56	High
9	Low	25	Normal	41	Medium	57	High
10	Low	26	Normal	42	Medium	58	High
11	Low	27	Normal	43	Medium	59	High
12	Low	28	Normal	44	Medium	60	High
13	Low	29	Normal	45	Medium	61	High
14	Low	30	Normal	46	Medium	62	High
15	Low	31	Normal	47	Medium	63	High

Apply

Configure Port QoS | Configure 802.1p Priority

Figure 31 — Map DiffServ DSCP and Priority Queue

4.2.5.7 Configure Fault Relay

The switch has a relay output that can be used to signal the occurrence of one or more events. The screen of Figure 32 has two panels — one for relay **Settings** and one to **Monitor Fault Condition** — and shows that the relay can indicate the *loss* of a link or *presence* of a link on one or many ports.

4.2.5.7.1 Settings (upper panel)

Port Monitoring (link monitoring of specific ports) can be either *Enabled* or *Disabled*.

Relay State determines the behavior of the relay. By selecting “Make on Fault”, it will *close* its contacts once a fault is detected. “Break on Fault” will cause the relay to normally keep its contacts closed and *open* them upon detection of a fault. A fault is active when the condition of a monitored link or port matches its monitored state (see 4.2.5.7.2).

Relay Automation Time After Startup specifies a port-monitoring delay which allows the switch to stabilize for 1 to 999 seconds after startup. This is provided because, after startup, several seconds may be required for the switch to complete auto-negotiation of the data rate and duplex mode for each port. If the relay were not inhibited during this time, it could repeatedly activate without a true fault existing.

Relay Reset Method is “Automatic” by default, but can be set to “Manual”. In “Manual” mode, clicking on “Clear Relay” will reset the relay.

4.2.5.7.2 Monitor Fault Condition (lower panel)

As displayed in Figure 32, the user can monitor three conditions on a port-by-port basis: *Ignore* (the default) removes the port from link monitoring.

No Link (read only) reports a fault, (relay activated) if the link for the port has been lost.

Link Present (read only) reports a fault (relay activated) if a link has been detected on the port. This option is commonly used as a security feature to detect unauthorized connections to the switch.

Current Faults (read only) displays conditions when the screen first appears. While on screen, this report is static; it will only be updated if the refresh option is selected.

After relay activation is noted, the fault should be corrected. If *No Link* monitoring is in force, this will require restoration of a broken link or repair of the defective device to which the switch is connected. On the other hand, if *Link Present* monitoring is in use, removing the offending cable or end device will be required.

On *all* models the **Status LED** on the front panel of the switch glows solid green when operation is fault-free. For EI-WPMR-M models, this LED flashes if a fault occurs.

Settings:	
Port Monitoring	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
Relay State	<input checked="" type="radio"/> Break on Fault <input type="radio"/> Make on Fault
Relay Automation Time After Power Up	<input type="text" value="1"/> (1-999 seconds)
Relay Reset Method	<input checked="" type="radio"/> Automatic <input type="radio"/> Manual <input type="button" value="Clear Relay"/>

Monitor Fault Condition:								
Port	1	2	3	4	5	6	7	8
Ignore	<input checked="" type="radio"/>							
No Link	<input type="radio"/>							
Link Present	<input type="radio"/>							
Current Faults	OK							

Figure 32 — Configure Fault Relay

4.2.5.8 Configure Redundancy

Each managed switch from Contemporary Controls offers you a choice between the standard protocols known as Spanning Tree Protocol or Rapid Spanning Tree Protocol and the proprietary redundancy protocol known as RapidRing[®]. By default, the screen of Figure 33 displays RSTP and its basic parameter values for this switch.

To choose RapidRing as your redundancy scheme, select the *RapidRing*[®] button in the upper panel of the screen. Once this is done, a new screen appears with the options of Figure 36 — but before adjusting these values as described in Section 4.2.5.8.3.4, you should be familiar with all of the material discussed in Section 4.2.5.8.3.

4.2.5.8.1 Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol (RSTP) provides network path redundancy but without data loops that are prohibited. If more than one *active* path exists between two stations, a confused forwarding algorithm could transmit duplicate frames — one along each path.

RSTP constructs a tree of all RSTP-compliant switches in the network. To avoid loops, it forces each redundant path into an inactive state. If a segment is interrupted, an algorithm reconfigures the tree by quickly activating a normally unused link to substitute for the failed link.

By the exchange of messages, each switch in the tree collects information on all other switches. This information includes switch and port priorities, Media Access Control (MAC) addresses and path merit figures called “port costs”. This exchange results in the election of one switch to perform as the **root switch** (the *logical* center of the tree) and also defines how ports are to be used on all other switches. The **root port** on a switch will send traffic to the root switch along the most efficient path. If the root port is disrupted, a **backup port** is activated as the substitute. A **designated port** provides the best path for root-bound traffic from outlying switches. If the designated port is disrupted, an **alternate port** is activated as the substitute.

If all RSTP-compliant switches in the network are enabled with default settings, the switch with the lowest MAC address is elected the root switch. But due to network traffic and architecture issues, the elected switch might not be the best to serve as the root device. You can manually force the switch of your choice to serve as the root device by increasing its priority so that the root-election algorithm chooses it as the root.

In general, you should configure your RSTP network so that the paths with the greatest bandwidth are those which support traffic for the root switch. For conveying root traffic, a fiber optic link would be preferred over a copper link and a 100 Mbps link would serve better than one operating at 10 Mbps. Also, the tree should consist of only of devices that are RSTP-compliant — non-compliant switches and hubs, if used at all, should only occupy the periphery of the tree because they will not forward the special messages needed for the construction and maintenance of the tree.

4.2.5.8.2 Configure Spanning Tree Protocol

The second panel of Figure 33 (Configure STP) allows you to either *Disable* (the default) or *Enable* the protocol.

The third panel (Bridge Settings) is explained in Section 4.2.5.8.2.1.

The fourth panel (STP Port Settings) is explained in Section 4.2.5.8.2.2.

The fifth panel (STP Port Settings) is explained in Section 4.2.5.8.2.3.

Configure Redundancy:	
Redundancy State	<input checked="" type="radio"/> STP/RSTP <input type="radio"/> RapidRing™

Configure STP	
STP State	<input type="radio"/> Enable <input checked="" type="radio"/> Disable

Bridge Settings:	
Hello Time (secs)	2 (1-10)
Version	<input type="radio"/> STP <input checked="" type="radio"/> RSTP
Priority	32768
Max Age Time	20 (Min 6, Max 40)
Forward Delay	15 (Min 4, Max 30)
Root ID	Self

Apply

STP Port Settings:		
Port	Priority	Path Cost(1-65535)
1	128	19
2	128	19
3	128	19
4	128	19
5	128	19
6	128	19
7	128	19
8	128	19

Apply

STP Port Settings:		
Port	Admin Edge State	Non STP State
1	<input type="radio"/> On <input checked="" type="radio"/> Off	<input type="radio"/> On <input checked="" type="radio"/> Off
2	<input type="radio"/> On <input checked="" type="radio"/> Off	<input type="radio"/> On <input checked="" type="radio"/> Off
3	<input type="radio"/> On <input checked="" type="radio"/> Off	<input type="radio"/> On <input checked="" type="radio"/> Off
4	<input type="radio"/> On <input checked="" type="radio"/> Off	<input type="radio"/> On <input checked="" type="radio"/> Off
5	<input type="radio"/> On <input checked="" type="radio"/> Off	<input type="radio"/> On <input checked="" type="radio"/> Off
6	<input type="radio"/> On <input checked="" type="radio"/> Off	<input type="radio"/> On <input checked="" type="radio"/> Off
7	<input type="radio"/> On <input checked="" type="radio"/> Off	<input type="radio"/> On <input checked="" type="radio"/> Off
8	<input type="radio"/> On <input checked="" type="radio"/> Off	<input type="radio"/> On <input checked="" type="radio"/> Off

Apply

Figure 33 — Configure Spanning Tree Protocol

TD070446-0MA

4.2.5.8.2.1 Bridge Settings (Figure 33, Third Panel)

Hello Time

This is the interval at which the root device transmits a configuration message. The default value is 2 seconds, and can be set from 1–10 seconds. But if the result of the equation

$$(Max\ Age / 2) - 1$$

is less than 10, the maximum Hello Time will be the calculated value.

Version

The original 1990 link-management specification in IEEE802.1D, Clause 8, described path redundancy via the Spanning Tree Protocol (STP). In 1999 STP was superseded by the Rapid Spanning Tree Protocol (RSTP) of IEEE802.1D, Clause 17. The RSTP interoperates with STP, but if RSTP-compliant switches are used in the same network with legacy STP-compliant switches, rapid reconfiguration may not be possible. RSTP is the default selection for managed switches from Contemporary Controls.

Priority

You can adjust the *switch* priority in steps of 4096 (the default is 32768) as follows:

0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440.

Max Age

This is the maximum number of seconds a device waits to receive a configuration frame before attempting to reconfigure. The default value is 20 seconds, and it can be set from 6–40 seconds. However, actual minimum and maximum limits may be imposed by calculations based on the Hello Time and Forward Time as follows:

If the result of the equation

$$2 \times (Hello\ Time + 1)$$

is more than 6, the minimum Max Age will be the calculated value.

If the result of the equation

$$2 \times (Forward\ Time - 1)$$

is less than 40, the maximum Max Age will be the calculated value.

Forward Delay

This is the time a device will wait before changing states. Each device must receive topology information before it forwards frames and each port needs time to listen for any information that might force it to a discarding state. The default value is 15 seconds, and can be set from 4–30 seconds. But if the result of the equation

$$(Max\ Age / 2) + 1$$

is more than 4, the minimum Forward Delay will be the calculated value.

Root ID

This read-only value is the MAC Address of the root switch. However, if the switch under consideration is the root switch, it reports as “Self”.

NOTE: The value limits for **Hello Time**, **Max Age** and **Forward Delay** will only be imposed if the RSTP/STP field has been Enabled.

4.2.5.8.2.2 STP Port Settings (Figure 33, Fourth Panel)

The example shows an 8-port switch. The screen for a 16- or 24-port device would simply extend the panel to include all appropriate ports. In this panel you can adjust:

Priority (of an individual port)

With the **Priority** field, you can modify the priorities of individual ports to affect RSTP path choices in the local vicinity of the switch. A lower value means a higher priority. Priority settings differ from path costs which are *cumulative* in calculating a total path from periphery to root. The **Priority** value only acts *locally* so you can force RSTP to favor a certain path emerging from the switch in question when two paths from the switch are otherwise equal. This field's default value is 128 and its value can be toggled in 16 steps from 0–240 where each increment has a value of 16.

Path Cost

In determining the most efficient path for conveying messages between the periphery of the tree and its root, one of the factors RSTP relies on is “path costs”. A typical link operates at either 10 or 100 Mbps and a port sending traffic to that link is assigned a “cost” derived from the link data rate. The default for the **Path Cost** field is 19 (the nominal RSTP port cost for a 100 Mbps link). Nominal values of port path costs and the *suggested* ranges through which these values might vary in most networks are listed in **Table 1**. **Path Cost** can be set from 1–65535. When should you set a port path cost to a very high value? Although the associated link might operate at 10 Mbps, for example, the non-RSTP end of the link might have a very slow device such as a dial-up modem that could slow traffic drastically. In such a situation, you would likely want to raise this port path cost value to force RSTP to only use this path as a last resort.

Data Rate	Cost Range	Cost Value
4 Mbps	100 – 1000	250
10 Mbps	50 – 600	100
16 Mbps	40 – 400	62
100 Mbps	10 – 60	19

Table 1 — Port Path Costs

4.2.5.8.2.3 STP Port Settings (Figure 33, Fifth Panel)

Admin Edge State

This option (set to *Off* by default) can be set *On* if the attached device falls outside the RSTP tree. Such a device could be an operator work station or a server. In this case, the port affected by the **Admin Edge** would be at the *edge* of RSTP *administration*. End nodes cannot cause loops, so they can pass directly into the *Forwarding* state. Setting this value to *On* provides several benefits:

- gives faster RSTP algorithm solution (convergence)
- reduces flooding needed for rebuilding address tables during reconfiguration
- eliminates an RSTP reconfiguration if the port changes state
- improves other RSTP-related timeout issues

Non STP State

By default, each port has this parameter set to *Off* so the port can participate in RSTP. Setting this value to *On* will remove the port from RSTP, making the port immediately and constantly available for packet switching — regardless of any RSTP-related activities.

4.2.5.8.3 RapidRing®

4.2.5.8.3.1 Characteristics of RapidRing

RapidRing technology from Contemporary Controls provides high speed redundancy in Ethernet networks. It allows recovery in under 300 ms.

RapidRing is wired in a simple ring structure (see Figure 34) using the two RapidRing Ports (RRPs) on each switch. For EI-WPMR-M models, these are ports 7 and 8. Every link in the ring must connect by one RRP of one switch to the alternate RRP on the adjacent switch. Thus, port 7 ties to port 8. A properly constructed ring will *never* have a link between identical ports of adjacent switches. The ring ports can be wired with copper or fiber optic cable, depending on the model of switch.

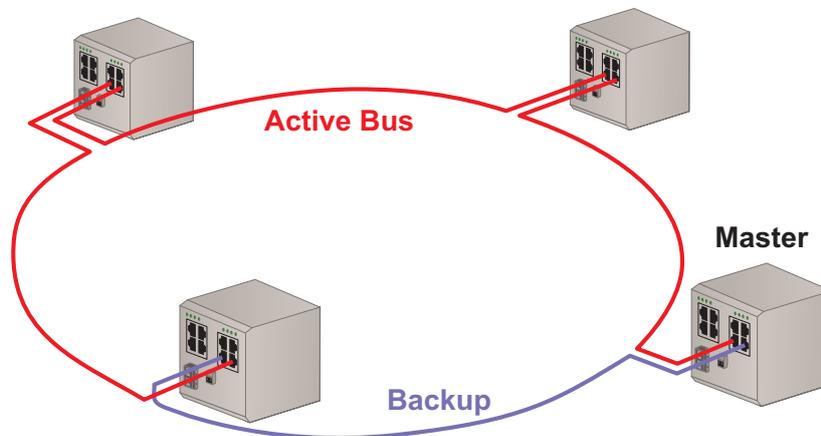


Figure 34 — RapidRing

As shown in Figure 34, one switch must be selected as the “master” that activates the backup link if a ring failure occurs. For EI-WPMR-M models, the backup link must tie to the master through **port 8** which remains inactive during normal network activity. If a cable failure is detected, the master will activate its backup port to maintain communications.

NOTE: Upon enabling RapidRing, the switch will *automatically reboot*.

Fault relays and status LEDs work independent of RapidRing. If RapidRing is enabled *and* the RRP's are being monitored, the following applies: When a break occurs, each switch losing connectivity will flash its Status LED and activate its Fault Relay (and *regardless of port monitoring*, it will transmit a link-down SNMP trap). The break is thus located to the link between the two fault-reporting switches. Once the cable is repaired by the user, the fault relays will disengage and the status LEDs will glow solid to indicate the ring network is properly connected.

NOTE: A flashing Status LED and Fault Relay activation do not necessarily indicate a *ring* failure. The user may also be monitoring *non-ring* ports — in which case the failure report might indicate a non-ring issue. However, ring failure in a properly constructed circuit will be reported by *two ring* switches — a non-ring issue will not be so reported.

4.2.5.8.3.2 RapidRing and Other Management Features

As a rule, when the RRP's are used for the RapidRing, they should not be involved with other management features. Specific issues are described below.

Port Parameters — If RapidRing is enabled, port options (Figure 11, upper panel) are not available for the RRP's. That is, the ring ports cannot be disabled nor can their configuration be changed.

Trunking — RapidRing and trunking *cannot* use the same ports. If a ring port exists as part of a trunk, the ring cannot be enabled. Either remove the ring port(s) from the trunk or disable the trunk that includes the ring port(s). Otherwise, the ring cannot be established. Once a ring is enabled, its ports cannot be added to a trunk.

Mirroring — RapidRing ports may be mirrored like any other port, but neither RRP should be designated as the Mirror Port (see Section 4.2.5.3).

VLANs — The RRP's must be included in the *same* VLAN or in none at all.

Multicast Filtering and Static Forwarding — RRP's may be used for these functions, but the backup RRP (of the ring master) is normally unavailable.

QoS — RRP's do support QoS.

4.2.5.8.3.3 Dual RapidRings

RapidRing will also support two interconnected rings (each having its own ID number) — allowing greater flexibility of wiring and network styles. Figure 35 shows two rings connected via two redundant network links — thus protecting against failure of a single cable within the trunk. (However, the two rings could be connected by just one link, if inter-ring redundancy were not required.) The two switches providing the inter-ring connection may be either master or slave and may use any available *non-RRP*.

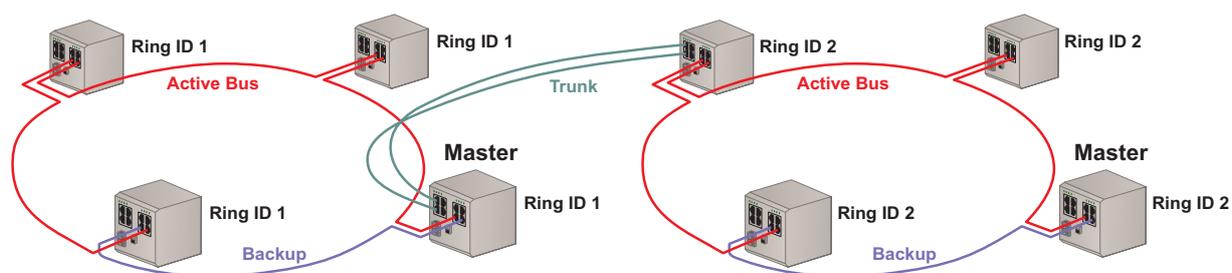


Figure 35 — Dual RapidRings

To setup the RapidRing, select **System Configuration**, then *Configure Redundancy* (described in Section 4.2.5.8) then select the *RapidRing*[®] option. Also, if the two rings are to have redundant lines between them, configure a trunk between the two switches that connect the rings together. (See Section 4.2.5.2 for configuring a trunk).

4.2.5.8.3.4 Configure RapidRing

Each switch in a ring must be configured for ring operation. The screen of Figure 36 displays the *default* screen settings. There are three options to consider:

Ring Status enables or disables RapidRing functionality. Every switch in a ring must have this option enabled. Otherwise, the presumed backup protection will not exist because a link failure might not be reported to the master.

Ring State sets the master/slave status of the switch being configured. Only one master is defined per ring. Otherwise, the ring will not be established, some signal paths may not exist and messages could be lost.

Ring ID must be 1 to 100. If only one ring is used, the default value of 1 is proper. If more rings are defined, each switch must be properly assigned to Ring ID values (from 1 to 100) — otherwise, messages might be lost. In summary, all switches in a ring must have matching Ring ID values.

Configure Redundancy:	
Redundancy State	<input type="radio"/> STP/RSTP <input checked="" type="radio"/> RapidRing™

Configure RapidRing™:	
Ring Status	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
Ring State	<input type="radio"/> Master <input checked="" type="radio"/> Slave
Ring ID	<input type="text" value="100"/> (1-100)
Network Status	Ring Not Available

Figure 36 — Configure RapidRing

In addition to the three options described above, a *Network Status* field reports the condition of the ring. The three following conditions are reported:

“Ring Not Available” is displayed when RapidRing is not enabled.

“Ring Incomplete” is displayed when RapidRing is enabled, but the master has invoked the backup link due to a primary ring failure.

“Ring Complete” is displayed when RapidRing is enabled and all links are intact.

4.2.5.9 Configure Rate Control

Figure 37 shows how ingress and egress port traffic can be controlled for the traffic rate and the traffic type. The default condition is for all messages to pass at 100 Mbps. At the top of the screen, select the traffic types for rate limiting (applies to both ingress and egress traffic). With rate control, bandwidth allocation can be finely controlled.

Configure Rate Control	
Configure Ingress Traffic Rate Control:	
Limited Type	<input checked="" type="checkbox"/> Broadcast <input checked="" type="checkbox"/> Multicast <input checked="" type="checkbox"/> Unicast <input checked="" type="checkbox"/> Destination Lookup Fail <input checked="" type="checkbox"/> MAC Control Frame
Port	Max Bit Rate (Mbps)
1	3
2	20
3	50
4	100
5	100
6	100
7	100
8	100

Apply

Configure Egress Traffic Rate Control:	
Port	Max Bit Rate (Mbps)
1	30
2	4
3	80
4	100
5	100
6	100
7	100
8	100

Apply

Figure 37 — Configure Rate Control

The selected Max Bit Rate is the maximum bandwidth level for the types of messages selected. Types not selected will be allowed to use 100% of the port's bandwidth. By selecting all the types, the full bandwidth of the port can be controlled. Selecting broadcast only creates a broadcast storm control with a selectable maximum bandwidth setting.

Rate control can be a useful feature for limiting communications from an unknown network. For example, when interconnecting the office and control networks.

4.2.5.10 Configure Port Security

Figure 38 illustrates the **Configure Port Security** default screen where each port has its security disabled. If security is enabled for a port, no further MAC addresses are learned for that port and future transmissions through the port will only succeed if the destination is listed in the address look-up table. Because static MAC addresses are not learned (aged out of the table), they are not affected by the applied security. A convenient link to **Add Static MAC Addresses** is provided at the bottom of the screen.

Configure Port Security:	
Port	Security
1	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
2	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
3	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
4	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
5	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
6	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
7	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
8	<input type="radio"/> Enable <input checked="" type="radio"/> Disable

Apply

[Add Static MAC Addresses](#)

Figure 38 — Configure Port Security

This is a useful feature when connecting to an unknown network — such as connecting the office network to the control network — and can be used to limit the office devices that can access the control network.

4.2.5.11 Configure IGMP Snooping

Traditionally, IP messages are either unicast or broadcast, but *multicasting* can deliver messages to a select group of devices on the network. IGMP (Internet Group Multicast Protocol) is a session-layer protocol for defining membership in a multicast group. IGMP Group Destination Addresses (GDA) range from 224.0.0.0 to 239.255.255.255.

Managed switches from Contemporary Controls support IGMPv2 and can provide the *querier function* in the event no router exists in the network. Our switches implement the *general query function* in which members of *all* multicast groups report.

With *IGMP Snooping*, managed switches can recognize packets that are for multicast groups and direct those packets to **only** the destination ports. These are ports which have received IGMP “join” messages from devices that seek membership in specific multicast groups. Figure 39 illustrates the **Configure IGMP Snooping** default screen. The available controls are listed below.

IGMP Snooping State enables or disables IGMP Snooping.

IGMP Querier Function allows the switch to initiate a query to discover multicast group membership. Only one querier is allowed on a network at a time. If a query is transmitted by a device with a lower IP address than this switch, this switch relinquishes the role of querier to that device. If another LAN device is preferred for the query function, then this option can be disabled.

IGMP Forwarding Map specifies the ports that will forward IGMP *join* or *leave* messages to other querying devices. Disabling ports that do not connect to querying devices can improve bandwidth.

IGMP Query Interval (secs) specifies how often querying occurs. This setting is meaningless if the IGMP Query Function is disabled.

Multicast Filtering Age Out (secs) specifies how long the switch waits before deleting an entry from its multicast group list. If no member of a group responds within this time, the group is deemed inactive and removed from the list. This time should exceed the Query Interval or else the entry may be deleted before another query occurs. The switch can accommodate up to 20 multicast groups.

Configure IGMP Snooping:	
IGMP Snooping State	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
IGMP Querier Function	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
IGMP Forwarding Map	<input checked="" type="checkbox"/> 1 <input checked="" type="checkbox"/> 2 <input checked="" type="checkbox"/> 3 <input checked="" type="checkbox"/> 4 <input checked="" type="checkbox"/> 5 <input checked="" type="checkbox"/> 6 <input checked="" type="checkbox"/> 7 <input checked="" type="checkbox"/> 8
IGMP Query Interval (secs)	<input type="text" value="125"/> (Default = 125, Max = 9999)
Multicast Filtering Age Out (secs)	<input type="text" value="125"/> (Default = 125, Max = 9999)

Apply

Figure 39 — Configure IGMP Snooping

4.2.6 SNMP Configuration.

In a large network, a network management system (NMS) based on the Simple Network Management Protocol (SNMP) is often used to keep track of operations. SNMP, created in 1988, is the standard protocol for managing network devices. Over TCP/IP, SNMP usually uses UDP ports 161 (SNMP) and 162 (SNMP-traps).

An SNMP implementation involves three areas of functionality: managed devices, SNMP agents and the NMS. SNMP agents reside in network devices where they use MIBs (information specific to the device) to interface the devices with the NMS — which then monitors and controls devices via these agents.

SNMP manages devices via a very small command set described in Section 5.1.6.

All EI-WPMR-M models are managed via SNMP. Figure 40 displays 3 panels where users can set **SNMP Configuration**, **Community** and **Trap Receivers**.

SNMP Configuration:	
System Name:	<input type="text" value="Managed Switch V4.16"/>
System Location:	<input type="text" value="(Type switch location here.)"/>
System Contact:	<input type="text" value="(Type contact's name here.)"/>

Community:		
Community String	Access	Status
<input type="text" value="public"/>	<input type="text" value="Read Only"/> ▼	<input type="text" value="Valid"/> ▼
<input type="text" value="private"/>	<input type="text" value="Read Write"/> ▼	<input type="text" value="Valid"/> ▼
<input type="text"/>	<input type="text" value="Read Only"/> ▼	<input type="text" value="Invalid"/> ▼
<input type="text"/>	<input type="text" value="Read Only"/> ▼	<input type="text" value="Invalid"/> ▼

Trap Receivers :	
IP Address	Status
<input type="text" value="10.0.0.138"/>	<input type="text" value="Valid"/> ▼
<input type="text" value="0.0.0.0"/>	<input type="text" value="Invalid"/> ▼
<input type="text" value="0.0.0.0"/>	<input type="text" value="Invalid"/> ▼
<input type="text" value="0.0.0.0"/>	<input type="text" value="Invalid"/> ▼

Figure 40 — SNMP Configuration

4.2.6.1 Configure System Information (Figure 40, upper panel)

In this panel, users can set three of the MIBs listed in Section 5.1.1.1:

System Name	(1.3.6.1.2.1.1.5.0)
System Location	(1.3.6.1.2.1.1.6.0)
System Contact	(1.3.6.1.2.1.1.4.0)

4.2.6.2 Configure SNMP Community (Figure 40, middle panel)

Managed devices are grouped into “communities” wherein every device has the same “community string” (aka “community name”) to be able to communicate via SNMP. This string assures authorized access to SNMP. Community strings can provide two types of access, read-only and read-write. Read-only access allows only *get* and *get-next* commands. Read-write access allows *get*, *get-next* and *set* commands.

Up to four SNMP community names can be defined — each specifying either read-only access or read/write access. Before an access can be used, it must be set to *Valid*. Each community has three parameters to be configured:

Community String is the name (up to 10 characters) created by the user. It functions as a password to be used by any SNMP management software which accesses the switch.

Access is chosen by the user to be either “Read Only” or “Read Write”.

Status is chosen to be either “Valid” (string enabled) or “Invalid” (string disabled).

Figure 40 shows that two community strings are, by default, pre-defined and valid: “public” (set for read-only access) and “private” (set for read-write access). (Several commonly-available SNMP manager applications use *public* and *private* as default strings.)

4.2.6.3 Configure SNMP Trap Receivers (Figure 40, lower panel)

An SNMP Trap is a message that is transmitted when a trap event occurs. The menu in Figure 40 allows up to four trap receiver IP addresses to be defined and each must be marked *Valid* for it to be used. Each valid trap receiver will receive a trap message upon a trap event occurring. The switch supports traps for:

- link-up*
- link-down*
- authentication failure*
- cold start*
- warm start*

The example of Figure 40 defines an **IP Address** for only one trap receiver — which will function because its status parameter has been set to *Valid*.

NOTE: For more information on SNMP support within the switch, see the Appendix.

4.2.7 Performance Monitoring

Switch performance can be monitored via SNMP, web page and console menus. The **Performance Monitoring** options of Figure 41 are discussed below.



Figure 41 — Performance Monitoring

4.2.7.1 Browse Address Table

The **Browse Address Table** appears in Figure 42. The entire table (up to 512 entries) can be displayed or a particular MAC address can be located. In the *Sort By* field, select the type of search (*Sequence* or *MAC Address*) then click the *Display* button.

If many screens are needed, scroll bars allow viewing all located MAC addresses by the *MAC Address* and its associated port. Under *Status*, "VALID" means the entry has not been *aged* from the table, "INVALID" means it has been aged or deleted by the user and may be replaced when a new entry is added, "STATIC" means it is controlled by the management CPU and automatic learning and aging of the entry will not occur.

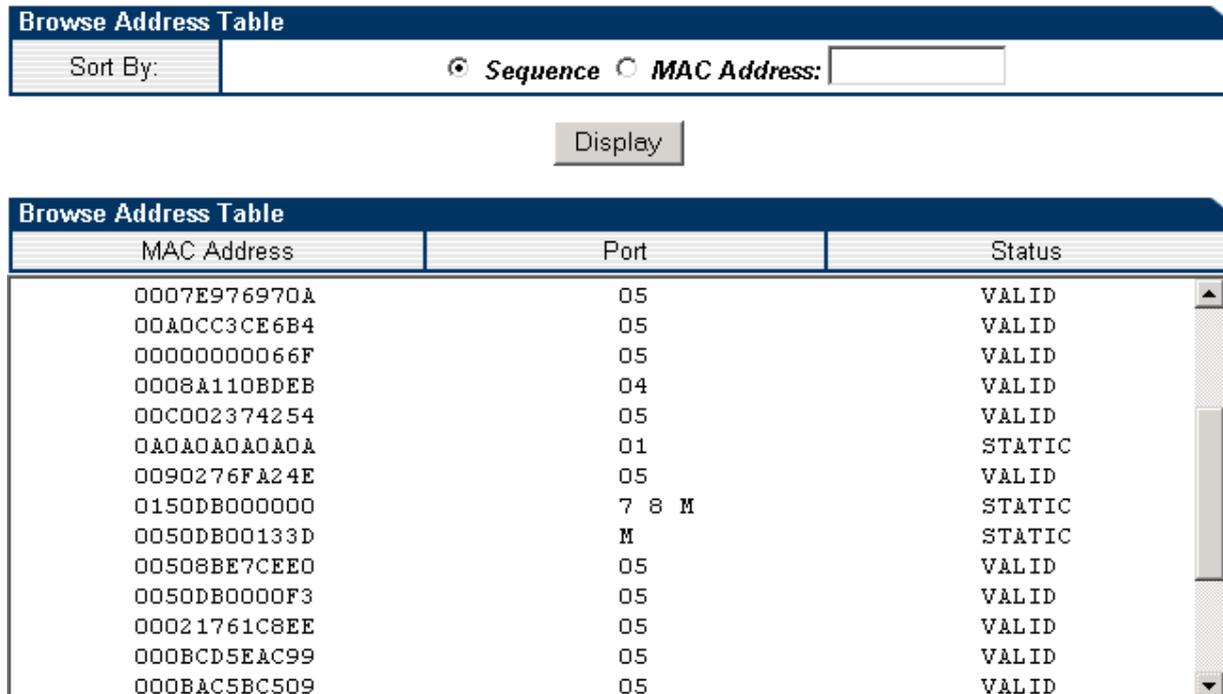


Figure 42 — Browse Address Table

4.2.7.2 Traps Log

The **Traps Log** of Figure 43 reports SNMP traps sent since the last Cold Start and the number of seconds that have passed since the Cold Start occurred (at time zero).

Traps Log :		
Secs Since Startup	Type	Port
0	Cold Start	
56711	Link Down	4
56712	Link Up	4
60886	Link Down	5
60892	Link Up	5

Figure 43 — Traps Log

4.2.7.3 Monitor STP Port Status

The screens of Figure 44 (where STP is enabled) or that of Figure 45 (where STP is disabled) report the **STP State** and **Link Status** for each port. These show an 8-port switch. A 16- or 24-port device screen would extend the panel to display the proper number of ports. The significance of the reported information is discussed on the next page.

Monitor STP Port Status:		
Port	STP State	Link Status
1	Discarding	Down
2	Discarding	Down
3	Discarding	Down
4	Forwarding	Up
5	Discarding	Down
6	Discarding	Down
7	Discarding	Down
8	Forwarding	Up

Figure 44 — Monitor STP Port Status (STP Enabled)

Monitor STP Port Status:		
Port	STP State	Link Status
1	Forwarding	Down
2	Forwarding	Down
3	Forwarding	Down
4	Forwarding	Up
5	Forwarding	Down
6	Forwarding	Down
7	Forwarding	Down
8	Forwarding	Up

Figure 45 — Monitor STP Port Status (STP Disabled)

STP State may report any of three states:

Forwarding indicates the RSTP port is actively participating in the tree — that is, it is not currently a backup or alternate port. The port forwards frames, and continues to learn new addresses. If RSTP is **disabled** as in Figure 53, each port is reported to be in the *Forwarding* state (whether the port is *Up* or *Down*) — but this has no meaning since the tree is disabled.

Learning is the brief port state (as it begins to learn addresses) before its forwarding delay expires. Actually, this state is rarely reported on screen, but could be if the user happens to refresh the monitor screen while a port is transitioning from Discarding to Forwarding. This Learning state will only occur as the tree is being restructured — as must be the case in the sample of Figure 52.

Discarding is reported if the port is not active in the tree. The port receives STP frames, but it does not forward frames. If the tree is stable and port states are consistent throughout the network, every root port and designated port will quickly transition through its Learning state to the Forwarding state. At the same time, all alternate and backup ports will stay in the Discarding state because they are not active in the tree.

Notes on Figures 52 and 53: A link may be reported as Up even if its associated port is not active in the tree. A look at Figure 52 reveals that although its link is Up, port 8 is Discarding — therefore port 8 is either a backup or an alternate port. Port 3 is listed as Up and Forwarding — thus, it is active in the tree and working properly. Port 2 is Up, but its state is momentarily Learning — so port 2 will soon become Forwarding.

Only RSTP port states (which superseded STP states) are shown in the monitor screen. For those wish to compare the port states of RSTP versus STP, refer to Table 2.

Status	STP State	RSTP State	Active In Tree?
Enabled	Blocking	Discarding	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	Yes
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

Table 2 — STP & RSTP Port States Compared

The **Link Status** column of the Monitor STP Port Status screen simply reports if a port is *Up* or *Down*. It does this even if the RSTP/STP function is disabled.

5 Appendix

5.1 SNMP

All EI-WPMR-M models provide an SNMP interface for management of the device. The switches currently support:

- RFC 1157 — SNMP protocol
- RFC 1213 — MIB-2
- RFC 1215 — Traps for SNMP
- RFC 1493 — Bridge MIB
- RFC 1573 — MIB-2 Extension (IF-MIB)
- RFC 1643 — Ethernet-like Interface MIB

The following MIBs are supported.

5.1.1 Managed Objects for TCP/IP Based Internet (MIB-II) — From RFC 1213

5.1.1.1 'System' group 1.3.6.1.2.1.1

oid = "1.3.6.1.2.1.1.1.0"

sysDescr: A textual description of the switch
Access: read-only

oid = "1.3.6.1.2.1.1.2.0"

sysObjectID: The vendor's authoritative identification
Access: read-only

oid = "1.3.6.1.2.1.1.3.0"

sysUpTime: The time since the last re-initialization
Access: read-only

oid = "1.3.6.1.2.1.1.4.0"

sysContact: The identification of the contact person
Access: read-write

oid = "1.3.6.1.2.1.1.5.0"

sysName: An administratively assigned name
Access: read-write

oid = "1.3.6.1.2.1.1.6.0"

sysLocation: The physical location of this node
Access: read-write

oid = "1.3.6.1.2.1.1.7.0"

sysServices: Indicates the set of services that this switch primarily offers
Access: read-only

5.1.1.2 'Interfaces' group 1.3.6.1.2.1.2

oid = "1.3.6.1.2.1.2.1.0"

ifNumber: The number of network interfaces (regardless of their current state) present on this system

Access: read-only

5.1.1.2.1 The Interfaces Table — 'ifTable' 1.3.6.1.2.1.2.2

oid = "1.3.6.1.2.1.2.2.1.1.ifIndex"

ifIndex: A unique value, greater than zero, for each interface

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.2.ifIndex"

ifDescr: Interface string contains information about the interface

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.3.ifIndex"

ifType: Interface type = 6 if Ethernet

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.4.ifIndex"

ifMtu: The size of the largest datagram that can be sent/received

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.5.ifIndex"

ifSpeed: Interface speed = 100000000 for 100Base-TX

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.6.ifIndex"

ifPhysAddress: Ethernet (MAC) address (only used for designated management port on a switch)

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.7.ifIndex"

ifAdminStatus: The desired state of the interface

1 = up (ready to pass packets)

2 = down

3 = testing

Access: read-write

oid = "1.3.6.1.2.1.2.2.1.8.ifIndex"

ifOperStatus: The current operational state of the interface

1 = up (ready to pass packets)

2 = down

3 = testing

4 = unknown

5 = dormant

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.9.ifIndex"
ifLastChange: The value of sysUpTime at the time the interface entered its current operational state
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.10.ifIndex"
ifInOctets: the total number of octets received on the interface, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.11.ifIndex"
ifInUcastPkts: The number of unicast packets received, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.12.ifIndex"
ifInNUcastPkts: The number of non-unicast packets received, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.13.ifIndex"
ifInDiscards: The number of inbound packets discarded with no error detected, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.14.ifIndex"
ifInErrors: The number of inbound packets with errors, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.15.ifIndex"
ifInUnknownProtos: The number of packets received via the interface which were discarded because of unknown or unsupported protocol. Returns 0.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.16.ifIndex"
ifOutOctets: The total number of packets transmitted, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.17.ifIndex"
ifOutUcastPkts: The number of packets transmitted to a unicast address, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.18.ifIndex"
ifOutNUcastPkts: The number of packets transmitted to a non-unicast address, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.19.ifIndex"
ifOutDiscards: The number of outbound packets discarded with no error detected, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.20.ifIndex"
ifOutErrors: The number of outbound packets with errors, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.22.ifIndex"
ifSpecific: A reference to MIB definition specific to the particular media being used to realize the interface
Access: read-only

5.1.1.3 'IP' group 1.3.6.1.2.1.4

oid = "1.3.6.1.2.1.4.1.0"

ipForwarding: The indication of whether this switch is acting as an IP gateway in respect to the forwarding of datagrams received by, but not addressed to, this switch

1 = forwarding

2 = not forwarding

Access: read-write

oid = "1.3.6.1.2.1.4.2.0"

ipDefault: The default value inserted into the Time-To-Live field of the IP header of datagrams originated at this switch, whenever a TTL value is not supplied by the transport layer protocol

Access: read-only

oid = "1.3.6.1.2.1.4.3.0"

ipInReceives: The total number of input datagrams received from interfaces, including those received in error

Access: read-only

oid = "1.3.6.1.2.1.4.9.0"

ipInDelivers: The total number of input datagrams successfully delivered to IP user-protocols (including ICMP)

Access: read-only

oid = "1.3.6.1.2.1.4.10.0"

ipOutRequests: The total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission

Access: read-only

oid = "1.3.6.1.2.1.4.15.0"

ipReasmOKs: The number of IP datagrams successfully re-assembled

Access: read-only

oid = "1.3.6.1.2.1.4.17.0"

ipFragOKs: The number of IP datagrams that have been successfully fragmented at this switch

Access: read-only

5.1.1.3.1 The IP Address Table — 'ipAddrTable' 1.3.6.1.2.1.4.20

When IP address is used as input, its format should be 4 decimal fields.

oid = "1.3.6.1.2.1.4.20.1.1.<ipAdEntAddr>"

ipAdEntAddr: The IP address

Access: read-only

oid = "1.3.6.1.2.1.4.20.1.2.< ipAdEntAddr>"

ipAdEntIfIndex: Physical port number associated with this particular subnet by IP address

Access: read-only

oid = "1.3.6.1.2.1.4.20.1.3.< ipAdEntAddr >"

ipAdEntNetMask: Subnet mask associated with this particular subnet by IP address

Access: read-only

oid = "1.3.6.1.2.1.4.20.1.4.< ipAdEntAddr >"
ipAdEntBcastAddr: The value of the least-significant bit in the IP broadcast address = 1 for Internet standard all-ones broadcast address
Access: read-only

oid = "1.3.6.1.2.1.4.20.1.5.< ipAdEntAddr >"
ipAdEntReasmMaxSize: The size of the largest IP datagram which this switch can re-assemble from incoming IP fragments
Access: read-only

5.1.1.4 'ICMP' group 1.3.6.1.2.1.5

oid = "1.3.6.1.2.1.5.1.0"
icmplnMsgs: The total number of ICMP messages which the switch received
Access: read-only

oid = "1.3.6.1.2.1.5.2.0"
icmplnErrors: The number of ICMP messages which the switch received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.)
Access: read-only

oid = "1.3.6.1.2.1.5.8.0"
icmplnEchos: The number of ICMP Echo (request) messages received
Access: read-only

oid = "1.3.6.1.2.1.5.9.0"
icmplnEchoReps: The number of ICMP Echo Reply messages received
Access: read-only

oid = "1.3.6.1.2.1.5.14.0"
icmplnOutMsgs: The total number of ICMP messages which this switch attempted to send
Access: read-only

oid = "1.3.6.1.2.1.5.15.0"
icmplnOutErrors: The number of ICMP messages which this switch did not send due to problems discovered within ICMP
Access: read-only

5.1.1.5 'TCP' group 1.3.6.1.2.1.6

oid = "1.3.6.1.2.1.6.10.0"
tcpInSegs: The total number of segments received, including those received in error
Access: read-only

oid = "1.3.6.1.2.1.6.11.0"
tcpOutSegs: The total number of segments sent, including those on current connections but excluding those containing only retransmitted octets
Access: read-only

oid = "1.3.6.1.2.1.6.12.0"
tcpRetransSegs: The total number of segments retransmitted
Access: read-only

5.1.1.6 'UDP' group 1.3.6.1.2.1.7

5.1.1.6.1 The UDP Listener Table — 'udpTable' 1.3.6.1.2.1.7.5

When an IP address is used as input, its format should be 4 decimal fields.

oid = "1.3.6.1.2.1.7.5.1.1.< udpLocalAddress >.< udpLocalPort >"

udpLocalAddress: The local IP address for this UDP listener.

Access: read-only

oid = "1.3.6.1.2.1.7.5.1.2.<Local IP address as 4 decimal fields>.<LocalPort>"

udpLocalPort: The local port number for this UDP listener.

Access: read-only

5.1.1.7 'Transmission' group 1.3.6.1.2.1.10

Based on the transmission media underlying each interface on a system, the corresponding portion of the Transmission group is mandatory for that system.

In this switch all interfaces are Ethernet-based — hence, the Ethernet-like interface is used.

5.1.1.8 'SNMP' group 1.3.6.1.2.1.11

oid = "1.3.6.1.2.1.11.1.0"

snmplnPkts: The total number of Messages delivered to the SNMP switch from the transport service

Access: read-only

oid = "1.3.6.1.2.1.11.2.0"

snmpOutPkts: The total number of SNMP Messages which were passed from the SNMP protocol switch to the transport service

Access: read-only

oid = "1.3.6.1.2.1.11.3.0"

snmplnBadVersions: The total number of SNMP Messages which were delivered to the SNMP protocol switch and were for an unsupported SNMP version

Access: read-only

oid = "1.3.6.1.2.1.11.4.0"

snmplnBadCommunityNames: The total number of SNMP Messages delivered to the SNMP protocol switch that used an SNMP community name not known to said switch

Access: read-only

oid = "1.3.6.1.2.1.11.5.0"

snmplnBadCommunityUses: The total number of SNMP messages delivered to the SNMP protocol switch and representing an SNMP operation that was not allowed by the SNMP community named in the message

Access: read-only

oid = "1.3.6.1.2.1.11.6.0"

snmplnASNParseErrs: The total number of ASN.1 or BER errors encountered by the SNMP protocol switch when decoding received SNMP Messages

Access: read-only

oid = "1.3.6.1.2.1.11.8.0"

snmplnTooBig: The total number of SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'tooBig'

Access: read-only

oid = "1.3.6.1.2.1.11.9.0"
snmplnNoSuchNames: The total number of SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'noSuchName'
Access: read-only

oid = "1.3.6.1.2.1.11.10.0"
snmplnBadValues: The total number of SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'badValue'
Access: read-only

oid = "1.3.6.1.2.1.11.11.0"
snmplnReadOnlys: The total number valid SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'readOnly'
Access: read-only

oid = "1.3.6.1.2.1.11.12.0"
snmplnGenErrs: The total number of SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'genErr'
Access: read-only

oid = "1.3.6.1.2.1.11.13.0"
snmplnTotalReqVars: The total number of MIB objects which have been retrieved successfully by the SNMP protocol switch as the result of receiving valid SNMP Get-Request and Get-Next PDUs
Access: read-only

oid = "1.3.6.1.2.1.11.14.0"
snmplnTotalSetVars: The total number of MIB objects that have been altered successfully by the SNMP protocol switch as the result of receiving valid SNMP Set-Request PDUs
Access: read-only

oid = "1.3.6.1.2.1.11.15.0"
snmplnGetRequests: The total number of SNMP Get-Request PDUs that have been accepted and processed by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.16.0"
snmplnGetNexts: The total number of SNMP Get-Next PDUs that have been accepted and processed by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.17.0"
snmplnSetRequests: The total number of SNMP Set-Request PDUs that have been accepted and processed by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.18.0"
snmplnGetResponses: The total number of SNMP Get-Response PDUs that have been accepted and processed by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.19.0"
snmplnTraps: The total number of SNMP Trap PDUs that have been accepted and processed by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.20.0"
snmpOutTooBig: The total number of SNMP PDUs that were generated by the SNMP protocol switch and for which the value of the error-status field is 'tooBig'
Access: read-only

oid = "1.3.6.1.2.1.11.21.0"
snmpOutNoSuchNames: The total number of SNMP PDUs that were generated by the SNMP protocol switch and for which the value of the error-status is 'noSuchName'
Access: read-only

oid = "1.3.6.1.2.1.11.22.0"
snmpOutBadValues: The total number of SNMP PDUs that were generated by the SNMP protocol switch and for which the value of the error-status field is 'badValue'
Access: read-only

oid = "1.3.6.1.2.1.11.24.0"
snmpOutGenErrs: The total number of SNMP PDUs that were generated by the SNMP protocol switch and for which the value of the error-status field is 'genErr'
Access: read-only

oid = "1.3.6.1.2.1.11.25.0"
snmpOutGetRequests: The total number of SNMP Get-Request PDUs that have been generated by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.26.0"
snmpOutGetNexts: The total number of SNMP Get-Next PDUs that have been generated by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.27.0"
snmpOutSetRequests: The total number of SNMP Set-Request PDUs that have been generated by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.28.0"
snmpOutGetResponses: The total number of SNMP Get-Response PDUs that have been generated by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.29.0"
snmpOutTraps: The total number of SNMP Trap PDUs that have been generated by the SNMP protocol switch
Access: read-only

oid = "1.3.6.1.2.1.11.30.0"
snmpEnableAuthenTraps: Indicates whether the SNMP agent process is permitted to generate authentication-failure traps
1 = enabled
2 = disabled
Access: read-write

5.1.2 Managed Objects for Bridges — From RFC 1493

Bridge MIB — 'dot1dBridge' 1.3.6.1.2.1.17

5.1.2.1 'dot1dBase' group 1.3.6.1.2.1.17.1

oid = "1.3.6.1.2.1.17.1.1.0"

dot1dBaseBridgeAddress: The MAC address used by this bridge
Access: read-only

oid = "1.3.6.1.2.1.17.1.2.0"

dot1dBaseNumPorts: The number of ports controlled by this bridging switch
Access: read-only

oid = "1.3.6.1.2.1.17.1.3.0"

dot1dBaseType: Indicates what type of bridging this bridge can perform
1 = unknown
2 = transparent-only
3 = sourceroute-only
Access: read-only

5.1.2.1.1 'dot1dBasePortTable' 1.3.6.1.2.1.17.1.4

oid = "1.3.6.1.2.1.17.1.4.1.1.port"

dot1dBasePort: The port number of the port for which this entry contains bridge management information
Access: read-only

oid = "1.3.6.1.2.1.17.1.4.1.2.port"

dot1dBasePortIfIndex: The value of instance of ifIndex object defined in Interface group of MIB-2 for the interface corresponding to this port
Access: read-only

oid = "1.3.6.1.2.1.17.1.4.1.3.port"

dot1dBasePortCircuit: For a port which has the same value of dot1dBasePortIfIndex as another port on the same bridge, this object contains the name of an object instance unique to this port. For a port which has a unique value of dot1dBasePortIfIndex, this object can have the value {0}.
Access: read-only

oid = "1.3.6.1.2.1.17.1.4.1.4.port"

dot1dBasePortDelayExceededDiscards: The number of frames discarded by this port due to excessive transmit delay through the bridge. Returns 0.
Access: read-only

oid = "1.3.6.1.2.1.17.1.4.1.5.port"

dot1dBasePortMtuExceededDiscards: The number of frames discarded by this port due to an excessive size
Access: read-only

5.1.2.1.2 'dot1dStp' group 1.3.6.1.2.1.17.2

oid = "1.3.6.1.2.1.17.2.1.0"

dot1dStpProtocolSpecification: Spanning Tree Protocol (STP) being run

1 = unknown(1)

2 = decLb100

3 = ieee8021d

Access: read-write

oid = "1.3.6.1.2.1.17.2.2.0"

dot1dStpPriority: Value of the write-able portion (first two octets) of the Bridge ID

Access: read-write

oid = "1.3.6.1.2.1.17.2.3.0"

dot1dStpTimeSinceTopologyChange: Time (in hundredths of a second) since the last time a topology change was detected by the bridge entity

Access: read-write

oid = "1.3.6.1.2.1.17.2.4.0"

dot1dStpTopChanges: Total number of topology changes detected by this bridge since the management entity was last reset or initialized

Access: read-write

oid = "1.3.6.1.2.1.17.2.5.0"

dot1dStpDesignatedRoot: Bridge identifier of the spanning tree root determined by the STP as executed by this node

Access: read-write

oid = "1.3.6.1.2.1.17.2.6.0"

dot1dStpRootCost: Path cost from this bridge to the root bridge

Access: read-write

oid = "1.3.6.1.2.1.17.2.7.0"

dot1dStpRootPort: Port number of the port with the lowest cost path from this bridge to the root bridge

Access: read-write

oid = "1.3.6.1.2.1.17.2.8.0"

dot1dStpMaxAge: Maximum age (hundredths of a second) of STP information learned from the network on any port before it is discarded

Access: read-write

oid = "1.3.6.1.2.1.17.2.9.0"

dot1dStpHelloTime: Interval (hundredths of a second) between transmissions of Configuration bridge PDUs by this node on any port when it is the root bridge or trying to become the root bridge

Access: read-write

oid = "1.3.6.1.2.1.17.2.10.0"

dot1dStpHoldTime: Interval (hundredths of a second) during which no more than two Configuration bridge PDUs shall be transmitted by this node

Access: read-write

oid = "1.3.6.1.2.1.17.2.11.0"

dot1dStpForwardDelay: Duration (in hundredths of a second) of a port's Listening and Learning states preceding the Forwarding state. This value is also used to age all Forwarding Database dynamic entries when a topology change has been detected and is underway.

Access: read-write

oid = "1.3.6.1.2.1.17.2.12.0"

dot1dStpBridgeMaxAge: MaxAge for all bridges when this bridge is the root

Access: read-write

oid = "1.3.6.1.2.1.17.2.13.0"

dot1dStpBridgeHelloTime: HelloTime for all bridges when this bridge is the root

Access: read-write

oid = "1.3.6.1.2.1.17.2.14.0"

dot1dStpBridgeForwardDelay: ForwardDelay for all bridges if this bridge is the root

Access: read-write

5.1.2.1.2.1 'dot1dStpPortTable' 1.3.6.1.2.1.17.2.15

oid = "1.3.6.1.2.1.17.2.15.1.1.port"

dot1dStpPort: Port number of the port for which this entry contains STP management information

Access: read-only

oid = "1.3.6.1.2.1.17.2.15.1.2.port"

dot1dStpPortPriority: Value of the priority field contained in the first octet (in network byte order) of the 2 octet Port ID

Access: read-only

oid = "1.3.6.1.2.1.17.2.15.1.3.port"

dot1dStpPortState: Port's current STP state

1 = disabled

2 = blocking

3 = listening

4 = learning

5 = forwarding

6 = broken

Access: read-only

oid = "1.3.6.1.2.1.17.2.15.1.4.port"

dot1dStpPortEnable: Port status

1 = enabled

2 = disabled

Access: read-only

oid = "1.3.6.1.2.1.17.2.15.1.5.port"

dot1dStpPortPathCost: Contribution of this port to the path cost of any paths toward the root bridge that include this port

Access: read-only

oid = "1.3.6.1.2.1.17.2.15.1.6.port"

dot1dStpPortDesignatedRoot: Unique Bridge Identifier of the Root Bridge recorded in the Configuration BPDUs transmitted by the Designated Bridge for the segment to which the port is attached

Access: read-only

oid = "1.3.6.1.2.1.17.2.15.1.7.port"

dot1dStpPortDesignatedCost: Path cost of the Designated Port of the segment connected to this port

Access: read-only

oid = "1.3.6.1.2.1.17.2.15.1.8.port"

dot1dStpPortDesignatedBridge: Bridge Identifier of the Designated Bridge for this port's segment

Access: read-only

oid = "1.3.6.1.2.1.17.2.15.1.9.port"

dot1dStpPortDesignatedPort: Port Identifier of the port on the Designated Bridge for this port's segment

Access: read-only

oid = "1.3.6.1.2.1.17.2.15.1.10.port"

dot1dStpPortForwardTransitions: Number of times this port has transitioned from the Learning state to the Forwarding state

Access: read-only

5.1.2.2 'dot1dTp' group 1.3.6.1.2.1.17.4

oid = "1.3.6.1.2.1.17.4.2.0"

dot1dTpAgingTime: The timeout period in seconds for aging out dynamically learned forwarding information

Access: read-write

5.1.2.3 'dot1dTpFdbTable' 1.3.6.1.2.1.17.4.3

oid = "1.3.6.1.2.1.17.4.3.1.1.<MAC address as 6 decimal fields>"

dot1dTpFdbAddress: A unicast MAC address for which the bridge has forwarding and/or filtering information

Access: read-only

oid = "1.3.6.1.2.1.17.4.3.1.2.<MAC address as 6 decimal fields>"

dot1dTpFdbPort: The port number where this MAC has been 'learned' and stored in the switch lookup table

Access: read-only

oid = "1.3.6.1.2.1.17.4.3.1.3. <MAC address as 6 decimal fields>"

dot1dTpFdbStatus: The status of this entry.

- 1 = other
- 2 = invalid
- 3 = learned
- 4 = self
- 5 = mgmt

Access: read-only

5.1.2.4 'dot1dTpPortTable' 1.3.6.1.2.1.17.4.4

oid = "1.3.6.1.2.1.17.4.4.1.1.port"

dot1dTpPort: The port number of the port for which this entry contains transparent bridging management information

Access: read-only

oid = "1.3.6.1.2.1.17.4.4.1.2.port"

dot1dTpPortMaxInfo: The maximum size of the INFO (non-MAC) field that this port will receive or transmit

Access: read-only

oid = "1.3.6.1.2.1.17.4.4.1.3.port"

dot1dTpPortInFrames: The number of frames received by this port

Access: read-only

oid = "1.3.6.1.2.1.17.4.4.1.4.port"

dot1dTpPortOutFrames: The number of frames transmitted by this port

Access: read-only

oid = "1.3.6.1.2.1.17.4.4.1.5.port"

dot1dTpPortInDiscards: The number of valid frames received which were discarded by the forwarding process

Access: read-only

5.1.3 Managed Objects for Ethernet-like Interface Types — From RFC 1643

Ethernet-like Interface MIB — ‘dot3’ 1.3.6.1.2.1.10.7

5.1.3.1 Ethernet-like Statistics Group — ‘dot3StatsTable’ 1.3.6.1.2.1.10.7.2

oid = “1.3.6.1.2.1.10.7.2.1.1.dot3StatsIndex”

dot3StatsIndex: An index that identifies an interface, same value as ifIndex
Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.2.dot3StatsIndex”

dot3StatsAlignmentErrors: The number of frames received on the interface that are not an integral number of octets in length and do not pass the FCS check. This count is incremented when the alignmentError status is returned by the MAC service to the LLC.
Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.3.dot3StatsIndex”

dot3StatsFCSErrors: The number of frames received on the interface that are not an integral number of octets in length and do not pass the FCS check. This count is incremented when the frameCheckError status is returned by the MAC service to the LLC.
Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.4.dot3StatsIndex”

dot3StatsSingleCollisionFrames: The number of successfully transmitted frame on the interface for which transmission is inhibited by one collision
Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.5.dot3StatsIndex”

dot3StatsMultipleCollisionFrames: The number of successfully transmitted frame on the interface for which transmission is inhibited by more than one collision
Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.6.dot3StatsIndex”

dot3StatsSQETestErrors: The number of times that the SQE TEST ERROR message is generated by the PLS sublayer for this interface. Returns 0.
Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.7.dot3StatsIndex”

dot3StatsDeferredTransmissions: The number of frames for which the first transmission attempt on the interface is delayed because the medium is busy
Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.8.dot3StatsIndex”

dot3StatsLateCollisions: The number of times that a collisions is detected on the interface later than 512 bit-times into the transmission of a packet
Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.9.dot3StatsIndex”

dot3StatsExcessiveCollisions: The number of frames for which transmission on the interface fails due to excessive collisions
Access: read-only

oid = "1.3.6.1.2.1.10.7.2.1.10.dot3StatsIndex"

dot3StatsInternalMacTransmitErrors: The number of frames for which transmission on the interface fails due to an internal MAC sublayer transmit error

Access: read-only

oid = "1.3.6.1.2.1.10.7.2.1.11.dot3StatsIndex"

dot3StatsCarrierSenseErrors: The number of frames that the carrier sense condition was lost or never asserted in attempting to transmit a frame on this interface. Returns 0.

Access: read-only

oid = "1.3.6.1.2.1.10.7.2.1.13.dot3StatsIndex"

dot3StatsFrameTooLong: The number of frames received on a particular interface that exceed the maximum permitted frame size

Access: read-only

oid = "1.3.6.1.2.1.10.7.2.1.16.dot3StatsIndex"

dot3StatsInternalMacReceiveErrors: The number of frames for which reception on the interface fails due to an internal MAC sublayer transmit error

Access: read-only

oid = "1.3.6.1.2.1.10.7.2.1.17.dot3StatsIndex"

dot3StatsEthernetChipSet: This object contains an OBJECT IDENTIFIER that identifies the chipset used to realize the interface.

Access: read-only

5.1.4 Evolution of the Interface Group of MIB-II — From RFC 1573

MIBs for generic objects for network Interface sub-layers — ‘ifMIB’ 1.3.6.1.2.1.31

- *MIB Objects — ‘ifMIBObjects’ 1.3.6.1.2.1.31.1*
 - *Extension to the Interface Table — ‘ifXTable’ 1.3.6.1.2.1.31.1.1*

oid = “1.3.6.1.2.1.31.1.1.1.ifIndex”

ifName: The textual name of the interface
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.2.ifIndex”

ifInMulticastPkts: The number of multicast packets received, 32 bit
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.3.ifIndex”

ifInBroadcastPkts: The number of broadcast packets received, 32 bit
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.4.ifIndex”

ifOutMulticastPkts: The number of multicast packets transmitted, 32 bit
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.5.ifIndex”

ifOutBroadcastPkts: The number of broadcast packets transmitted, 32 bit
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.6.ifIndex”

ifHCInOctets: The total number of octets received on the interface, 64 bit
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.10.ifIndex”

ifHCOctets: The total number of octets transmitted by the interface, 64 bit
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.14.ifIndex”

ifLinkUpDownTrapEnable: Indicates whether or not linkup/linkDown traps should be generated for this interface
1 = enabled
2 = disabled
Access: read-write

oid = “1.3.6.1.2.1.31.1.1.1.15.ifIndex”

ifHighSpeed: An estimate of the interface current bandwidth in units of 1M bits/sec
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.16.ifIndex”

ifPromiscuousMode: Indicates whether this interface only accepts packets/frames addressed to this station (false) or accepts all packets/frames transmitted on the media (true). The value does not affect broadcast and multicast packets/frames.
1 = true
2 = false
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.17.ifIndex”

ifConnectorPresent: Indicates whether the interface has a physical connector
1 = true
2 = false
Access: read-only

5.1.5 Private Managed Objects

MIBs for Contemporary Controls — 1.3.6.1.4.1.17384

- **Switch Family — 1.3.6.1.4.1.17384.1**
 - **Switch Series — 1.3.6.1.4.1.17384.1.1**

oid = "1.3.6.1.4.1.17384.1.1.1.0"

temperature: Indicates the internal temperature of the switch in degrees Celsius

Type: string

Access: read-only

5.1.5.1 Relay Group – 1.3.6.1.4.1.17384.1.1.2

oid = "1.3.6.1.4.1.17384.1.1.2.1.0"

switchFaultStatus: The fault status of this switch

Type: integer

1 = fault has occurred

2 = no fault has occurred

Access: read-only

5.1.5.1.1 Fault Table – ‘faultTable’ 1.3.6.1.4.1.17384.1.1.2.2

oid = "1.3.6.1.4.1.17384.1.1.2.2.1.ifIndex"

portMonitoringFaultStatus: Fault status for each port

Type: integer

1 = fault has occurred

2 = no fault has occurred

Access: read-only

5.1.5.2 RapidRing Group – 1.3.6.1.4.1.17384.1.1.3

oid = "1.3.6.1.4.1.17384.1.1.3.1.0"

ringEnableStatus: The RapidRing is enabled or disabled

Type: integer

1 = RapidRing is enabled

2 = RapidRing is disabled

oid = "1.3.6.1.4.1.17384.1.1.3.2.0"

switchMode: The switch is a Master or a Slave

Type: integer

1 = Master

2 = Slave

oid = "1.3.6.1.4.1.17384.1.1.3.3.0"

ringID: The group number of the RapidRing

Type: integer

1 = group number 1

2 = group number 2

oid = "1.3.6.1.4.1.17384.1.1.3.4.0"

networkStatus: The network topology

Type: integer

1 = Ring Not Available (because RapidRing is not enabled)

2 = Ring Complete

3 = Ring Incomplete

5.1.6 Message Format for SNMP Operations

Five SNMP operations are used in SNMP version 1: *get*, *get-next*, *set*, *get-response*, *trap*. The first four commands are used to send and receive information for managed objects and use the same message format. *Trap* uses a different format discussed in Section 5.1.6.2.

5.1.6.1 Format of Command Messages

Each command message contains a header and a protocol data unit (PDU).

5.1.6.1.1 Message Header

The fields in the message header contain:

Version — 0, indicating SNMP version 1

Community string — The community string, which authorizes NMS access to the switch

5.1.6.1.2 Message Protocol Data Unit (PDU)

SNMPv1 PDUs contain a specific command (*get*, *set* etc.) and operands that indicate the object instances involved in the transaction.

The fields in the PDU contain:

PDU type — Indicates the command type: *get(0xA0)*, *get-next(0xA1)*, *set(0xA3)*, *get-response(0xA2)*

Request ID — A 4-octet integer used to match response to queries

Error Status — A single octet integer containing a value of zero in a request and the following error status in a response

noError (0): No problem

tooBig (1): The response to your request was too big to fit into one response.

noSuchName (2): An agent was asked to get or set an OID that it can't find; i.e., the OID doesn't exist. It can be used for an unsupported object.

badValue (3): A read-write or write-only object was set to an inconsistent value.

readOnly (4): This error is generally not used.

genErr (5): None of the previous errors

Error Index — A single octet integer that associates an error with a particular object identifier (OID). Only the response operation sets this field. Other operations set this field to zero.

Variable binding — contains a sequence of OIDs and values

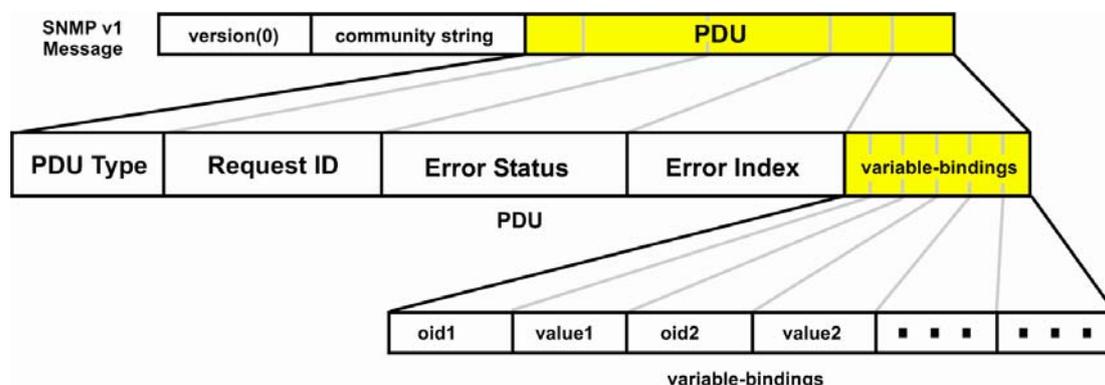


Figure 46 — Format of the Command Message for SNMP Version 1

5.1.6.2 Traps for SNMPv1

5.1.6.2.1 Format of Trap Messages

Each trap message contains a header and a protocol data unit (PDU).

5.1.6.2.2 Trap Header

The fields in the trap message header contain:

Version — 0, indicating SNMP version 1

Community string — The community string, which authorizes NMS access to the Trap Protocol Data Unit (PDU)

The fields in the PDU contain:

PDU type — 4 (indicates version 1 trap PDU)

Enterprise — Identifies the type of managed object generating the trap. For switch traps, the value is as follows:

Generic Trap — value is *SNMP* (1.3.6.1.2.1.11)

Agent-address — The IP address of the originating agent

Generic-trap — 0 to 6, indicating the generic trap type. See Section 5.1.6.2.3 for descriptions of the generic-trap types.

Specific-trap — Indicates the specific trap type. This field is only interpreted when the generic trap type is 6, *enterpriseSpecific*.

Time-stamp — Seconds since last power cycle

Variable bindings — One or more OIDs (object identifiers) paired with the corresponding values. A variable is an instance of a managed object. These pairings provide more information about the event.

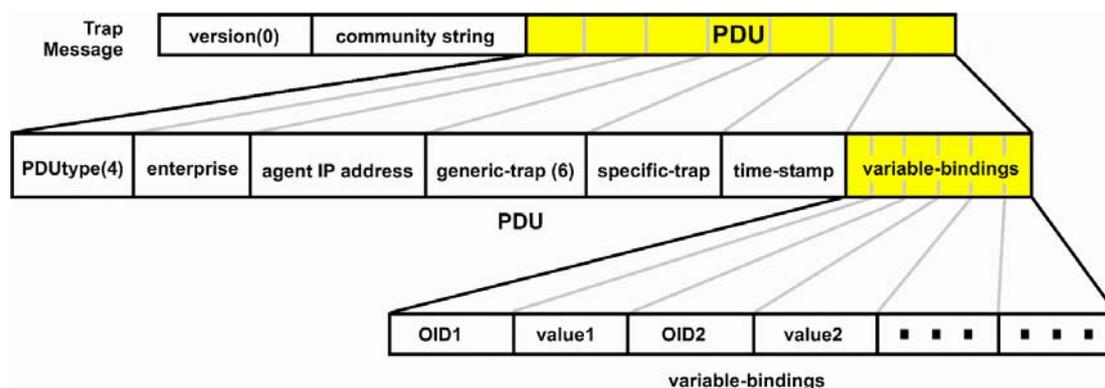


Figure 47 — Format of the Trap Message for SNMP Version 1

5.1.6.2.3 SNMP Generic Traps

enterprise = 1.3.6.1.2.1.11

Generic-trap = 0

coldStart: Signifies that the sending protocol switch is reinitializing itself such that the agent's configuration or the protocol switch implementation may be altered

Generic-trap = 1

warmStart: Signifies that the sending protocol switch is reinitializing itself such that neither the agent configuration nor the protocol switch implementation is altered

Generic-trap = 2

linkDown: Signifies that the sending protocol switch recognizes a failure in one of the communication links represented in the agent's configuration

Generic-trap = 3

linkUp: Signifies that the sending protocol switch recognizes that one of the communication links represented in the agent's configuration has come up

Generic-trap = 4

authenticationFailure: Signifies that the sending protocol switch is the addressee of a protocol message that is not properly authenticated