Managed Switch Software Setup

In This Chapter...

Main Settings **4–2**
System Settings 4–2
Remote Access Security 4–4
Port Settings 4–6
Port Mirroring 4–8
Set IP per Port 4–9
Switch Time Settings 4–10
Manage Firmware 4–11
Install Firmware 4–12
Redundancy Settings **4–14**
Spanning Tree Settings 4–18
Spanning Tree Port Settings 4–21
Real-Time Ring Settings 4–23
RSTP Examples 4–24
Traffic Priority (Priority Queuing QoS, CoS, ToS/DS) **4–29**
QoS / CoS Settings 4–30
802.1p Tag Settings 4–31
Message Rate Limiting 4–32
QoS Example 4–33
Multicast Filtering (IGMP) **4–36**
IGMP Protocol Settings 4–37
Port Settings 4–38
IGMP Example 4–39
Virtual LANs (VLANs) **4–40**
VLAN Settings 4–41
VLAN Port Settings 4–43

VLAN with RSTP 4–44
VLAN Examples 4–46
Security Settings 4–51
Remote Access Security 4–51
Port Security Enables and Port Security MAC Entries 4–52
IPsec Settings 4–54
IKE Policy 4–57
IKE Pre-shared Keys 4–59
IKE Certificates 4–60
Monitoring Settings 4–62
Alarm (OK) Output 4–62
Modbus 4–63
Register Mapping: 4–64
SNMP Notifications 4–65

This page intentionally left blank

Main Settings

The Main Settings section is where the general network settings of the switch are configured, such as the IP address and security access User name and password.

System Settings

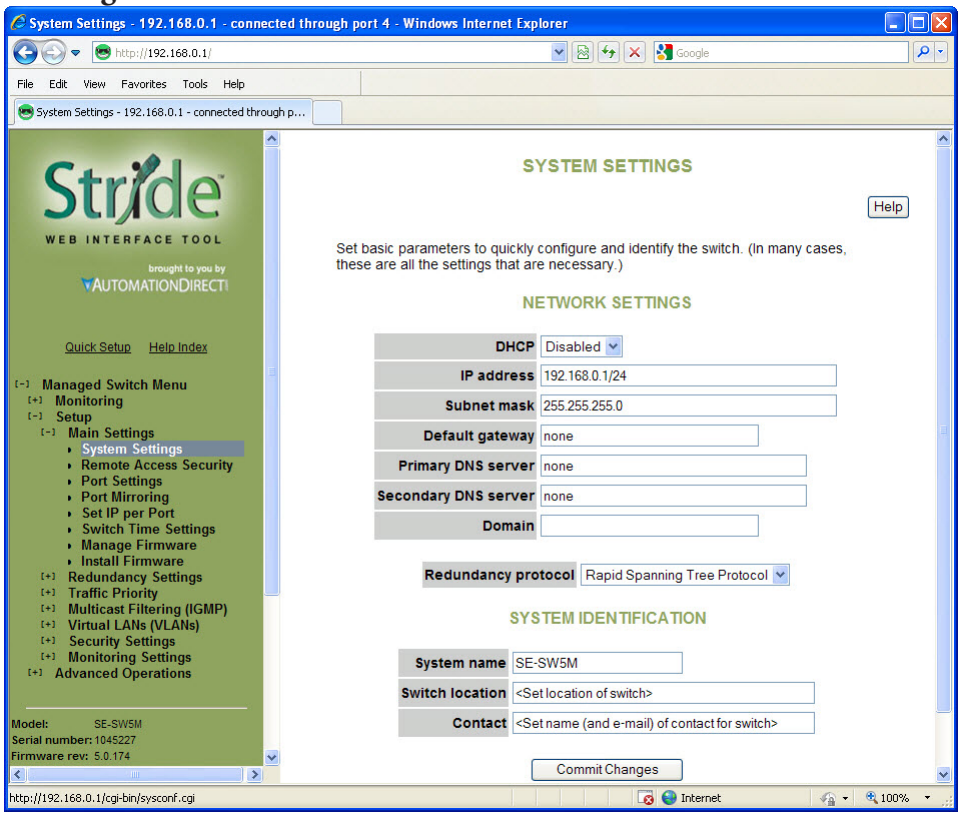

To control and monitor the switch via the network, it must be configured with basic network settings, including an IP address and subnet mask. Refer to the quick start guide in Chapter 2 to learn how to initially access your switch.

To configure the switch for network access, select Quick Setup from the Main menu to reach the System Settings menu. The settings in this menu control the switch's general network configuration.

DHCP Enabled/Disabled: The switch can automatically obtain an IP address from a server using the Dynamic Host Configuration Protocol (DHCP). This can speed up initial set up, as the network administrator does not have to find an open IP address.

NOTE: *If DHCP has been enabled, it will be necessary to connect to the console port serially or via USB in order to ascertain which IP address has been assigned so that you may be able to access the Switch using the web browser.*

IP address and Subnet Mask: The IP address of the switch can be changed to a user-defined address along with a customized subnet mask.

NOTE: *For additional security, advanced users can set the IP address to 0.0.0.0 to disable the web browser access . However, any features requiring an IP address (i.e., web interface, etc.) will no longer be available.*

Default gateway: A Gateway address is the address of a router that connects two different networks. This can be an IP address or a Fully Qualified Domain Name (FQDN) such as "domainname.org".

Primary DNS server: A DNS server address will be required if domain names are used in the switch settings. A Domain Name System Server converts a name, such as "domainname.org", into an IP address that is usable in the Ethernet messaging. Consult your network administrator for the proper DNS address for your network.

Secondary DNS server: A secondary DNS server can be configured in the case that the Primary DNS server is unreachable.

Remote Access Security

This screen allows you to set your remote access security settings.

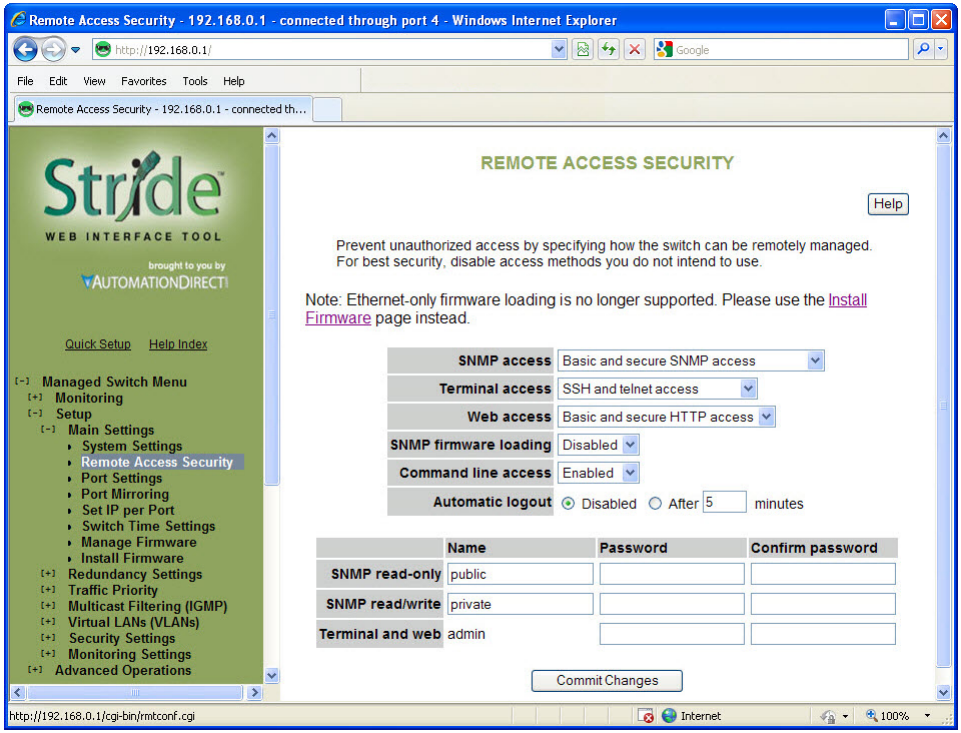

SNMP Access: Choose the level of SNMP access to allow.

- **None:** No SNMP access allowed.
- **SNMPv1 & SNMPv2 access (no passwords):** SNMP v1 and SNMPv2 access with community string (None) sent in clear text and no password required.
- **SNMPv3 access:** SNMPv3 access with encrypted password.
- **Basic and secure SNMP access:** SNMPv1, SNMPv2 and SNMPv3 access allowed.

Terminal access: Choose the type of terminal access to allow

- **None:** No terminal access to the switch will be allowed.
- **Non-secure access via telnet:** Non-secure access via telnet protocol. Remote access is possible through this protocol, although all information being transacted between server and client will be sent as clear text. Should security be of concern, use the Secure Shell protocol instead.
- **Secure access via SSH:** Secure access can be achieved through the use of the Secure Shell protocol (SSH), which implements strong authentication and secure communications using encryption. Using this protocol will ensure that your login information never gets sent as clear text, keeping the switch protected against possible attacks coming from the network.

- **SSH and telnet access:** The switch can be accessed through secure (SSH) and non-secure (telnet) terminal access.

The switch supports these encryption algorithms for SSH:

- 3DES
- Blowfish
- AES
- Arcfour

To take advantage of the SSH capability in the switch, you will need to use an SSH client program. There are many SSH client programs available for you to log onto the host (the switch).

Two open source SSH client programs are available on the Internet:

- Program name: OpenSSH for Windows: *http://sshwindows.sourceforge.net/*
- Program name: PuTTY: *http://www.chiark.greenend.org.uk/~sgtatham/putty/*

The SSH protocol requires some way for clients to be sure they are communicating with the intended host. The host computes a "fingerprint" based on its key and provides that to the client for verification. The first time a client program sees a fingerprint, it typically displays it and asks something like "The host is offering me these credentials, should I trust it?"

If you agree, the fingerprint is stored for later reuse.

For the system to be secure, the fingerprint used for comparison must be transmitted "out of band" (by a means other than the channel that is being secured by the fingerprint). In this case, via documentation.

The RSA fingerprint for the managed switch's encryption key is:
1e:0f:31:39:26:3f:23:8c:ba:7e:e9:d1:56:ff:98:f6

Web access: Choose the level of web access to allow.

- **No web access:** No web access allowed.
- **HTTP access:** Basic HTTP access allowed.
- **Secure HTTP (HTTPS) access:** Secure HTTP (HTTPS) required. Attempts to access the switch via http will be redirected to the secure protocol.
- **Basic and secure HTTP access:** Basic and secure HTTP access allowed.

SNMP firmware loading: Allows or disallows loading firmware via the SNMP protocol.

Command line access: Allows or disallows Command Line (CLI) access.

Automatic Logout: Specify the number of minutes of inactivity before terminal sessions automatically logout to prevent unauthorized access. The default is 5 minutes.

SNMP read-only name: This parameter sets the SNMPv2 community string and SNMPv3 user name that may be used by SNMP clients for read-only access of settings. Enter your own value if you wish to secure read-only access. (Default is "public").

SNMP read-only password: These parameters set the password for secure SNMPv3 access by the read-only user. SNMP passwords must be at least eight characters long. The default read-only password is 'publicpwd' (w/out quotes).

SNMP read/write name: This parameter sets the SNMPv2 community string and SNMPv3 user name that may be used by SNMP clients for read/write access to settings. Enter your own value if you wish to secure read/write access. (Default is "private").

SNMP read/write password: These parameters set the password for secure SNMPv3 access by the read/write user. SNMP passwords must be at least 8 characters long. The default read/write password is 'privatepwd' (w/out quotes).

Terminal and web: Password set here is used for Telnet and web access. To change the administrative password, select this option. (Default password is 'admin').

Port Settings

The switch comes with default port settings that allow you to connect to the Ethernet Ports without any configuration. Should there be a need to change the name of the ports, negotiation settings or flow control settings, you can do this in the Port Settings menu.

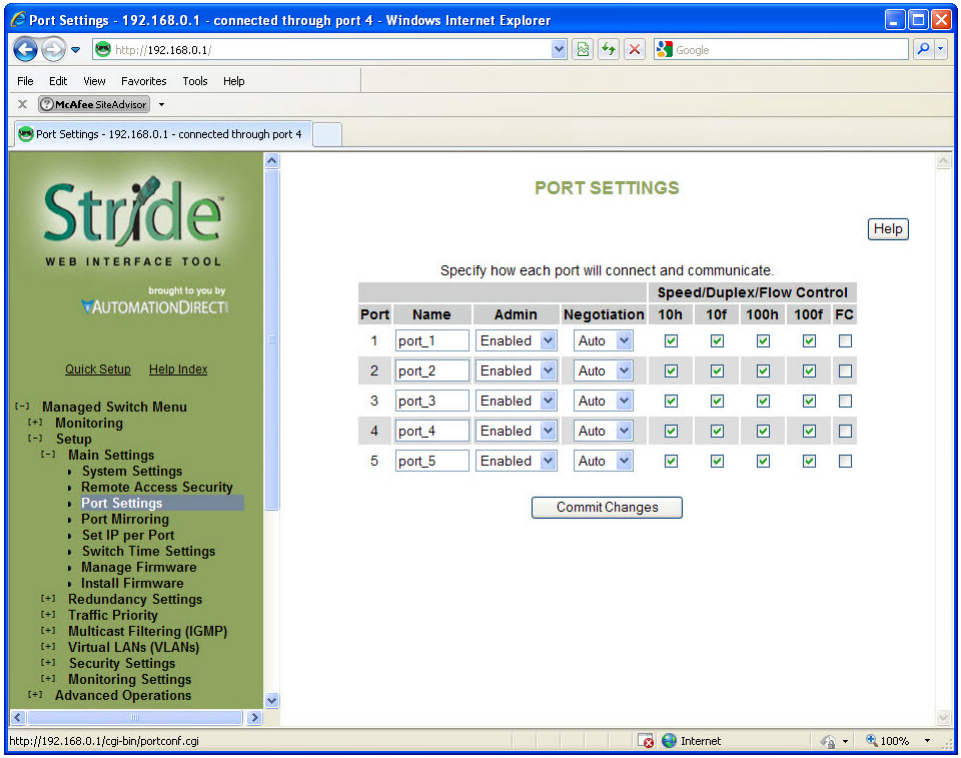

Port Name: Each port in the managed switch can be identified with a custom name. Specify a name for each port here.

Admin: Ports can be enabled or disabled in the managed switch. For ports that are disabled, they are virtually non-existent (not visible in terms of switch operation or spanning tree algorithm). Choose to enable or disable a port by selecting Enabled or Disabled, respectively.

Negotiation: All copper ports and gigabit fiber ports in the managed switch are capable of auto-negotiation such that the fastest bandwidth is selected. Choose to enable auto-negotiation or use fixed settings. 100Mbps fiber ports are fixed speed only.

Speed/Duplex/Flow Control: Each port can be set to allow speed and duplex to be negotiated to any or all Speed/Duplex/Flow control options. Network performance can be optimized by using Fixed Negotiation and restricting Speed/Duplex/Flow Control to a single value if network traffic is known.

These options are available:

- 10h – 10 Mbps, Half Duplex
- 10f – 10 Mbps, Full Duplex
- 100h – 100 Mbps, Half Duplex
- 100f – 100 Mbps, Full Duplex
- 1000f – 1000 Mbps, Full Duplex

On managed switches with gigabit combination ports, those ports will have two rows, a standard row of check boxes and a row labeled "SFP" with radio buttons. The SFP setting independently sets the speed at which a transceiver will operate if one is plugged in. Otherwise, the switch will use the fixed Ethernet port and the corresponding settings for it.

NOTE: *The SFP settings are NOT automatically sensed or negioated. If a 100 Mbps SFP is installed in the switch, that port must be manually set on the port settings page to 100 Mbps.*

Flow Control: Flow control can also be enabled or disabled, and is indicated by 'FC' when enabled. Devices use flow control to ensure that the receiving devices takes in all the data without error. If the transmitting device sends at a faster rate than the receiving device, then the receiving device will eventually have its buffer full. No further information can be taken when the buffer is full, so a flow control signal is sent to the transmitting device to temporarily stop the flow of incoming data.

Port Mirroring

In an unmanaged switch, each port is filtered to only send and receive Ethernet packets to devices physically connected to that port. This makes it impossible to view the messages occurring between two other devices from a third device (such as a PC running a tool like "Wireshark").

The mirroring option is ideal for performing diagnostics by allowing traffic that is being sent to and received from one or more source ports to be replicated out the monitor port.

Choose a monitor port.

Choose the source ports to be mirrored (monitored). For each source port choose the data to monitor: choose to monitor messages being sent (select Egress), messages being received (select Ingress) or messages being sent and received (select Both).

NOTE: *The Ingress Only option is not supported on SE-SW5M-xxx and SE-SW8M-xxx models.*

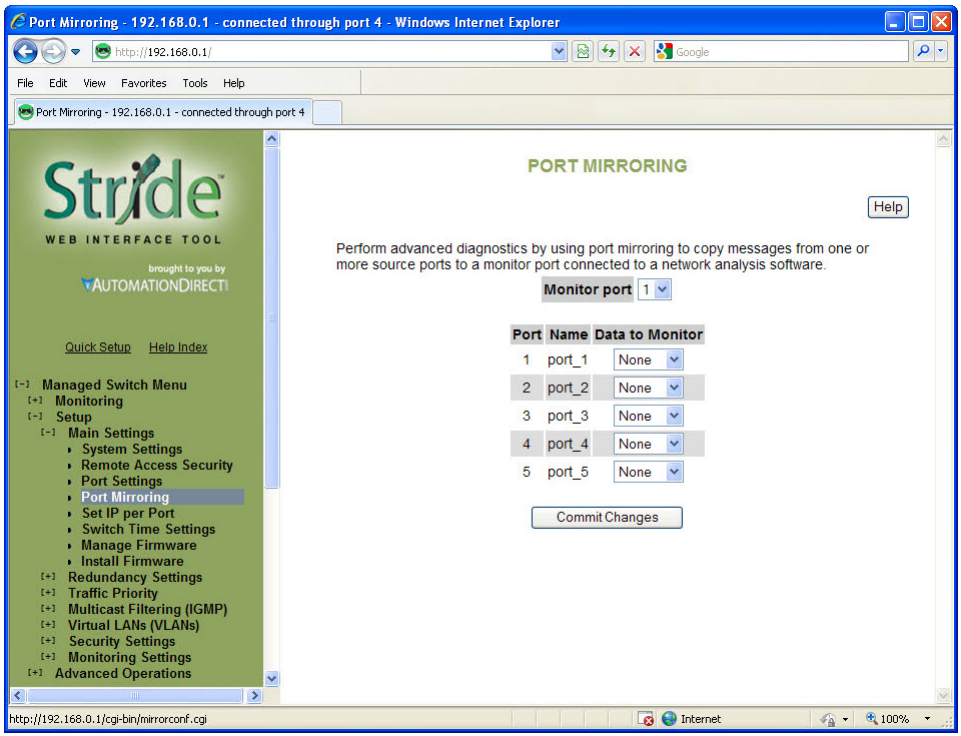

To view the traffic, connect a PC running network monitoring software (such as Wireshark) to the Monitor port.

Set IP per Port

The switch may provide an IP address to one device on each network port. This feature may be turned on and off for the whole switch and individually controlled for each port.

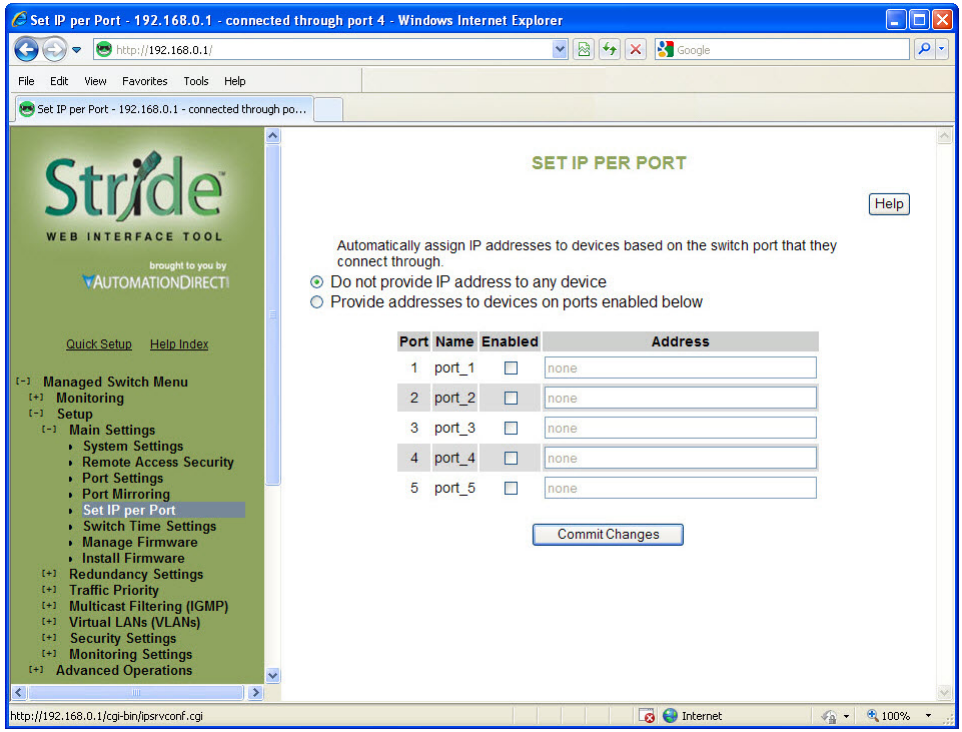

This feature is not a DHCP service. With Set IP per Port enabled on a port, the switch will respond to a DHCP request on that port with an IP address only.

For the feature to function properly, the host and network must meet the following criteria:

1. A single host must be directly connected to the switch port.
2. The host must not require a Subnet Mask to be offered.
3. The host and network must not require a Default Gateway to be offered.
4. There must be no other DHCP server on the network.
5. VLAN's must not be configured on the network.

NOTE: *This feature will not provide DHCP service required for Productivity CPUs, ECOM/ERM modules, and C-more panels.*

Switch Time Settings

This screen allows you to configure the switch's time settings, including time zone, current date and time as well as an NTP (Network Time Protocol) time server.

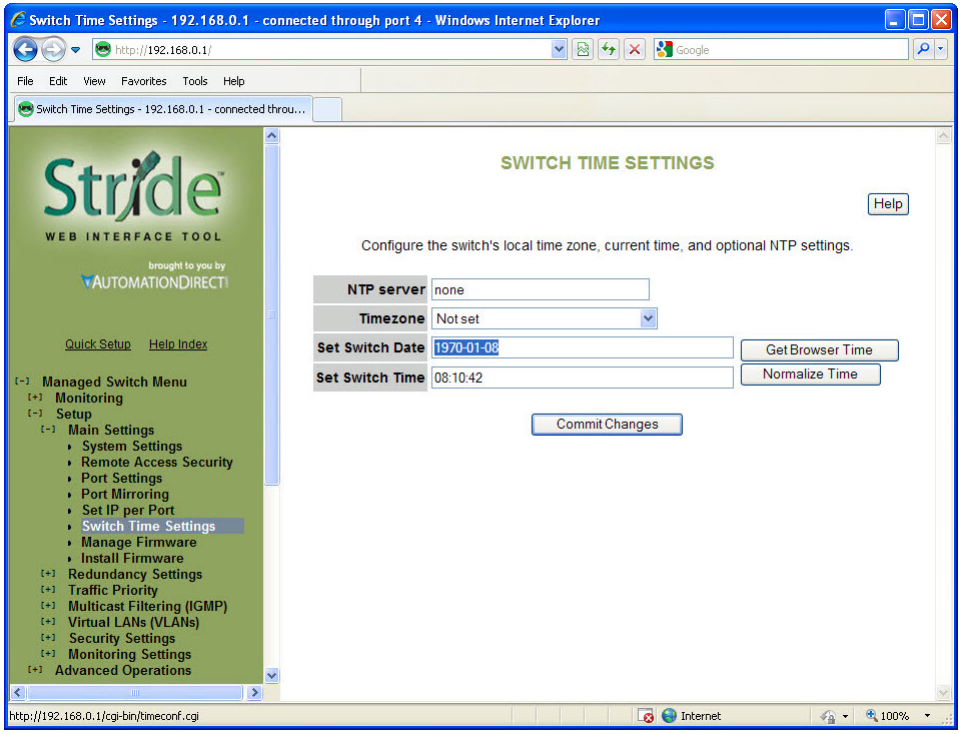

NTP server: You may specify an NTP server to automatically set the switch's clock. If a DNS server is configured, you may use a fully qualified domain name; otherwise, you must specify an IP address.

Timezone: This is the local timezone where the switch is installed. The switch will offset accordingly from the current time configured in the switch.

Set Switch Date: This is where the date is set for the switch. The format of the date is Year-Month-Day (YYYY-MM-DD).

Set Switch Time: This is where the time is set for the switch. The format of the time is hour:minute:second (HH:MM:SS).

There is also a "**Get Browser Time**" button to synchronize the switch's clock to your local browser's time and a "Normalize Time" button to format the time in a manner that the switch will view it. In other words, if the seconds are left out in time field, the normalize button will show the seconds field that will be set when the Commit button is pressed.

Manage Firmware

The Manage firmware page displays the current status of each of the two firmware images on a switch, and allows for changing which one will run the next time the switch is reset.

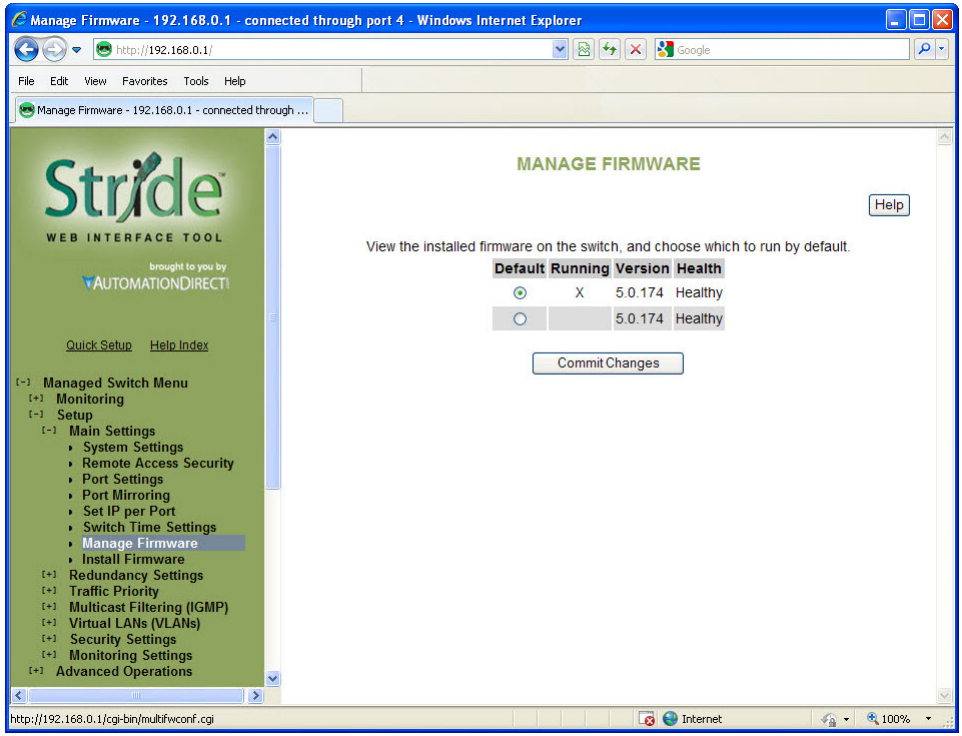

Default: Shows the current default firmware image to run when the switch is reset. May be changed to run a different firmware on the next reset.

Running: Shows the current running firmware image. This may be different from the current default firmware image if the switch failed to boot recently.

Version: Displays the firmware version number for each installed firmware. If the version cannot be determined, this will report "Unknown".

Health: Shows the health of each firmware image. The health can be one of the following:

- **Healthy:** The firmware is running or is expected to be in good enough shape to run.
- **Broken:** The firmware is known to be in a state that would prevent it from booting. The Default column will not allow this image to be selected for booting.
- **Unknown:** The firmware may be bootable, but the switch cannot be certain. This will happen if the switch is running the non-default firmware. This can happen if the default firmware somehow became corrupt, or if the switch lost power part way through booting.

If the firmware that is currently running is not the default, and the switch is reset without explicitly saving the default ("Commit Changes"), the current firmware will be run again. To

boot the firmware marked as the default, first, **commit this page and then reset the switch.**

Install Firmware

The Install Firmware page allows the inactive firmware (the selection not marked Running on the Manage Firmware page) to be replaced with a new version. To make the new version be the running version on the switch, after uploading the new version, you must:

- Go to the Manage Firmware page
- Select the new version as default
- Reset the switch

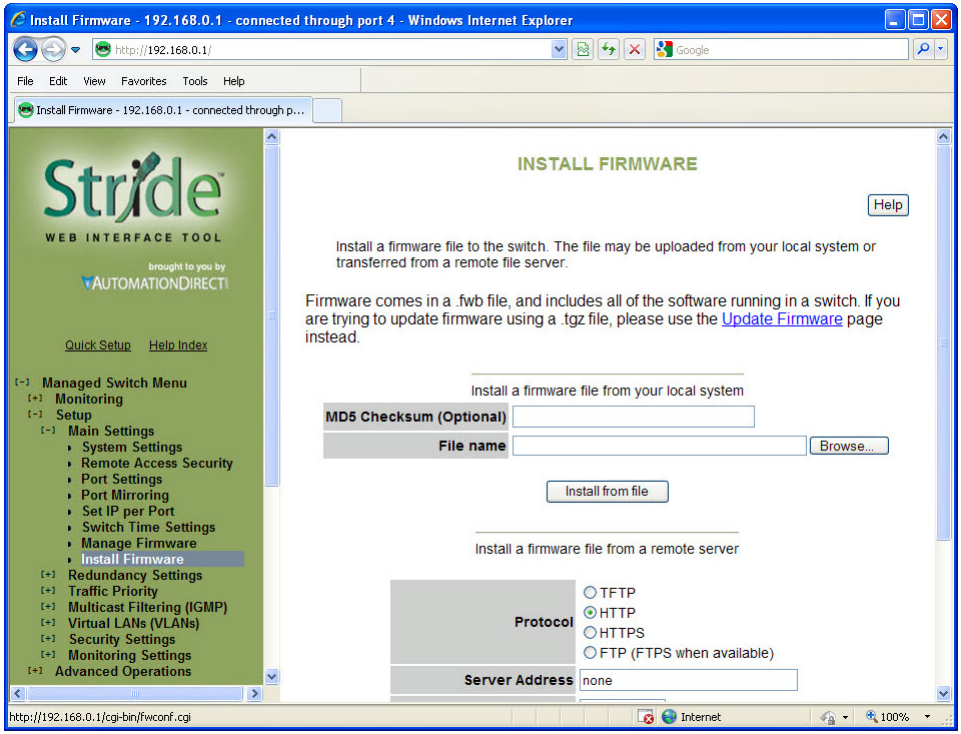

Firmware may be directly uploaded to the switch from the local system.

MD5 Checksum (Optional): If an MD5 checksum of the file is available, it may be entered into this field. Providing a checksum will ensure the firmware arrives at the switch intact and without any glitches. An MD5 checksum is not required.

File name: Use the "Browse" button to locate the .fwb firmware file.

Firmware may also be uploaded to the switch from a remote machine serving the .fwb firmware file. The server must be providing the file via TFTP, HTTP, HTTPS, FTP or FTPS.

Protocol: Choose one of the following protocols to retrieve the .fwb firmware file: TFTP,

HTTP, HTTPS, or FTP (FTPS when available).

Server Address: Enter the address of the server in this field. This may be an IP address, or a domain name if a DNS server has been configured on the System Settings page. Literal IPv6 addresses must be surrounded with square brackets. Example: the address fdda:2301: :2 must be entered as [fdda:2301: :2].

User Name: Enter the user name in this field if required by the server. Note that this is not available for TFTP.

Password: Enter the password in this field if required by the server. Note that this is not available for TFTP.

Anonymous Download: Check this box if no User Name and Password are required by the remote server.

Remote Filename: Enter the remote .fwb firmware file name into this field. The full path is required.

MD5 Checksum (Optional): If an MD5 checksum of the file is available, it may be entered into this field. Providing a checksum will ensure the firmware arrives at the switch intact and without any glitches. An MD5 checksum is not required.

Redundancy Settings

Another benefit of using managed switches over unmanaged switches is their redundancy capabilities. This allows you to have an Ethernet network with extra connections, so if one path between two points on the network fails, another path can be used to deliver messages. If one link or switch fails, another link or switch can take over transparently to prevent unnecessary down time. So why not just physically connect each of the switches in your network in various loop configurations such that there are always at least two paths going to and from each switch? That would create a broadcast loop that will bring a network to its knees very quickly.

In an unmanaged Ethernet network there can be only one path between any two ports on the network. If there is more than one path from one switch to another a broadcast message (and in some cases other messages) sent by the network will be forwarded until it completes a loop by returning on the second path. Since the switches forward all broadcasts and do not keep track of the messages they have sent, the returning message will be sent around the loop again and again. A single message circulating forever around a loop at high speed is clearly not a good thing, so no loops are allowed.

The limitations of having only one path are even simpler to see. If the one and only path fails for any reason, such as a broken cable or power failure at one of the switches, there are no paths left and no network traffic can get through. We need a way to add alternate paths without creating loops. A redundancy protocol such as RSTP, a loop prevention protocol, is used such that switches can communicate with each other to discover and prevent loops.

There are four methods of accomplishing redundancy in the Stride managed switches:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)
- Real-Time Ring

The Spanning Tree Protocols (STP, RSTP and MSTP) are an industry standard and are thus compatible with other manufacturer's managed switches for situations where both need to coexist and communicate. The recovery time, however, is slower with the Spanning Tree Protocols than with the proprietary Real-Time Ring protocol. The merits of both will be discussed in more detail below.

Spanning Tree Protocols:

In the diagram below all the links are the same speed, 100 Mbps. The root ports are those connected directly to the root bridge because they have the lowest path cost (only one hop). The paths that must go through another bridge (switch) have a higher path cost (two hops) and are designated as backup ports (decisions made internal to the switch by the Spanning Tree Protocol). The ports connected directly to end stations are assigned as edge ports (manually assigned on the Spanning Tree Port Settings page) so that RSTP doesn't waste time considering them.

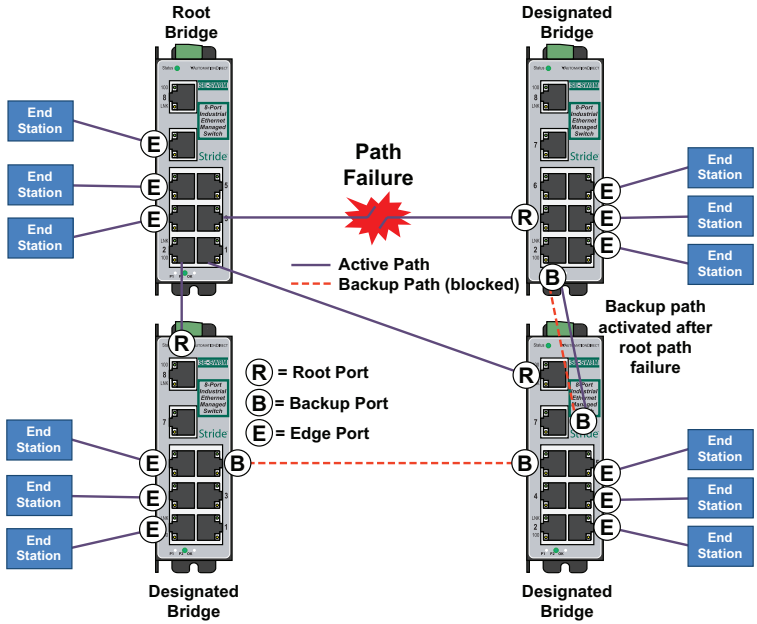

The Rapid Spanning Tree Protocol provides a standardized means for intelligent switches (also called bridges) to enable or disable network paths so there are no loops, but there is an alternative path if it is needed. Why is it called Rapid Spanning Tree Protocol?

- '**Rapid**' – it is faster than the previous (and completely compatible) version called Spanning Tree Protocol (STP).
- '**Spanning**' – it spans (connects) all of the stations and switches of the network.
- '**Tree**' – its branches provide only one connection between two points.

In a Spanning Tree network, only one bridge (managed switch) is responsible for forwarding packets between two adjacent LAN segments to ensure that no loops exist in a LAN. To ensure that only one bridge is responsible, all other bridges on the network must cooperate with each other to form a logical spanning tree that defines the pathways that packets should take from bridge to bridge.

The logical spanning tree has exactly one bridge that is assigned the role of root. All of the other bridges need to have exactly one active path to the root. The job of the root bridge

is to notify all bridges connected in the tree that there has been a topology change and restructuring of the tree is in progress (due to a communications link failure somewhere in the network or a new switch added in the network). The root bridge is determined by the bridge priority assigned to it and the MAC address.

By default, it is the bridge with the lowest MAC address that gets assigned the role as "root", but a specific bridge can be forced to be the root bridge by changing its bridge priority setting (a lower number with respect to other bridges means higher priority, set on the Spanning Tree Settings page).

Every communication path between each bridge (managed switch) on the network has an associated cost. This "path cost" may be determined by the speed of each segment, because it costs more time to move data at a slower speed, or the path cost can be manually configured to encourage or discourage the use of a particular network. For example, you may not want to use a particular high-speed link except when absolutely necessary because you pay a fee to a service providor for data using that path, while another path is free (no monetary cost).

The path cost is the cumulative cost of all the hops from the root bridge to a particular port on the network. A Spanning Tree network always uses the lower cost path available between a port and the root bridge. When the available network connections change, it reconfigures itself as necessary.

See the RSTP examples topic in this section for an example of how the path cost can be utilized to establish the primary and backup connections.

During the start-up of a Spanning Tree Network, all bridges (managed switches) are transmitting configuration messages (BPDUs) claiming to be the root. If a switch receives a BPDU that is "better" than the one it is sending, it will immediately stop claiming itself as the root and send the "better" root information instead. Assuming the working network segments actually connect all of the switches, after a certain period of time there will be only one switch that is sending its own root information and this switch is the root. All other switches transmit the root bridge's information at the rate of the root bridge's "hello time" or when the root bridge's BPDU is received on one of their ports.

The factor for determining which switch is the root (has the "best" root information) is the bridge priority and its tie-breaker, the switch MAC address. If a switch has more than one path to get messages from the root, other information in the configuration message determines which path is the best.

Once the root bridge is determined, all other switches see the root bridge's information and information about path (or paths) to the root. If more than one port provides a path to the root the non-root switches must decide which port to use. They check all of their ports to select the port that is receiving messages indicating the best path to the root.

The selected port for each bridge is called the root port. It provides the best path to communicate with the root. The best path is determined first by the lowest total path cost to the root (root path cost). Each port is assigned a cost (usually based on the speed) for messages received on that port. The root path cost for a given path is just sum of the individual port costs for that path. The lowest path cost indicates the shortest, fastest path to the root. If more than one path has the same cost the port priority assigned to each port, and its tie-breaker the port number pick the best path.

Recovery Time, Hops and Convergence:

The typical RSTP recovery time (time to start forwarding messages on the backup port) on a link-loss failure is <50ms per "hop". A hop is defined as a link between two switches. A link to an end station is not considered a hop.

The Max Age setting controls how long RSTP messages may circulate in the network. Since the largest value allowed for Max Age is 40, the largest RSTP network hop-diameter is also 40.

See the RSTP Examples topic in this section for a more detailed explanation about hops and recovery time.

The time it takes for all of the switches to have a stable configuration and send network traffic is called the convergence time. STP was developed when it was acceptable to have a convergence time of maybe a minute or more, but that is not the case anymore. Due to the increased demand for better convergence times, Rapid Spanning Tree Protocol was developed, bringing the normal convergence time for a properly configured network down to a few seconds. The RSTP takes advantage of the fact that most modern Ethernet links between switches are point-to-point connections. With a point-to-point link, the switches can quickly decide if the link should be active or not.

Spanning Tree Settings

The Spanning Tree Settings enable you to choose the redundancy protocol and set parameters related to that protocol.

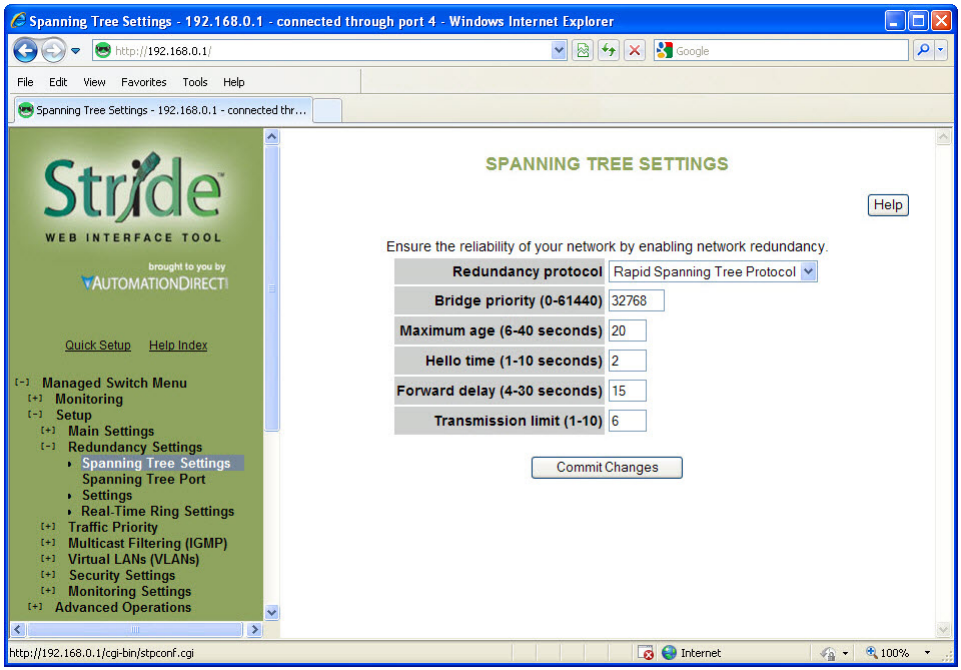

Redundancy Protocol: Choose the protocol by selecting STP (Spanning Tree Protocol), RSTP (Rapid Spanning Tree Protocol) or MSTP (Multiple Spanning Tree Protocol). A selection of None will disable this advanced feature. Choosing STP, RSTP or MSTP will allow the wiring of redundant networks (such as rings) for automatic failover. RSTP is compatible with STP so in most cases you should choose RSTP. Only choose STP if you want to force the switch to only use this protocol. STP/RSTP/MSTP use BPDUs (Bridge Protocol Data Units) to keep bridges informed of the network status.

MSTP is compatible with RSTP and STP but adds the ability to route VLANs over distinct spanning trees within an MSTP region. In order to configure spanning trees, you must create spanning tree instances using the STP configuration page and assign VLANS to them using the VLAN configuration page.

MSTP falls back to RSTP behavior outside of an MSTP region. A region is identified by the unique combination of Region Name, Configuration Revision and VLAN to MSTI mapping for each switch in that region. If those values match for linked switches running MSTP, those switches consider themselves to be in the same region.

CAUTION: If VLANS and redundancy (STP/RSTP/MSTP) are both enabled, situations can arise where the physical LAN is intact but one or more VLANs are being blocked by the redundancy algorithm and communication over those VLANS fails. The best practice is to make all switch-to-switch connections members of all VLANs to ensure connectivity at all times. Should you intend to use RSTP and VLANs at the same time, please see the "VLAN with RSTP" section for important information concerning the setup of your network. Otherwise, communication failures may occur.

Select none if you do not require the switch to manage redundant network connections. All ports will forward network traffic just as an unmanaged switch would. Otherwise RSTP (Rapid Spanning Tree Protocol) should usually be selected. A selection of STP or RSTP will allow redundant links between switches so those links can keep the network connected even when a primary link fails. RSTP is compatible with switches that only implement STP, an older version of the protocol. If STP is selected only the original STP format messages will be generated. Selecting STP reduces the chances of network packets being duplicated or delivered out of order, but at the expense of much longer reconfiguration time.

Bridge Priority (0 to 61440; Default = 32768): The bridge priority is used to determine the root bridge in the spanning tree. For MSTP, the bridge priority is used to determine the CIST root. The priority ranges from 0 to 61440 (default 32768) and must be a multiple of 4096. Lower numbers indicate a better priority.

By default, the bridge with the lowest bridge priority is selected as the root. In the event of a tie, the bridge with the lowest priority and lower MAC address is selected.

There are two ways to select a root bridge (switch). The first is to leave all the bridge priority settings at the default setting of 32768. When all the switches are set at the default priority, the managed switch with the lowest MAC address is selected as the root. This may be adequate for networks with light or evenly distributed traffic.

The second way to select a root bridge is to customize priority settings of each bridge. Customizing the bridge priority settings allows the network to select a root bridge that gives the best network performance. The goal is generally to have the network traffic pass through the network as directly as possible, so the root should be central in the network. If most messages are between one central server and several clients, the root should probably be a switch near the server so messages do not take a long path to the root and another long path back to the server.

Once you decide which switch should be the root, it should be given the best (numerically lowest) bridge priority number in the network.

Maximum Age (6 to 40; Default = 20): For STP, the max age indicates the maximum time (in seconds) that the switch will wait for configuration messages (BPDUs) from other managed switches. If that time expires, the switch assumes that it is no longer connected to the root of the network. If a link goes down in a way that the switch can detect the loss of link, it does not wait before reconfiguring the network.

RSTP waits 3 times the Hello Time instead of Max Age before assuming that it is no longer connected to the root of the network. However, Max Age is used to limit the number of hops Spanning Tree information may travel from the root bridge before being discarded as invalid. Furthermore, MSTP only counts hops that take place to or from switches outside the

MSTP region for this check. The value of Max Hops (below) is used to limit hops within an MSTP region.

NOTE: *Assign all Switches in an RSTP/STP network the same Max Age.*

The maximum age must satisfy the following constraints:

2 X (Hello Time + 1.0 seconds) < max message age < 2 X (forward delay – 1.0 seconds)

Hello Time (1 to 10; Default = 2): Configuration messages (BPDUs) are sent periodically to other bridges based on a time period labeled hello time. Decreasing the hello time gives faster recovery times; increasing the hello time interval decreases the overhead involved.

The hello time must satisfy the following constraints:

2 x (hello time + 1.0 seconds) < max message age < 2 x (forward delay - 1.0 seconds)

Forward Delay (4 to 30; Default = 15): The forward delay is a time (in seconds) used by all switches in the network. This value is controlled by the root bridge and is used as a timeout value to allow ports to begin forwarding traffic after network topology changes. If a port is not configured as an edge port and RSTP cannot negotiate the link status, a port must wait twice the forward delay before forwarding network traffic. In a properly configured network using RSTP (not STP) this setting has very little effect. For STP networks, setting the time too short may allow temporary loops when the network structure changes (switches turn on or off or links are added or broken). A longer time will prevent temporary loops, but network traffic will be disrupted for a longer time.

The default value for the forward delay is 15 seconds. If you change this setting, the switch will not allow a value unless it satisfies the following formula:

2 × (hello time + 1.0 seconds) < max message age < 2 x (forward delay - 1.0 seconds)

Transmission Limit (1 to 10; Default = 6): The transmission limit controls the maximum number of BPDUs that can be sent in one second.

The transmission limit can range from 1 to 10 messages/second (6 messages/second default). Increasing Transmission limit can speed convergence of the network but at the cost of configuration messages using a larger share of the available network bandwidth.

Region Name (MSTP): The region name is used together with the configuration revision and VLAN to MSTI (MST Instance) mapping to define an MSTP region.

Configuration Revision (MSTP; 0 – 65535): The configuration revision is used together with the region name and VLAN to MSTI (MST Instance) mapping to define an MSTP region.

Max Hops (MSTP; 6 to 40; Default = 20): Max Hops determines the maximum number of switches a BPDU will be propagated through within an MSTP region. This value is used to prevent old data from endlessly circulating within a region.

MST Instances: For MSTP, you can configure multiple spanning tree instances. Add an instance by clicking Add MSTI. For each MSTI, you can configure a name, the MST ID, and this bridge's priority in that spanning tree.

Spanning Tree Port Settings

Each port can be configured to tune the STP/RSTP/MSTP spanning tree. With MSTP, each spanning tree instance can be tuned independently.

Using MSTP, you can configure separate port settings for the CIST (Common Internal Spanning Tree) and for every spanning tree created by MSTP. Settings for individual MSTIs (Multiple Spanning Tree Instances) only affect ports connected to switches within the same MSTP Region.

By default, MSTIs inherit their settings from the CIST. To configure an MSTI individually, you must select it from the drop-down box and click the Customize button for the instance. Click Inherit if you want a spanning tree's values to be inherited from the CIST again.

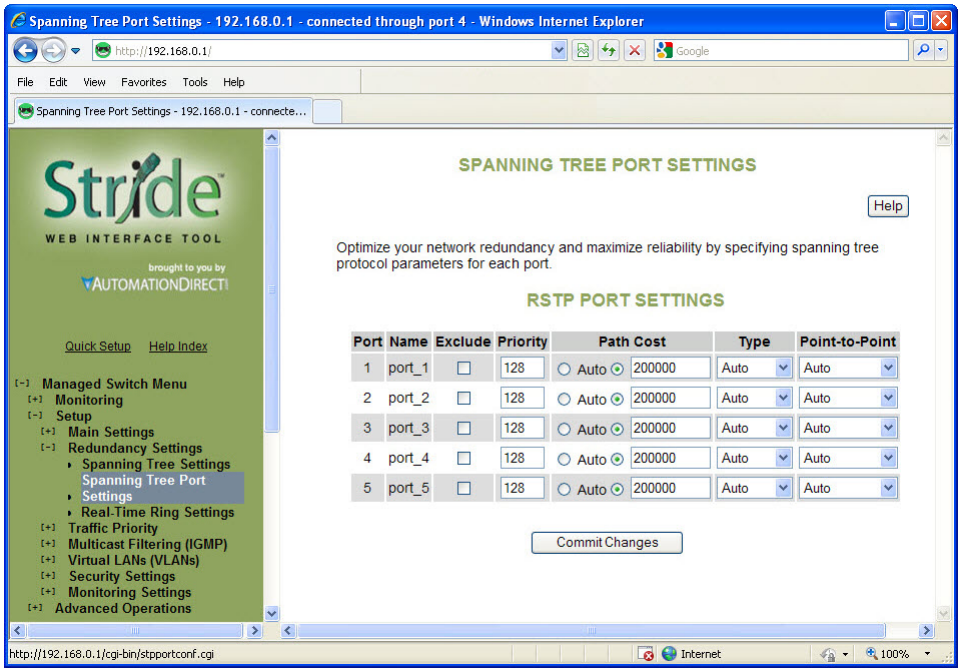

Exclude (Default = Included): Normally all ports should be included in determining the Spanning Tree network topology, either as a normal port or an edge port. It is possible to completely exclude a port, so that it will always forward network traffic and never generate or respond to network messages for RSTP or STP. Excluding a port is an advanced option that should be used only if absolutely necessary. The pair of ports assigned to a Real-Time Ring should be excluded from Spanning Tree.

This option excludes the port from all spanning tree instances and appears with the other CIST settings.

Port Priority (0 to 240; Default = 128): Selection of the port to be assigned "root" if two ports are connected in a loop is based on the port with the lowest port priority. If the root bridge fails, the bridge with the next lowest priority then becomes the root.

This option may be set per port per MSTI.

If the switch has more than one port that provides a path to the root bridge and they have the same root path cost, the selection of which port to use is based on the port priority. The port with the best (numerically lowest) priority will be used. If the port priority is the same, the switch will use the lowest numbered port. The port priority can range from 0 to 240 seconds (128 second default).

Path Cost (1 to 200,000,000; Default = 20,000 for 10 / 100 / 1000 ports and 200,000 for 10 / 100 ports): As with any network, there is an associated cost to go from a source location to a destination location. For RSTP, the root path cost is calculated based on the bandwidth available for that particular connection to the root bridge. The port with the lowest cost for delivering messages to the root is used to pass traffic toward the root.

The path cost can be assigned automatically based on the port speed, using the IEEE standard values of 200,000 for 100Mbps links and 2,000,000 for 10Mbps links, or the value can be specified in the range 1 to 200,000,000.

The default value depends on the capabilities of the port: 200,000 for 100 Mbps and 20,000 for 1000 Mbps ports.

This option can be set per port per MSTI.

See RSTP Examples for an illustration of how the path cost can be utilized to establish the primary and backup connections.

Type (Default = Auto): A port that connects to other switches in the network may be part of a loop. To ensure such loops do not occur, the switch will not put a port in the Forwarding state until enough time has passed for the spanning tree to stabilize (twice the forwarding delay, 30 seconds by default). However, if a port connects directly to a single device at the edge of the network, it may safely be put in Forwarding state almost immediately. The port Type controls the switch's assumptions about what is connected to the port.

- **Auto:** The port will initially be assumed to be an Edge port and go to Forwarding quickly. It will automatically adjust to being a Network port if BPDUs are received and revert to being an Edge port any time no BPDUs are received for 3 seconds.
- **Network:** The port will always wait a safe time before going to the Forwarding state.
- **Edge:** The port will initially be assumed to be a direct connection to a single device but will change to being a Network port if any BPDUs are received. Thereafter, it will always wait a safe time before going to Forwarding whenever a link is reestablished on the port.

This option can be set per port per MSTI.

Point-to-Point (Default = Auto): A port is part of a point-to-point network segment when there can be no more than one other network port connected to it. RSTP can decide whether it is safe to forward network traffic very quickly on point-to-point links to other managed switches, otherwise the port must wait many seconds (30 seconds by default, twice the forward delay) before forwarding network traffic. When set to Auto, full-duplex links are assumed to be point-to-point; half-duplex ports are not. This setting can be forced true or false if the automatic determination would be wrong.

Real-Time Ring Settings

The Real-Time Ring Settings page, accessed through the Redundancy Settings, allows configuration of Real-Time Ring protocol in supported switches.

A real-time ring increases network reliability by providing an alternative path for message flow in the event of a network segment failure. When a ring port detects a communications break, it quickly notifies the other switches in the ring. Messages are automatically rerouted through the alternative ring path within milliseconds.

STP (Spanning Tree Protocol) is more flexible than a ring configuration, but recovery times for spanning trees may be in the hundreds of milliseconds. The real-time ring protocol exchanges topological flexibility for recovery times in the tens of milliseconds.

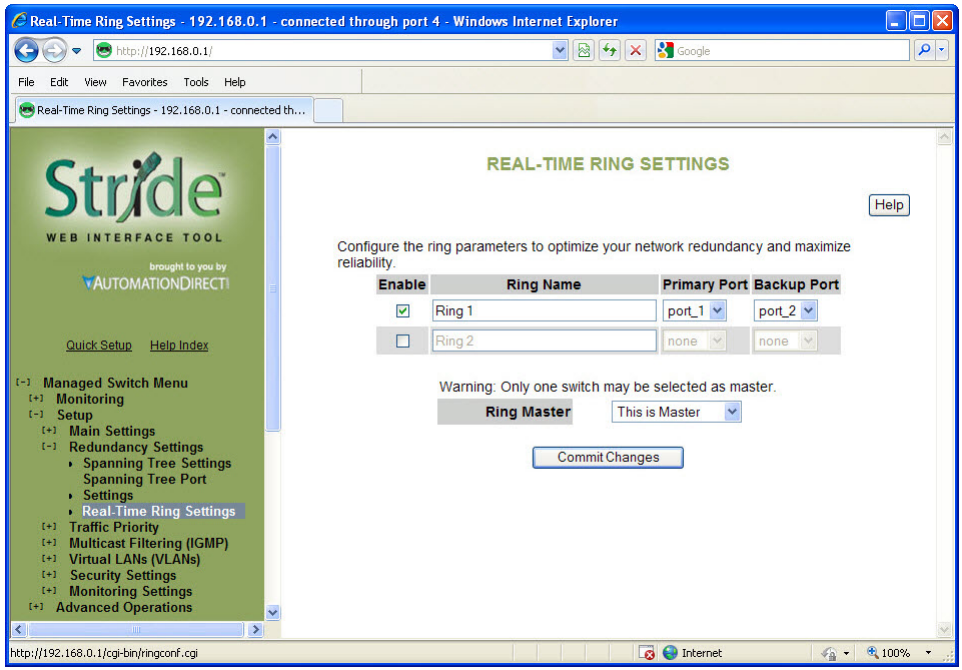

Activate a ring by selecting the appropriate Enable check box. You can configure one ring for every two ports on the switch.

When a ring is enabled, be sure to choose the two ports being used to connect the switch into that particular ring. To do so, pick ports from the Primary Port and Backup Port dropdown lists. Each port should be assigned to only one ring.

The pair(s) or ports assigned to ring(s) should be excluded from Spanning Tree on the Spanning Tree Ports Setup page.

The port defined as Backup will be blocked under normal operating conditions. By default, the switch with the lowest numbered MAC address in a ring will be the master switch, meaning that the communication in the ring will be blocked from one of the two ring ports of that switch. Only the master switch in a ring does this. You may designate a different

switch as the master switch by choosing “This is Master” from the Ring Master dropdown list for the desired switch. All other switches in the ring should be set to the default “Automatic” setting.

NOTE: *When a port is configured as a Ring port, that port cannot be used for communication to or through the Switch. It can ONLY be connected to another Ring port on a managed Switch or Real-Time Ring Switch.*

RSTP Examples

Example 1: Maximum “Hops” and Switches in a Redundant Ring:

The Max Age setting controls how long RSTP messages may circulate in the network. When a switch receives a message, it compares the age of the message with the Max Age (also carried in the message) and if the age has reached the Max Age, the message is discarded. Otherwise, the message age is incremented before the message is forwarded. Therefore, the maximum diameter of a RSTP network is controlled by Max Age. Since the largest value allowed for Max Age is 40, the largest RSTP network hop diameter is also 40.

Number of Hops vs. Recovery Time:

The diagram below shows a typical redundant ring network with 6 managed switches and 5 hops between stations.

The overall recovery time when there is a network segment failure is dependent on the number of hops. The recovery time is typically less than 50 ms per hop. Therefore, in the diagram below of a typical ring with 6 managed switches the overall recovery time would be less than 250 ms (5 hops x <50 ms).

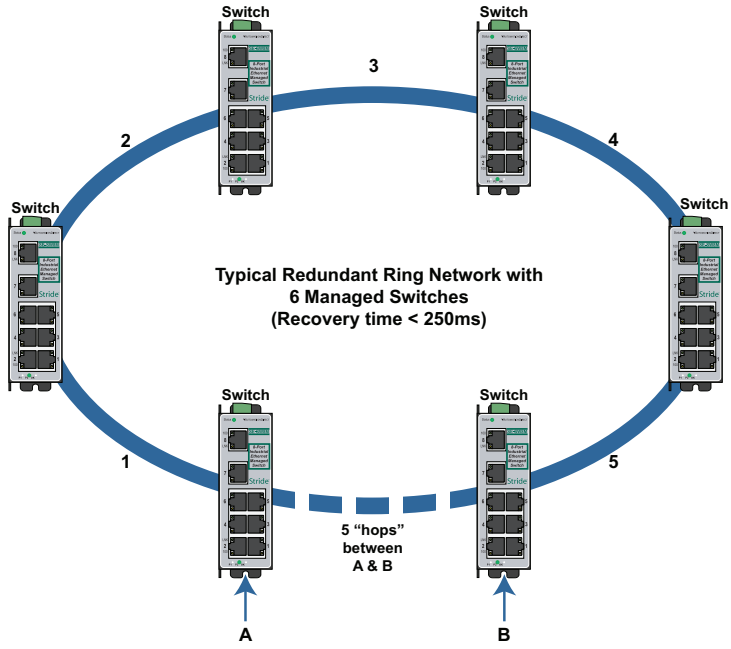

Example 2: Using Path Costs to Establish Primary & Backup Connections:

The path cost can be used to distinguish the best connections to use. You can assign a higher cost to pathways that are more expensive, slower or less desirable in any way. The managed switches will then add up the path costs to determine the best route back to the root switch. See the example below.

NOTE: *In most networks you may leave the path costs set to the default settings and allow the Switches to automatically determine the best paths.*

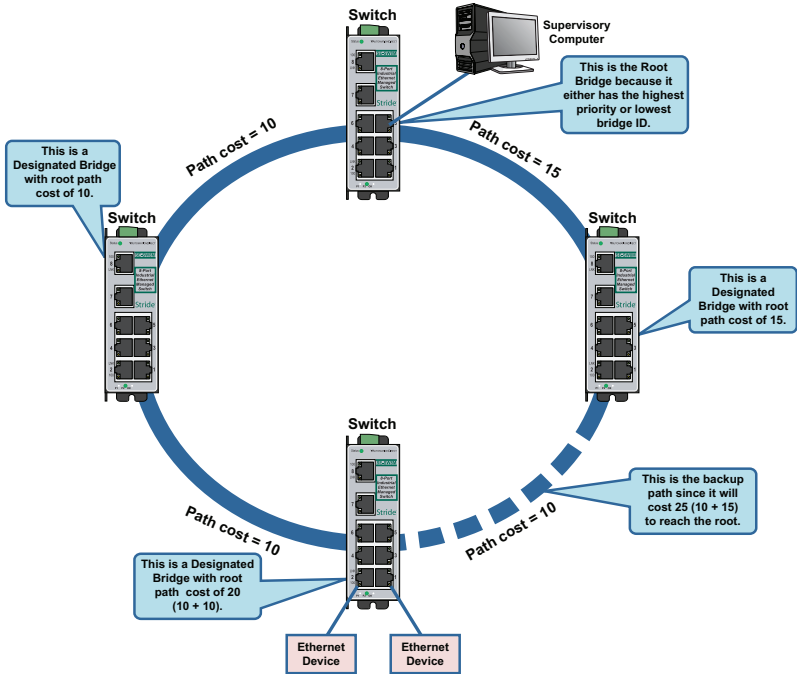
Switch
Supervisory Computer
This is the Root Bridge because it either has the highest priority or lowest bridge ID.
Path cost = 10
Path cost = 15
This is a Designated Bridge with root path cost of 10.
Switch
Switch
This is a Designated Bridge with root path cost of 15.
Switch
This is the backup path since it will cost 25 (10 + 15) to reach the root.
Path cost = 10
Path cost = 10
This is a Designated Bridge with root path cost of 20 (10 + 10).
Ethernet Device
Ethernet Device

Example 3: Ring Topology with only 1 Managed Switch (Bad idea!):

Implementing a ring topology with a single managed switch and several unmanaged switches is occasionally considered to try to save money. The topology is legal only if that single managed switch is a member of each ring. Although it is legal, it is not recommended, as the hypothetical scenario indicated below will explain.

Hypothetical Scenario:

An integrator wishes to implement a single Ethernet ring topology for the proposed network. Only one managed switch is used to connect to three or more unmanaged switches in the loop (Figure below).

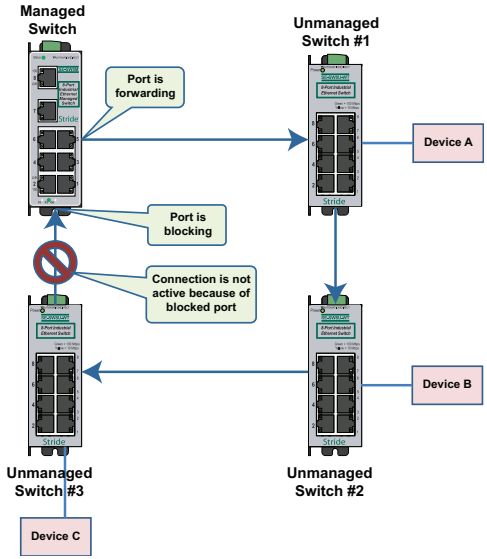

Initially, everything is working fine in the network. The managed switch detects the loop by seeing its own configuration messages and based on STP parameters, chooses one port to be in the forwarding state, and the other port to be in the blocking state. No loop is formed and device A can talk to device B.

Somewhere in the plant, a construction vehicle accidentally cuts the connection between unmanaged switch #1 and unmanaged switch #2. The managed switch in the network notices (typically around 6 seconds when connected to an unmanaged switch) that the port in blocking mode is not receiving configuration messages and transitions through the listening, learning, and forwarding states (Figure below).

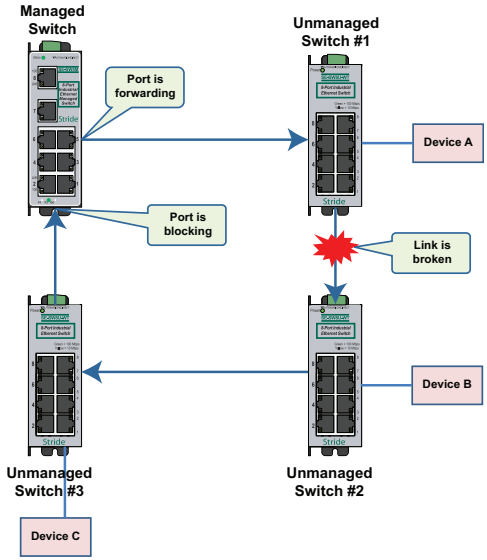

This would seem to have solved the problem as both ports in the managed switch are in forwarding mode, but it is not the case. Due to the fact that the other three switches are unmanaged, they do not have the intelligence to know that there has been a change in the network topology. switch #1 still points to switch #2 when device A is trying to talk to device B (across the broken Ethernet link). The bottleneck has been discovered, as we have to wait until the MAC table in switch #1 ages out its entries of device A and device B. The same applies for devices connected to switch #2 (B talking to A) and switch #3 (C talking to A).

As a result of this “money saving” configuration, the network redundancy performance is traded off and left at the mercy of the time it takes to age out MAC table entries in switches 1, 2, and 3. Depending on the model of unmanaged Ethernet switch, entries in the MAC table are usually aged out in a time period of 5 minutes or more.

This introduces at least 5 minutes of downtime for the plant, which could have a very detrimental cost with respect to the operation of the plant. By replacing switches 1, 2, and 3 with managed switches, the network convergence time is reduced to a less than a second. An additional benefit is that the network is not limited to only one redundant loop and can have a “mesh” of connections for a truly redundant network scheme at all points in the network.

Traffic Priority (Priority Queuing QoS, CoS, ToS/DS)

Without enabling special handling, a network provides a "best effort" service to all applications. This means that there are no assurances regarding the Quality of Service (QoS) for any particular application because all packets are treated equally at each switch or router. However, certain applications require deterministic response from the network to assure proper operation.

Consider a drilling machine in a plant that is controlled by a computer on a local network. The depth of the machine's drill is critical; such that if the hole is drilled is too deep, the material will have to be thrown out. Under normal conditions, the drill process is running smoothly (controller and computer are communicating efficiently over the network) but when another user on the network accesses records from an online database, the large volume of traffic can interfere with timely communication with the drill. A delay in communications between the drill and controller causes the drill to go too far and the material has to be thrown away. To prevent this from happening, we need to provide a certain QoS for all drill-controller communications so delay is avoided.

Numerous mechanisms exist to help assure reliable and timely network communication. The managed switch supports two common means of prioritizing messages: IP header and 802.1p user priorities.

The IP header is present in all frames and contains a priority field, which defaults to 0 and may be set as high as 255. This field is sometimes referred to as the Type of Service (ToS) field, or the Differentiated Services (DS or DiffServ) field.

Applications may add IEEE 802.1p tags, which contain a priority field that may be set from 0 to 7. Each value has a traffic type associated with it. For example, a tag of 5 is prescribed for video data.

The switch provides four priority queues for expediting outbound data. The 256 IP priorities and the 7 IEEE priorities are mapped into these ports in a way that optimizes throughput of high priority data.

Scheduling:

When choosing how to handle lower priority data, the switch can use strict or fair scheduling. This choice affects all queues on all ports.

Send All Priority Frames before any others: With strict scheduling, all data in the highest priority queue will be sent before any lower priority data, then all data from the second highest priority, and so on. This assures that high-priority data always gets through as quickly as possible.

Allow Lower Priority Frames through, a few at a time: With fair scheduling, a round-robin algorithm is used, weighted so that more high-priority than low priority data gets through. Specifically, the switch will send eight frames from the urgent queue, then four from the expedited queue, two from the normal queue, and one from the background queue, then start over with the urgent queue. This assures that the lower priority queues will not be starved.

QoS / CoS Settings

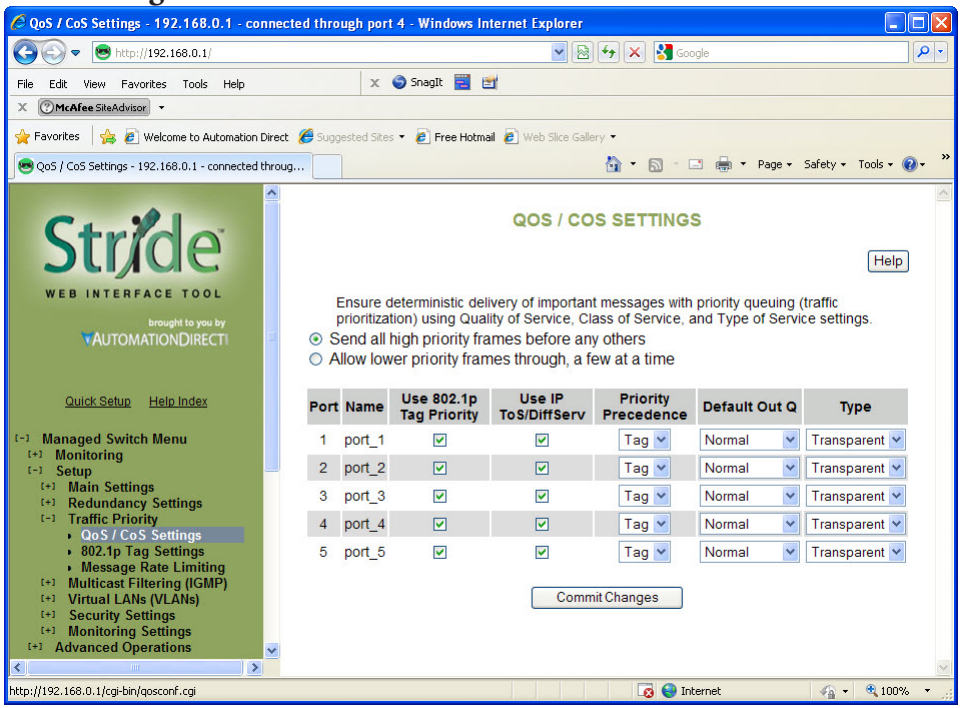

Use 802.1p Tag Priority: This setting controls whether the switch will honor IEEE tags if present in frames. When enabled, tagged data will routed to an outbound priority queue based on the configure tag mapping (See below). Disable this setting to ignore IEEE tags on all in-coming frames.

Use IP ToS/DiffServ: This setting controls whether the switch will honor priority fields in the IP header. When enabled, and not overridden by an IEEE tag, data will be routed to an outbound priority queue based on IPv4 Type of Service or Ipv6 Traffic Class. The priority queue will be the IP priority field value divided by 64. Disable this setting to ignore IP priority fields.

Priority Precedence: This setting controls which priority mark – IEEE tag or IP header – takes precedence if both are present and enabled. It has no effect if either Use Tags or Use IP is disabled.

Default Out Q: This setting controls the default priority to be assigned to frames when it cannot otherwise be determined. For example, if a frame without an IEEE tag arrived at a port where Use IP was disabled. Select an out-bound priority queue from the list.

Port Type: This setting controls how IEEE tags are handled in out-going data:

- **Transparent:** Maintains any tag that may have been present in a frame when it entered the switch.
- **Edge:** Removes tags from all out-going frames.

- **Network:** Adds a tag if none is present. The value of the tag is the queue number times two (six for queue 3, etc…)
- **Core:** All frames exiting this port will be tagged, in some cases double-tagged.

802.1p Tag Settings

The managedswitch has four Output Queues: Background, Normal, Expedited and Urgent with Background being the lowest priority and Urgent being the highest priority. In the IEEE 802.1p specification, there are eight different priorities that are carried in the tag. Configure each of the 802.1p priorities for the output queue that is appropriate. More than one 802.1p priority may be configured for a given output queue.

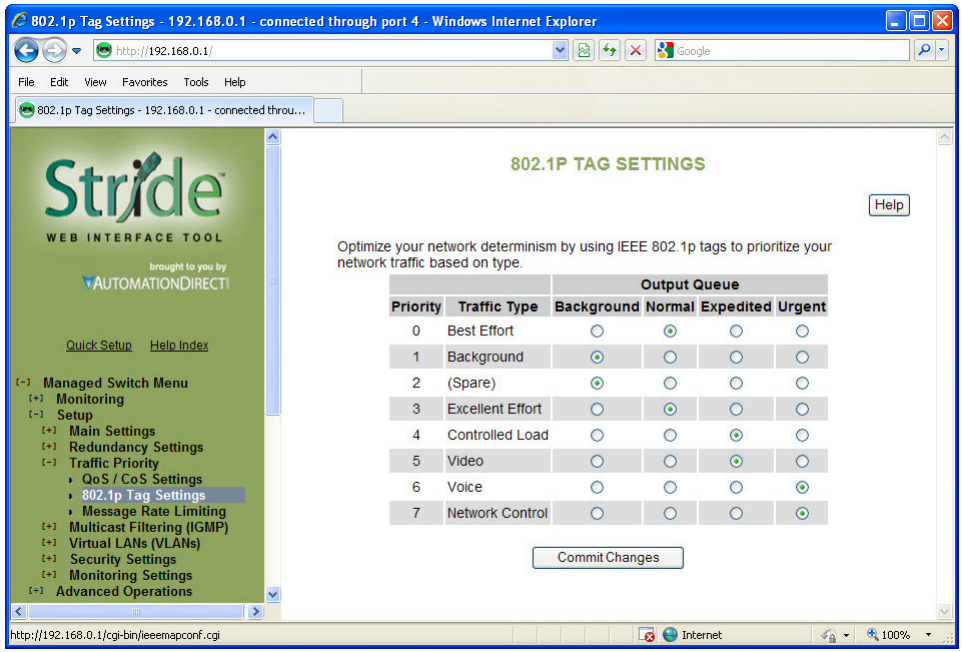

The table below indicates the defaults:

Managed Switch Output Queue					
Priority	**Traffic Type (802.1p priority)**	**Background**	**Normal**	**Expedited**	**Urgent**
0	Best Effort		**X**		
1	Background	**X**			
2	(Spare)	**X**			
3	Excellent Effort		**X**		
4	Controlled Load			**X**	
5	Video			**X**	
6	Voice				**X**
7	Network Control				**X**

Message Rate Limiting

Message Rate Limiting can prevent your switch and network from being overwhelmed by high volumes of broadcast and multicast messages. When enabled on a port, message rate limiting controls the percentage of messages which are allowed to be broadcast or multicast. Messages over the limit are dropped.

Poorly configured applications and devices or malicious users can flood your network with broadcast packets that are forwarded to all ports and can quickly consume most of a network's bandwidth. The managed switch provides some protection from such "broadcast storms" by allowing you to limit the rate at which these messages are accepted by the switch.

For each port, you may choose to limit the rate of broadcast and multicast messages accepted. Messages over the preset limit will be discarded.

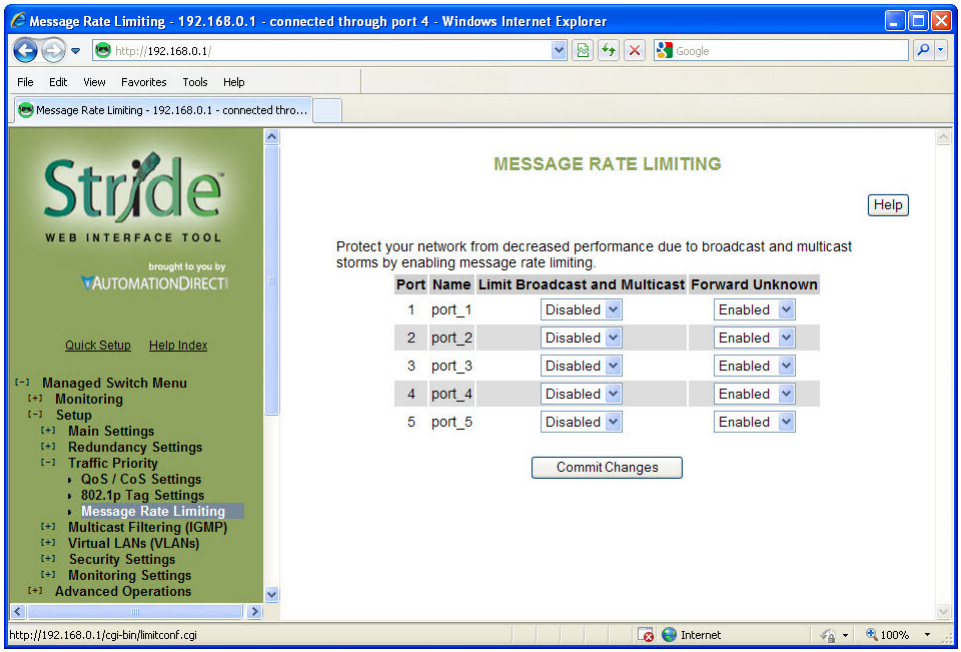

Limiting is done based on message type and priority. Broadcast and multicast messages are prioritized (by IP ToS) then limited to approximately the following rates:

Priority	Limit
Background	10% of link capacity
Normal	20% of link capacity
Expedited	40% of link capacity
Urgent	80% of link capacity

The exact limit depends on link speed.

Messages directly addressed to a single station (unicast messages) are not affected by message rate limiting.

Forward Unknown: By default, messages addressed to unicast addresses that have not yet been learned by the switch are flooded to all other ports. This is important for some protocols that transfer data primarily in one direction, but it can overwhelm smaller systems that do not expect a large amount of traffic. Forwarding of unknown unicast messages can be disabled on a port-by-port basis by disabling "Forward Unknown".

QoS Example

Let us investigate a detailed example of how to manage a network such that critical real time data will not be interrupted by data that is not as urgent

Hypothetical Scenario:

There is a power plant that is controlled by a central control system. In addition, because of security concerns, cameras have been mounted and installed at each location of mechanical control. The mechanical control devices and video cameras at each site communicate via Ethernet to their own switch. (For reasons of simplicity and clarity, we will assume that only video and control data reside on the network).

- **Problem:** Should any of the mechanical control devices receive delayed control data from the central control system, the power plant can't generate the maximum energy that it is capable of. Customers will experience brown outs, and the plant will be looked upon with negative scrutiny. It is therefore very important that the video traffic created by the cameras not delay critical data.
- **Goal:** To optimize the forwarding of critical real-time control data and minimize or eliminate the impact of video data traversing the network.
- **Solution:** Configure the switch such that video data has lower priority than control data by adjusting the priority queuing settings in the switch.

Configuration of the Switch:

As mentioned earlier in this manual, some applications require a certain Quality of Service (QoS) from the network to achieve a desired level of service. In this example, it is important that we achieve timeliness for control data. Without taking advantage of the switch's priority queuing abilities, we are using the best-effort network model. This means that the network will try to deliver all packets of information, but will not make any sort of promise or guarantees with respect to the timeliness of data for specific applications. Considering our control/video example, there is no guarantee that we can get the response time needed for control data if the video cameras are sending data at the same time.

A way to achieve the QoS desired is to prioritize network traffic. Prioritization of network traffic can be achieved even if the devices (video cameras and control systems) do not support selection or configuration of Quality of Service parameters.

Configure all the ports used to interconnect the switches as follows:

- Use 802.1p Tag Priority = Checked
- Use IP ToS/DiffServ = Checked
- Default Priority Precedence = Tag
- Output Tag = Add Tag

Where the data originates (the camera or control system), configure the QoS/CoS settings for the video camera ports as follows:

- Use 802.1p Tag Priority = Unchecked
- Use IP ToS/DiffServ = Unchecked
- Default Priority Precedence = Expedited
- Output Tag = Remove Tag

Also, configure the control system ports as follows:

- Use 802.1p Tag Priority = Unchecked
- Use IP ToS/DiffServ = Unchecked
- Default Priority Precedence = Urgent
- Output Tag = Remove Tag

In this way, the switches will handle the packets appropriately and tag them for handling elsewhere in the network.

At the destination, configure the control system port as follows:

- Use 802.1p Tag Priority = Checked
- Output Tag = Remove Tag

Also, configure the video concentrator port as follows:

- Output Tag = Remove Tag

Result:

Configuring the video data to have a lower priority than control data results in the QoS required for the control data.

In the following diagram, we have an IPm controlling a turbine and some torque converters. In addition, we have a video concentrator device that is collecting video data. Since the switch was configured such that video data (Triangles) has lower priority than control data (circles), we see that the control data gets sent out more often than the video data. For clarity, the diagram notes that untagged data in the network consists of open triangles and circles, while tagged data in the network consists of filled triangles and circles. This achieves the QoS needed for the control application.

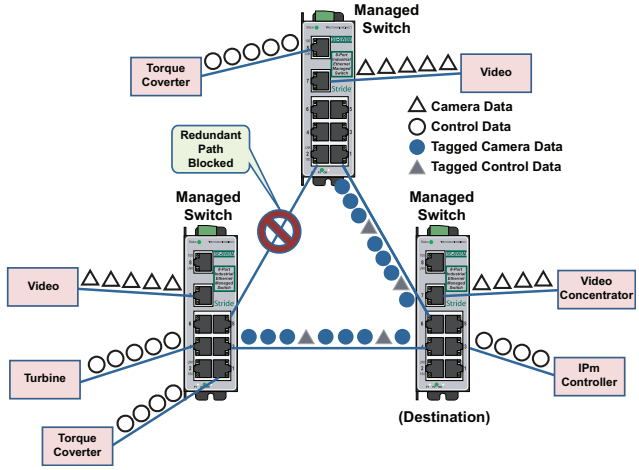
Managed Switch
Torque Coverter
Video
Camera Data
Control Data
Tagged Camera Data
Tagged Control Data
Redundant Path Blocked
Managed Switch
Managed Switch
Video
Video Concentrator
Turbine
IPm Controller
Torque Coverter
(Destination)

Multicast Filtering (IGMP)

IGMP (Internet Group Management Protocol) allows hosts and routers to work together to optimize forwarding of multicast traffic on a network. Without IGMP, all multicast packets must be forwarded to all network segments. With IGMP, multicast traffic is only forwarded to those network segments which connect interested hosts.

IGMPv1 provides a basic mechanism for hosts and routers to communicate about multicast groups. Routers send Query messages and hosts respond with group membership Report messages.

IGMPv2 adds a maximum response time to the Query and adds a Leave message to the protocol. IGMPv1 and IGMPv2 should not coexist on the same network. Also, IGMPv2 routers are expected to perform IGMPv1 on segments where IGMPv1 hosts are found.

An IGMP snooping switch performs many of the functions of an IGMP router. In passive mode, such a switch processes IGMP protocol messages sent by hosts and routers to configure efficient forwarding of multicast traffic. In active mode, a switch will also send its own queries to speed network convergence.

Periodically, routers and IGMP snooping switches in active mode send an IGMP Query on each attached network. (The query interval is generally around 1-2 minutes.) A host that wishes to be a member of a group sets a timer for a short, random delay when it sees the Query. If it sees a Report from another host before its timer expires, it cancels the timer and takes no further action until another Query is seen. If no other Report is seen, a Report is sent when the timer expires. The router or switch uses the Report to configure multicast forwarding.

The router or switch keeps track of how long it has been since the last Report on each port for each group. When the group expires, the router or switch stops forwarding multicast data to that port. Since the query interval is less than the expiration time, data for active groups continues to be forwarded without interruption.

IGMP Protocol Settings

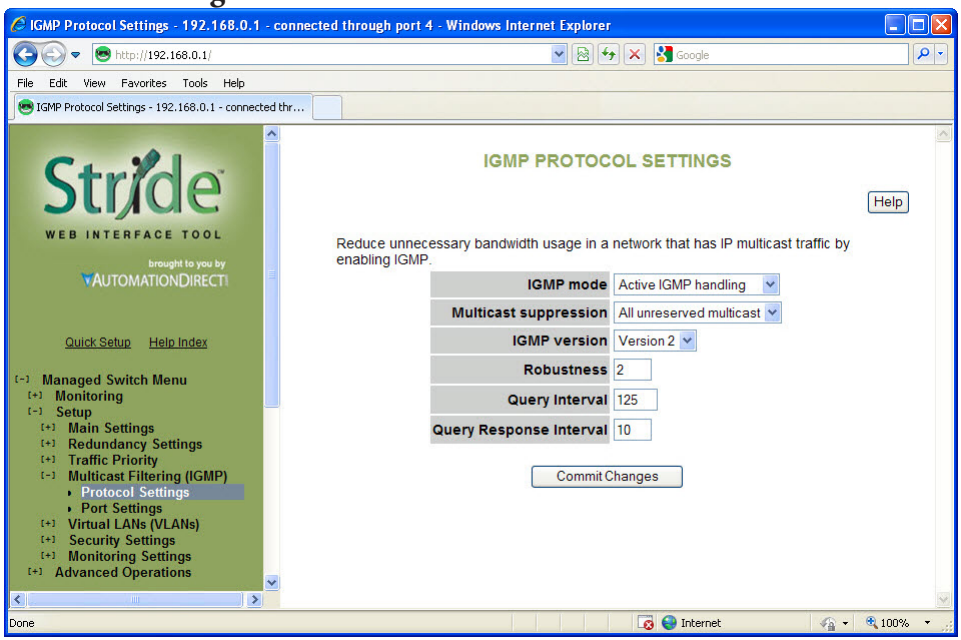

The default settings will allow the switch to regognize members of a multicast group and forward the multicast message to only members of that group.

IGMP Mode: This setting controls how the switch handles IGMP messages to determine how to forward multicast traffic.

- IGMP Disabled: Causes the switch to ignore IGMP messages. All multicast traffic will be sent to all ports.
- Passive IGMP handling: Causes the switch to listen to IGMP messages and configure forwarding of multicast traffic accordingly.
- Active IGMP handling: Causes the switch to act as an IGMP router, sending queries when needed and configuring multicast forwarding according to IGMP membership reports.

Multicast suppression: This enhanced feature can intelligently suppress multicast packets that no host has requested with IGMP.

- None: Multicast packets will be sent to all ports unless IGMP is enabled and one or more clients have sent IGMP Report requests.
- IP multicast groups: Multicast packets corresponding to IP multicast groups (with MAC addresses starting 01:00:5e) will be suppressed unless one or more clients have sent IGMP Report messages. Multicast packets with other addresses will be sent to all ports.
- All unreserved multicast: Multicast packets with reserved multicast addresses (01:80:c2:00:00:0x where x is 0..f) will be sent to all ports. All other multicast packets will be suppressed unless one or more clients have sent IGMP Report messages.

IGMP Version: This setting controls the highest IGMP version that the switch will use. All IGMP routers and snooping switches on a network should be configured for the same IGMP version. Select 1 or 2 as appropriate for your installation.

Robustness: This setting specifies how many queries may be lost without impacting forwarding as the switch tries to find IGMP hosts.

Query Interval: This setting specifies how often the switch will send IGMP queries in seconds.

Query Response Interval: This setting specifies the maximum time for hosts to respond to IGMP queries. (For IGMPv1, this is fixed at 10 seconds).

Port Settings

Like the default IGMP Protocol Settings, the default IGMP port settings will allow a switch to function in a network with multicast groups. Generally, the switch will dynamically learn which ports have IGMP routers attached to them by listening for IGMP Query messages. Under some circumstances, it is necessary to statically configure ports as leading to IGMP routers. Force the switch to forward IGMP messages to a specific port by choosing Static as the router type.

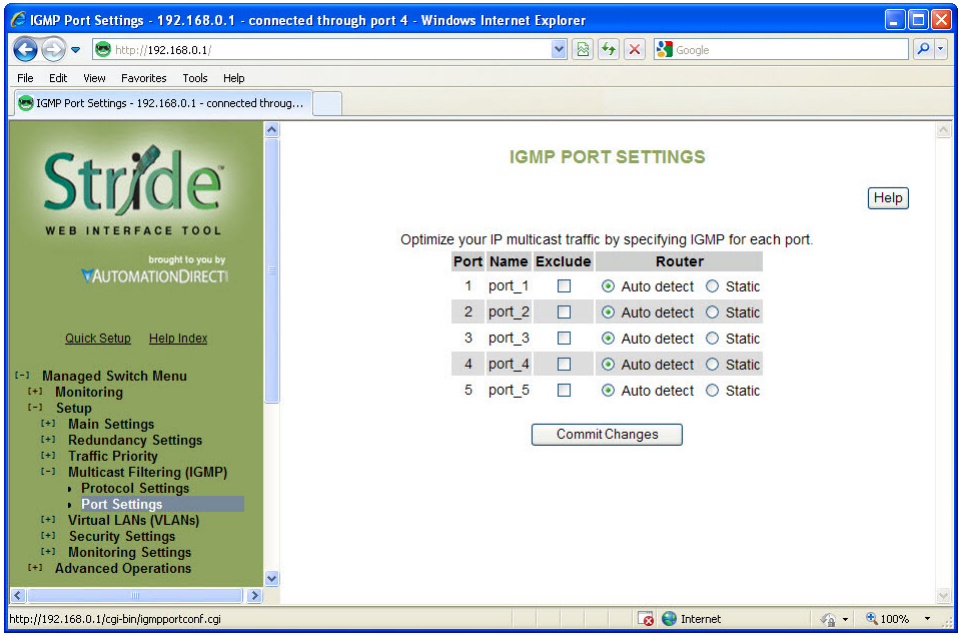

Exclude: A port may be excluded from IGMP processing. IGMP queries and reports received on an excluded port are ignored so devices reached via the excluded port cannot join multicast groups filtered by the switch. IGMP queries and reports will not be forwarded to the excluded port so IGMP routers reached via the excluded port will not know of memberships for devices reached by other ports.

Static Router: Specifies whether the switch should assume there is an IGMP router on this port even if no IGMP query messages are received.

IGMP Example

The Benefits of Enabling IGMP:

Take an already established control network that has an Ethernet device sending multicast data to several other Ethernet devices. Between the source of the multicast data, and the destination Ethernet devices that are interested in the multicast data, multicast packets might pass through a number of switches or routers.

To make this control network more efficient, the switches or routers should know how to handle the flow of multicast data by means of IGMP (Internet Group Management Protocol). Switches or routers that are not capable of supporting IGMP will not know what to do with the multicast data and forward multicast data out all ports. This will slow down the network.

Take a look at the following diagram, where the IGMP server is the source of the multicast data, and the IGMP hosts are the devices interested in receiving multicast data. On the network are two switches, where one has IGMP enabled and the other has IGMP disabled. We see that the switch with IGMP enabled only forwards multicast data to the interested host (Ethernet Station 2). The switch with IGMP disabled will not know where to send the multicast data; thus Ethernet Stations 4 and 6 unnecessarily receive multicast data even though only Station 5 is the interested host.

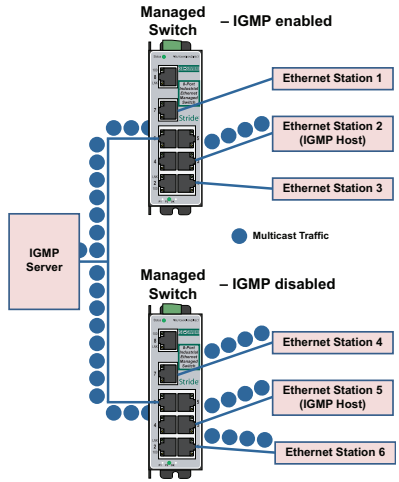

Virtual LANs (VLANs)

VLANs can segregate traffic flowing through a switch to improve bandwidth utilization or security. Segregation is done based on membership in a group of ports (port-based VLANs) or on IEEE 802.1Q tags which include a VLAN ID (tag-based VLANs).

A port-based VLAN limits forwarding traffic coming in a port to the group of ports to which that port belongs. For example, on a 10-port switch if ports 1, 3, 5, 7, and 9 were placed in a port-based VLAN, broadcast frames coming in port 3 would be sent to ports 1, 5, 7, and 9 (which are members of port 3's VLAN) but not to ports 2, 4, 6, and 8 (which are not members).

A port may be a member of two port-based VLANs, although results of this configuration are not always desirable or easily predictable. When initializing port-based VLANs the switch configures each port to be able to send data to all ports in all the port-based VLANs in which it is a member. For example, if one VLAN had ports 1-5 and another had ports 5-9, traffic from port 1-4 could go to ports 1-5, traffic from ports 6-9 could go to ports 5-9, and traffic from port 5 could go to all ports.

A tag-based VLAN is more common. A tag-based VLAN limits traffic based on the VLAN ID in a 'tag' associated with the frame. VLAN tags may be explicitly placed in frames by applications or switching equipment, or implicitly assigned to frames based on the switch port where they arrive.

VLAN IDs are 12-bits long providing 4096 possible IDs but several values are reserved:

- **0** = Indicates that the tag is not being used for VLAN routing but only to carry priority information. (See QoS/CoS topic).
- **1** = Used for switch configuration and management.
- **4095** = Not allowed by the 802.1Q standard.

VLAN Settings

There are several VLAN modes, which provide varying levels of flexibility and security.

Configuring VLANs requires creating VLANs on the VLAN Settings page and configuring ports for participation in the VLAN on the VLAN Port Settings page.

The VLAN settings page identifies which traffic a port can "listen" to. The VLAN Port Settings page identifies traffic a port can "talk" to. For ports to participate effectively in a VLAN, each port should be assigned to one VLAN on the VLAN settings page, then configured with that VLAN ID on the VLAN Port Settings page.

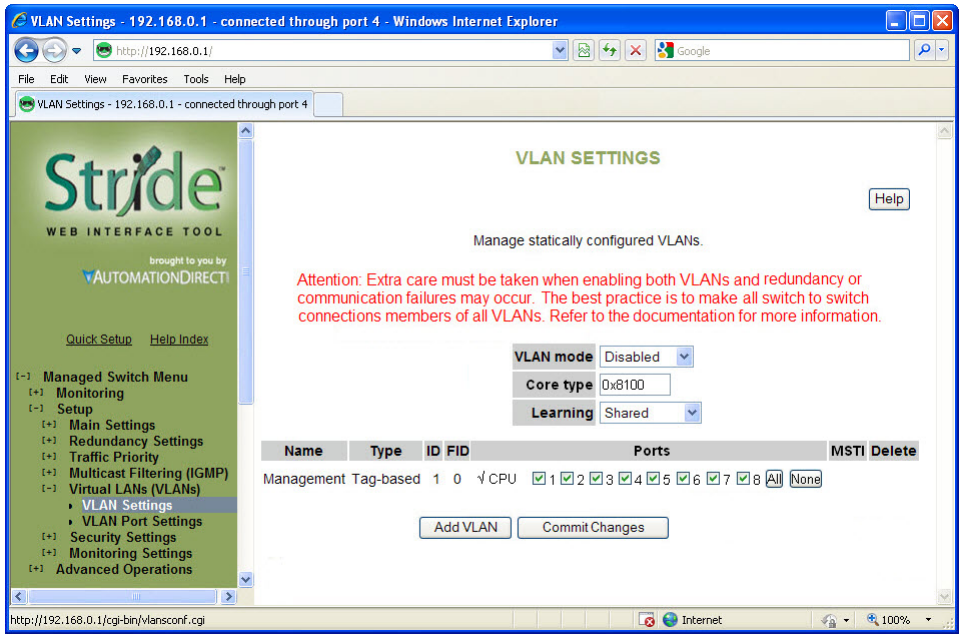

VLAN Mode:

- **Disabled:** No VLAN processing is done. VLAN IDs and port-based VLANs are ignored.
- **Port-Based:** Only port-based VLANs are used to route frames. VLAN IDs are ignored.
- **Standard:** (Most commonly configured) Port-based VLANs are ignored; all routing is done by VLAN ID. The source port of a frame need not be part of a VLAN for the frame to be forwarded.
- **Secure:** All routing is done by VLAN ID; however, if the source port of a frame is not a member of the target VLAN, then the frame is dropped. For example, if a tag-based VLAN for ID 1024 was configured to include ports 1-5 and a frame with VLAN ID 1204 in its tag arrived at port 6, the frame would not be forwarded.

Core Type: (gigabit switch only) Specify the Ethertype for double-tagged ("Q-in-Q") frames exiting ports of type Core. The value may be specified in hexadecimal with a 0x prefix.

Learning: This setting describes how different addresses on different VLANs are 'learned' by the switch.

- **Shared:** All VLANs (if MSTP is enabled, all VLANs assigned to the same MSTI) use the same forwarding database.
- **Independent:** The forwarding database used by each tag-based VLAN can be configured independently.

The switch supports up to 64 configurable VLANs including the management VLAN. To configure additional VLANs, click the "Add VLAN" button to create an empty row in the table. Then choose the name, ID information and ports for your VLAN. For tag based VLANs, the CPU should not be included in any VLAN other than the default management VLAN (1). The CPU should be included in port based VLANs.

To remove a VLAN, simply click the "X" in the delete column for that VLAN. When your settings have been changed as needed, click "Commit Changes" to save them.

Name: A mnemonic name for a VLAN such as "Cell 7", "Line 4", "Building 58". This is used for display only.

Type: The VLAN's type, port-based or tag-based.

ID: For tag-based VLANs, this is the ID to look for in the tag. This ID identifies the individual VLANs you create on your network. The VLAN ID must be specified in the range from 2 to 4094.

NOTE: *Take care when setting the management VLAN ID. If the device you are configuring from cannot work with VLANs and the port it is connected to does not have the proper PVID and port type setting the management VLAN may make the Switch inaccessible and require a local serial connection to reconnect.*

FID: For tag-based VLANs, the forwarding database to use when independent learning is enabled. If MSTP is running, all VLANs in the same MSTI must be configured to use the same forwarding database in independent learning mode. Shared learning automatically assigns a different forwarding database to each MSTI.

This filtering ID allows multiple VLANs to be grouped for easy filtering in the MAC address monitoring page.

Ports: The ports included in this VLAN. For tag based VLANs, the CPU should not be included in any VLAN other than the default management VLAN (1). The CPU should be included in port based VLANs.

To select the ports to include in this VLAN, check the box for each port you wish to include. Remember that if the "CPU" box is not checked, you will be unable to communicate with the switch from within this VLAN.

NOTE: *When working with tag-based VLANs, ports included in a VLAN may lead to other network devices (which require tags to properly route data) or to end devices, which cannot process VLAN tags. Use the VLAN Port Settings page to configure the appropriate type for each port.*

Delete: Select to delete the corresponding VLAN when changes are committed. When selected, this VLAN will be deleted when changes are committed.

VLAN Port Settings

Each switch port can be configured to control how VLAN tags are handled for frames coming in and going out of the port.

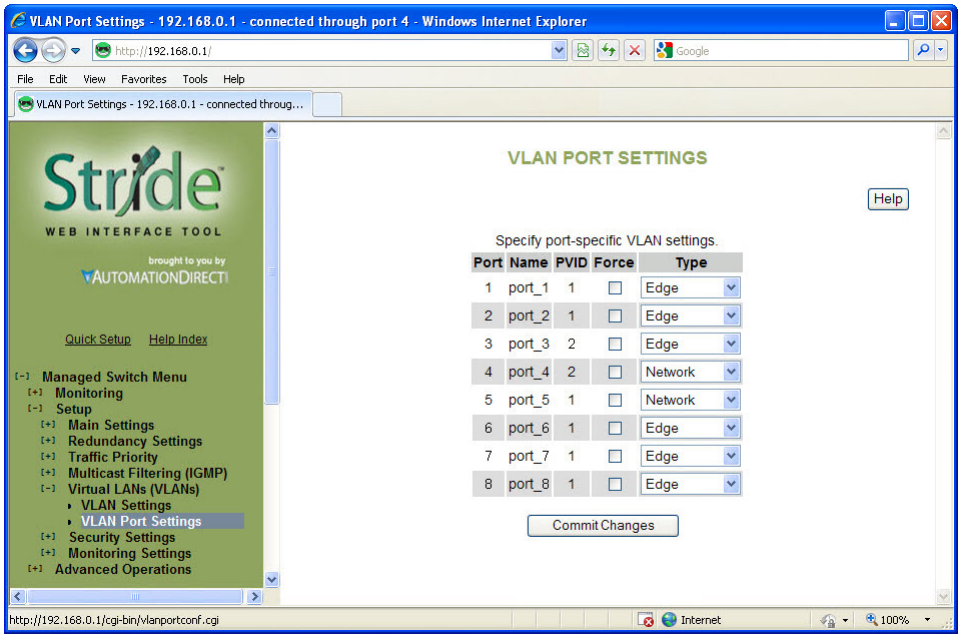

PVID: This is the port's default VLAN ID. It is applied to frames which arrive at the port without a VLAN tag or with a priority-only VLAN tag (one which contains the special VLAN ID 0). Set the desired PVID to make sure your untagged packets for the port get forwarded to other ports in the desired VLAN.

NOTE: *Switch management and configuration is only possible through the port if the PVID is set to 1 (the default). Setting the PVID to another value prevents the Switch from being managed/configured via that port (unless the system you are using to configure the Switch can explicitly tag frames for VLAN 1, the management VLAN).*

Force: The Force option is not usually configured. When this is checked, the PVID is forced on all frames coming in this port regardless of any existing tag.

Type: The port type controls how tags are handled on frames exiting this port.

- **Network:** This is a Trunking port that connects to another switch. All frames exiting this port will be tagged. If no tag was present when the frame entered the switch, the source port's PVID will be used. Typically, a Network port will be a member of many or all tag-based LANs on a switch and is used to forward VLAN traffic to another switch which then distributes it to other network segments based on the tags. A Network port can only send packets for VLANs in which it is a member.
- **Edge:** This is an Access port that typically connects to an end device or perhaps an unmanaged switch. No frames exiting this port will be tagged. (Use this setting for ports leading to legacy or end devices without VLAN support.)
- **Transparent**: Transparent is a useful setting for ISP use. Ordinarily, only the Edge and Network port types are configured in a network. At a Transparent type port, the existing tag is not stripped

from a frame, but a tag is still added if the port has a PVID other than 1. So when the tag is ultimately stripped at its destination, the original tag remains. If no VLANs are set (all in the default VLAN, which is untagged) and all traffic is untagged coming through, frames will be forwarded unchanged.

VLAN with RSTP

Extra care must be taken when enabling both VLANs and redundancy, or communications failures may occur.

The example shown in the following diagram depicts the problem with running the Rapid Spanning Tree Protocol (RSTP) and VLANs at the same time. The IEEE 802.1D based RSTP is not aware of the VLAN configuration. Therefore, in the example, one of the Network Ports for VLAN 3 is being blocked (see VLAN Port Settings topic in this section about Network type ports). This prevents VLAN 3 from being able to forward data to all its members.

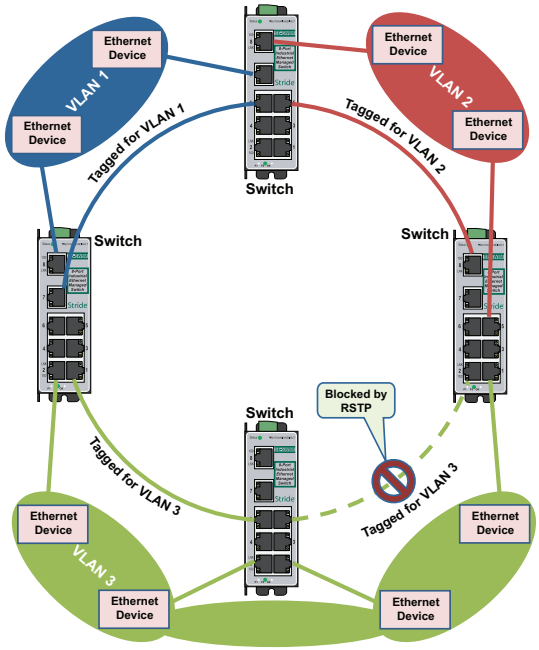

The solution to the problem above is to configure all “Network” type ports to carry all VLANs in the network. In other words, the Network Port should be a member of all VLANs defined in the switch. As seen from the example shown in the following diagram, VLAN 3 can forward to all its members through the other Network Port connections and is not affected by the block RSTP connection.

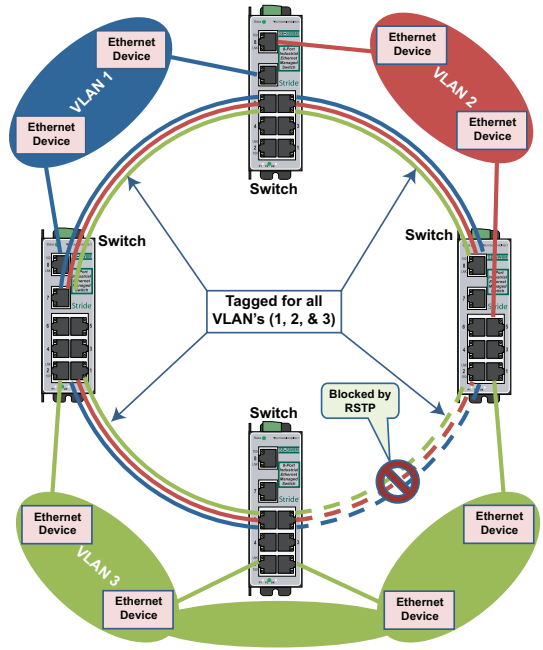

VLAN Examples

Shown below are two examples of using VLANs and how they can solve common network problems found in a factory automation facility. Note that the end devices used in these examples do not recognize nor originate VLAN tags.

Problem #1: The process requires a PLC, Remote I/O, Frequency Drive control, HMI access as well as a PC for Data Logging and a PC for configuration management. The Remote I/O device and drive communicate via Multicast and Broadcast messaging which an unmanaged switch cannot filter out. The PLC and the Remote I/O and Drive are remotely located from each other. Running multiple Ethernet connections would be costly and logistically complex so the customer wants to utilize existing wiring connections.

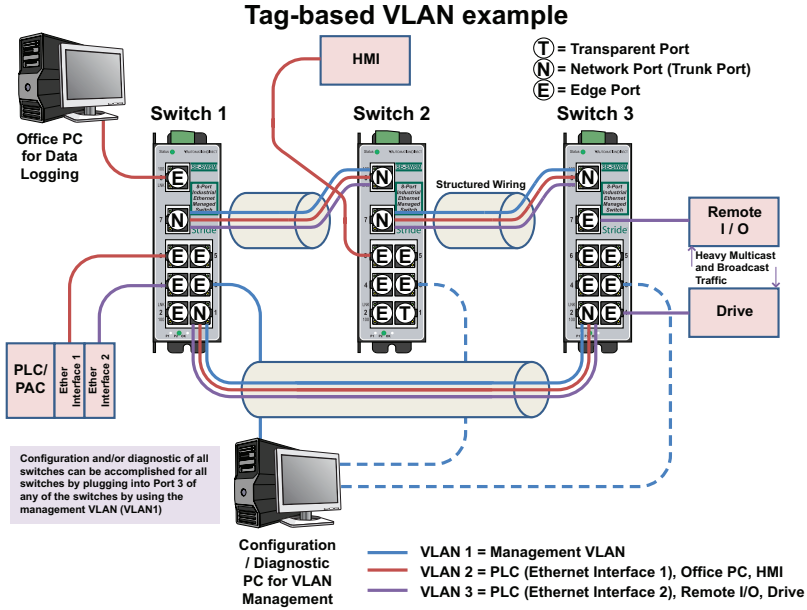

Solution: Use Stride managed switches, utilizing the VLAN feature to separate the broadcast and multicast traffic from all the devices except for the PLC. We will also wire the three switches into a Ring configuration so that we can take advantage of the redundancy feature of the switch. In this situation, we need to use Tag-based VLANs since the Ethernet packets will be traversing across multiple switches.

How to configure this setup:

We created 3 VLANs:

- VLAN 1 is the default VLAN and we leave it there and enable it on what we will call a 'management port' for each switch. In this way, we can plug our laptop into the management port of any switch and be able to access the other switches across this VLAN to tweak the configuration or view the diagnostics.
- VLAN 2 will contain one of the Ethernet interfaces of the PLC, the HMI and the Office PC.
- VLAN 3 will contain the other Ethernet interface of the PLC, the Remote I/O drop and the Drive.

Switch Setup:

Switch1:

Switch 2:

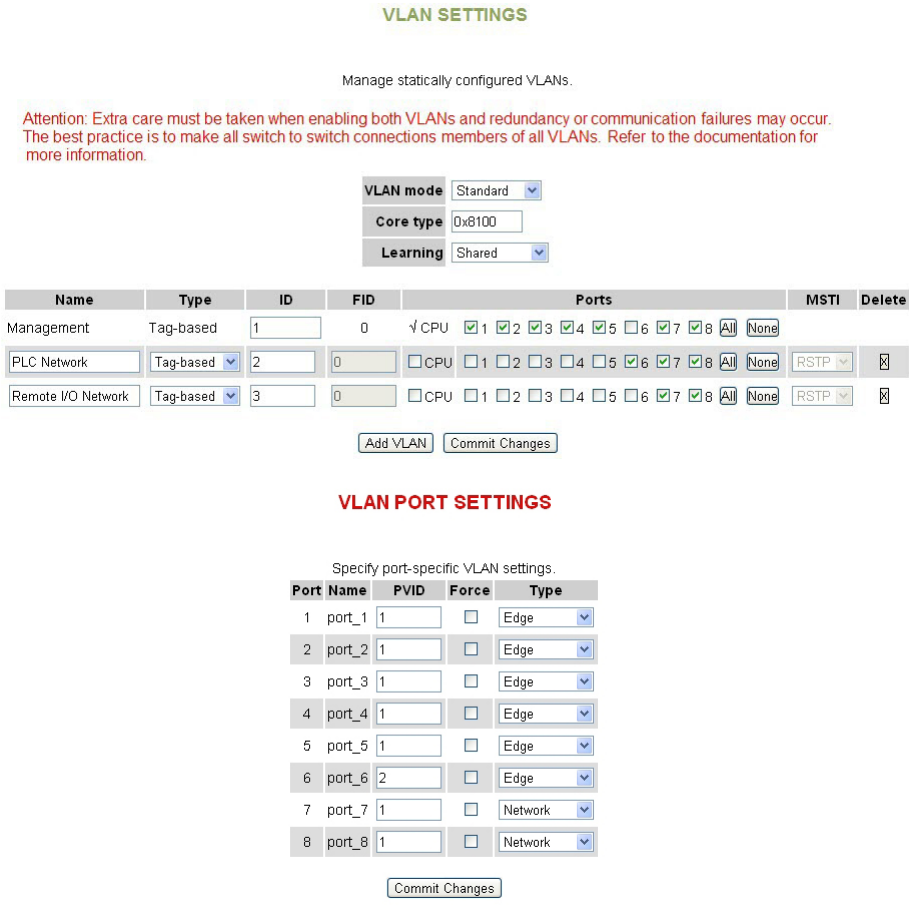
VLAN SETTINGS
Manage statically configured VLANs.
Attention: Extra care must be taken when enabling both VLANs and redundancy or communication failures may occur. The best practice is to make all switch to switch connections members of all VLANs. Refer to the documentation for more information.
VLAN mode Standard
Core type 0x8100
Learning Shared
Name Type ID FID Ports MSTI Delete
Management Tag-based 1 0 √CPU 1 2 3 4 5 6 7 8 All None
PLC Network Tag-based 2 0 CPU 1 2 3 4 5 6 7 8 All None RSTP
Remote I/O Network Tag-based 3 0 CPU 1 2 3 4 5 6 7 8 All None RSTP
Add VLAN
Commit Changes
VLAN PORT SETTINGS
Specify port-specific VLAN settings.
Port Name PVID Force Type
1 port_1 1 Edge
2 port_2 1 Edge
3 port_3 1 Edge
4 port_4 1 Edge
5 port_5 1 Edge
6 port_6 2 Edge
7 port_7 1 Network
8 port_8 1 Network
Commit Changes

Switch 3:


VLAN SETTINGS
Manage statically configured VLANs.
Attention: Extra care must be taken when enabling both VLANs and redundancy or communication failures may occur. The best practice is to make all switch to switch connections members of all VLANs. Refer to the documentation for more information.
VLAN mode Standard
Core type 0x8100
Learning Shared
Name Type ID FID Ports MSTI Delete
Management Tag-based 1 0 √CPU 1 2 3 4 5 6 7 8 All None
PLC Network Tag-based 2 0 CPU 1 2 3 4 5 6 7 8 All None RSTP
Remote I/O Network Tag-based 3 0 CPU 1 2 3 4 5 6 7 8 All None RSTP
Add VLAN
Commit Changes
VLAN PORT SETTINGS
Specify port-specific VLAN settings.
Port Name PVID Force Type
1 port_1 3 Edge
2 port_2 1 Network
3 port_3 1 Edge
4 port_4 1 Edge
5 port_5 1 Edge
6 port_6 1 Edge
7 port_7 3 Edge
8 port_8 1 Network
Commit Changes

Problem #2: This scenario is very similar to the first. We have the same problem to solve but the logistics are simpler, in that all of the devices are local and can be wired into the same switch.

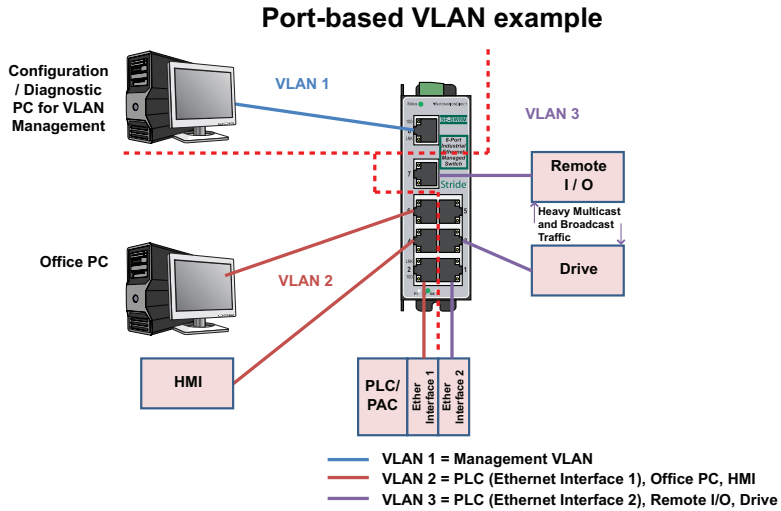

Solution: We will use a Stride managed switch, utilizing the Port-based VLAN feature. The question could be posed, "Why not just use two unmanaged switches?" While this would work, the customer wants to use as few components in the system as possible to minimize points for possible equipment faults and he would like the enhanced diagnostic capabilities that a managed switch provides.

Switch Setup:

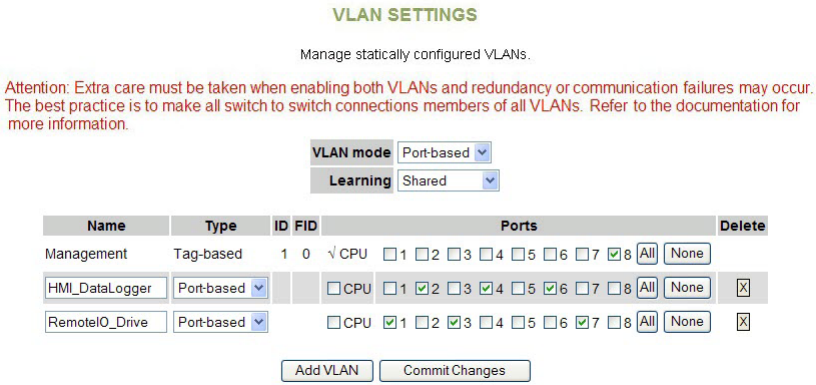

When using port-based VLANs, VLAN tags don't determine which VLAN a port is in so it is not necessary to configure the ports.

Security Settings

The managed switch offers several ways to secure access to the management functions. It can be remotely managed (monitored and configured) via the following methods:

- **Telnet:** This accesses the terminal or CLI interface (same as you would get through the console serial port) but over the Ethernet network. This type of access offers only password protection (authentication) but no encryption.
- **SSH:** Secure Shell, like Telnet, accesses the terminal or CLI interface over the Ethernet network. It offers both password protection and encryption.
- **SNMP/SNMPv3:** This method accesses the Management Information Bases (MIBs) using an SNMP server or master utility. Standard SNMPv1 or SNMPv2 has password security. SNMPv3 adds encryption.
- **HTTP/HTTPs:** This method accesses the web interface. Standard HTTP has password security. The more secure HTTPS adds encryption through SSL (Secure Socket Layers) or TLS (Transport Layer Security).

NOTE: The best security policy is to turn off or disable any access methods that you are not using.

Remote Access Security

See the “Remote Access Security” selection under the “Main Settings”.

Port Security Enables and Port Security MAC Entries

NOTE: *This feature is not supported in the 5-port models*

Port Security Enables and Port Security MAC Entries settings must be used in conjunction with one another.

The Port Security feature will drop packets from devices that are NOT entered in to the Port Security MAC Entries table. The security can be enabled for each port individually. The "Global Port Security Enable" selection must be enabled for the switch to start using the MAC Entries table.

First, on the Port Security MAC Entries page, create the table of MAC addresses allowed on each port and enter Commit Changes.

- The MAC address must be entered in the format 12:34:56:78:9A:BC.
- If a MAC address is configured to be allowed on one port AND that port is enabled on the Port Security Enables page, that MAC address is disallowed access on any other port, including ports for which security is not enabled on the Security Enables page. For example: If the MAC address for Device A has been configured for Port 1 in the MAC Entries table and Device A is plugged in to Port 5, the messages for Device A will be dropped even if Port 5 does not have security enabled.
- More than one MAC address may be configured for a port.
- A MAC address may be configured for only one port.
- If no MAC addresses are entered on the Port Security MAC Entries page AND that port is enabled on the Port Security Enables page, the port is effectively shut down and all packets will be dropped at that port.

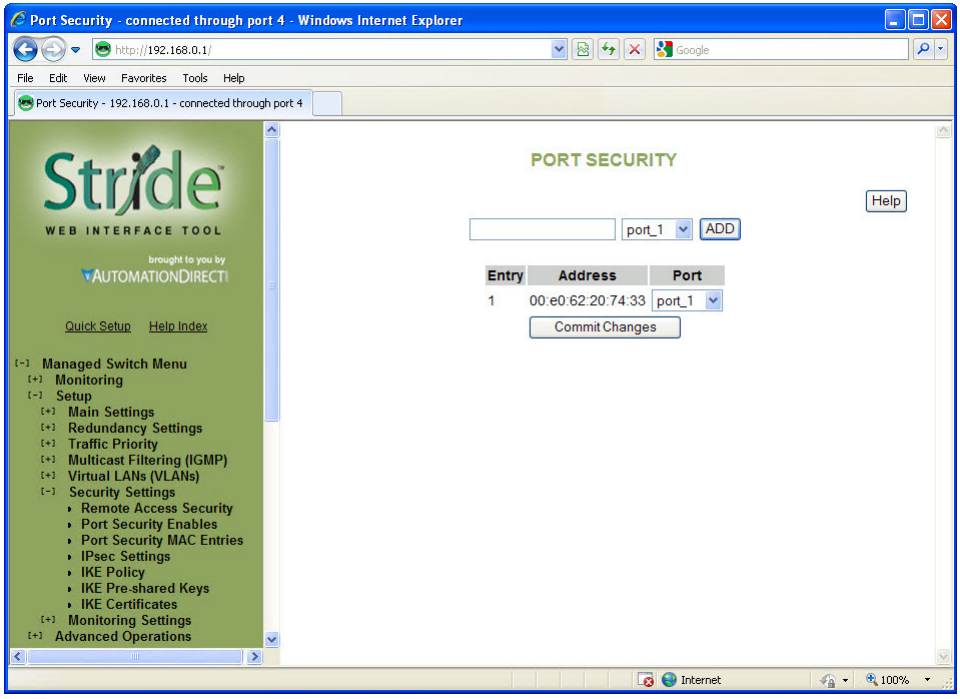

Second, to enable the MAC address security for the ports configured, select the ports and the Global Port Security selection on the "Port Security Enables" page.

Finally enter Commit Changes to write the configuration to the switch. The switch will then begin limiting access according to the configuration on these two pages.

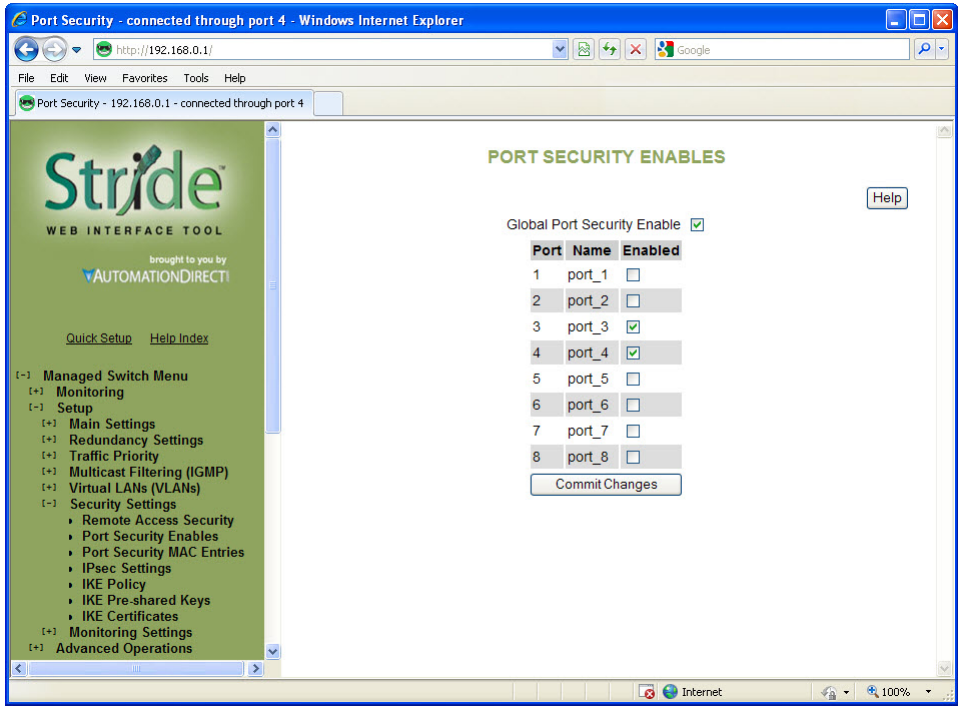

Once an entry has been configured and committed to the switch, a power cycle will be necessary after deletion of an entry in order for that security to be removed.

IPsec Settings

IPsec can authenticate, encrypt or compress IPv6 traffic to or from a switch. The IPsec software in this switch only affects management traffic addressed to or sent from the switch.

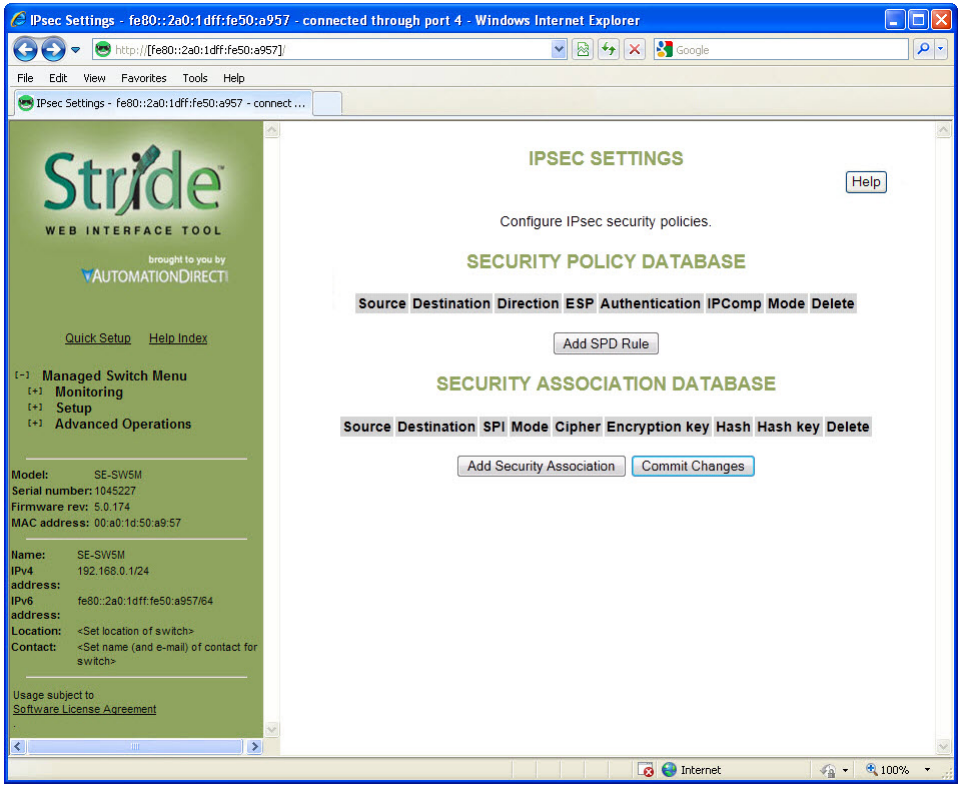

NOTE: *IPsec can only be used when the Switch's primary access address is configured with an IPv6 address. To connect to the switch via IPv6 with Internet Explorer, you must surround the address with http://[...]. Example: http://[fe80:0000:0000:0000:02a0:1dff:fe50:bfca]*

Warning: Misconfiguration on this screen may block network access to the Switch's configuration interface.

Configuration is done via two databases. The SPD (Security Policy Database) sets the required IPsec protocols for traffic going from or to configured hosts or networks. The SAD (Security Association Database) contains the encryption, compression and hash parameters needed to implement the policies required by the SPD for traffic between specific hosts.

The AH IPsec protocol is used for authentication. It uses cryptography to detect that the sender has the same hash key the receiver does. It does not provide any secrecy in transit.

The ESP protocol is used for encryption. It uses cryptography to hide the contents of traffic in transit from anyone who does not have the secret key it was encrypted with.

IPComp is used to compress traffic. It does not provide any secrecy or authenticity guarantees.

Security Policy Database: This section is used to create, delete, and modify SPD entries.

CAUTION: Take care when configuring SPD entries. If you do not configure appropriate SAD entries to go along with them and an SPD entry affects the host you are using to configure the Switch, you may find yourself unable to communicate with the Switch

To create an SPD entry, click "Add SPD Rule" and set the source, destination, direction, and protocol requirements as appropriate. To save your changes, click Commit Changes.

To delete an SPD entry, click the 'X' button at the end of the row and click Commit Changes.

To modify an SPD entry, change parameters as desired and click Commit Changes.

NOTE: *SPD entries will not apply to ICMPv6 Neighbor Discovery traffic. This allows Neighbor Discovery to function together with IKE. (Internally, the system adds high-priority rules bypassing IPsec for Neighbor Advertisement and Neighbor Solicitation packets.)*

- **Source:** An address of the form address, address/prefixlen, address/prefixlen[port], or address[port]. This specifies the source host or hosts that this policy will affect.
- **Destination:** An address in one of the same forms accepted by the Source field. This specifies the destination host or hosts that this policy will affect.
- **Direction:** The direction traffic is traveling through the switch. If the switch's address is specified in the source field, the direction should be Out. If the switch's address is in the destination field, the direction should be In.
- **ESP:** Whether to require encryption for communication between the specified hosts.
- **Authentication (AH):** Whether to require authentication for communication between the specified hosts.
- **IPComp:** Whether to require compression for communication between the specified hosts.
- **Delete:** When the button is clicked, this SPD entry will be deleted when changes are committed.

Security Association Database:

CAUTION: Take care when configuring SAD entries. If the keys and SPI values are not the same on two communicating hosts and their security policies require encryption or authentication they will be unable to successfully communicate. You may find yourself unable to communicate with the Switch.

To create an SAD entry, click "Add Security Association" and set the source, destination, SPI, mode, cipher, hash algorithm, and keys as appropriate. To save your changes, click Commit Changes.

To delete an SAD entry, click the 'X' button at the end of the row and click Commit Changes.

To modify an SAD entry, change parameters as desired and click Commit Changes.

- **Source:** An address of the form address or address[port]. This specifies the source host (and optionally port) for the security association.
- **Destination:** An address of the form address or address[port]. This specifies the destination host (and optionally port) for the security association.
- **SPI:** A locally unique value identifying this security association. This is assigned locally and may be specified in hex or decimal formats. This should be at least 0x100 (256 decimal) and must be the same on both peers in an association.
- **Mode:** The IPsec mode to use: ESP, AH, ESP and AH, or IPComp.
- **Cipher:** The cipher to use when an ESP mode is selected.
- **Encryption key:** The key to use when ESP is enabled. This must be specified in hexadecimal (beginning with 0x) and should be 24 bytes (48 digits) long for 3DES or 16, 24 or 32 bytes (32, 48, or 64 digits) long for AES.
- **Hash:** The hash algorithm to use when an AH mode is selected. MD5 is not recommended.
- **Hash key:** The hash key to use when AH is enabled. This must be specified in hexadecimal (beginning with 0x) and should be 20 bytes (40 digits) long for SHA1 or 32 bytes (64 digits) long for SHA256.
- **Delete:** When the button is clicked, this SAD entry will be deleted when changes are committed.

IKE Policy

This screen allows you to configure IKE policy for auto negotiating IPsec Security Associations over IPv6.

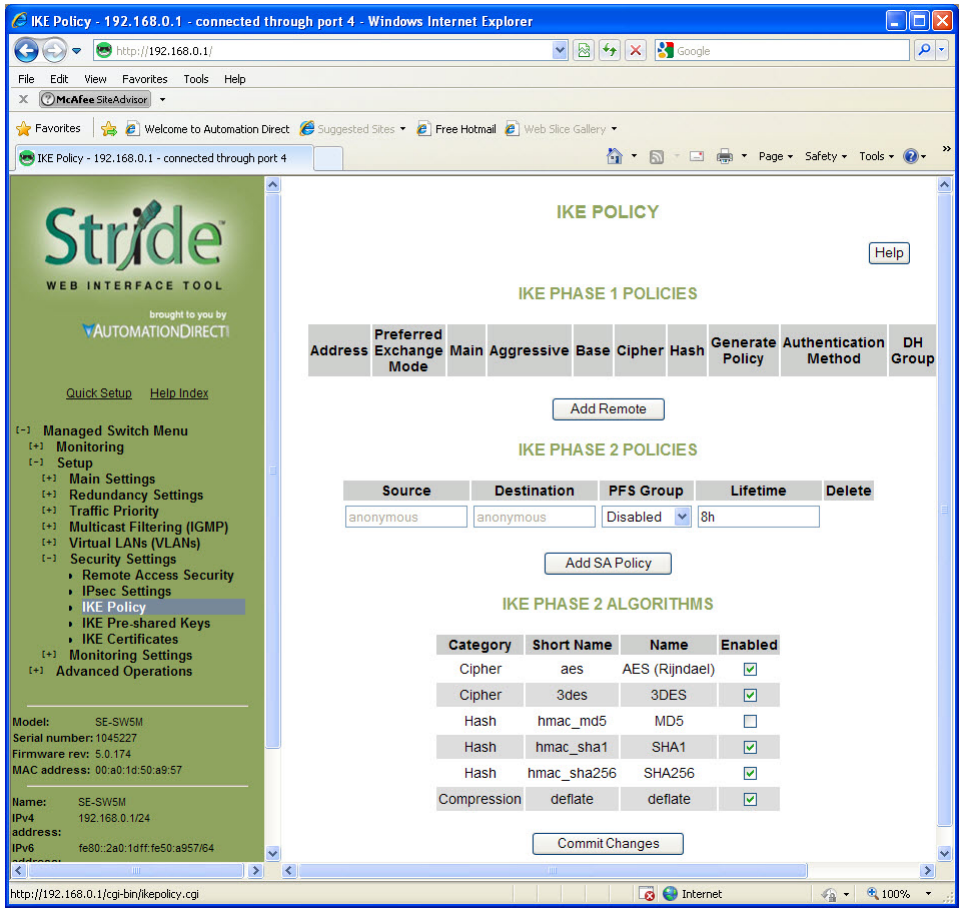

Warning: Misconfiguration on this screen may block network access to the Switch's configuration interface.

IKE Phase 1 Policies: This section may be used to create, delete, and modify ISAKMP (IKE phase 1) policies. Phase 1 is used to securely authenticate peers.

- **Address:** The address of the peer the policy will apply to. A policy for "anonymous" will apply to all peers without a more specific policy.
- **Preferred Exchange Mode:** The preferred exchange mode is the one that will be sent in any proposal to a peer. If other exchange modes are specified, they will be accepted in received proposals. With Aggressive, the DH Group in the sent proposal must exactly match the peer's configuration.

- **Cipher:** The cipher used to encrypt proposal exchanges. You must choose a cipher.
- **Hash:** The hash used to authenticate proposal exchanges. You must choose a hash algorithm.
- **DH Group:** The Diffie-Hellman group used for exponentiations. Larger groups should be more secure, but may take so long to compute that completing negotiation becomes impossible due to timeouts, preventing connectivity to the switch management interface. This should generally be set to the same value on both peers in a connection.

IKE Phase 2 Policies: This section, together with IKE Phase 2 Algorithms is used to configure the parameters used to establish Security Associations between peers once they have authenticated each other in phase 1.

The policy to use is selected using the source and destination selectors from the Security Policy Database entry or the ID payload from the received IKE packet which triggered the negotiation. The match for any values other than "anonymous" must be exact.

- **Source:** The source address to match against. The address specified should exactly match the Destination address field in a phase 2 policy on the peer, unless either value is "anonymous". The value "anonymous" matches sources not handled by other rules.
- **Destination:** The destination address to match against. The address specified should exactly match the Source address field in a phase 2 policy on the peer, unless either value is "anonymous". The value "anonymous" matches the destinations not handled by other rules.
- **PFS Group:** The Diffie-Hellman exponentiation group used for Perfect Forward Secrecy. This may be disabled if not required, but any proposal suggesting it will still be accepted. Larger groups may require an excessive amount of processing time during negotiation, causing timeouts.

IKE Phase 2 Algorithms: This section is used to configure the algorithms which may be used for phase 2. The exact algorithms chosen will be an intersection between the sets specified here and on a peer.

You must enable at least one algorithm from each category (cipher, hash, and compression), even if the switch's IPsec policies do not require one of the given protocols to be used.

The default values should be compatible with most installations.

AES (default = Enabled) Cipher

3DES (default = Enabled) Cipher

SHA1 (default = Enabled) Hash

SHA256 (default = Enabled) Hash

MD5 (default = Disabled) Hash MD5 is known to be insecure and is included only for compatibility with old implementations.

Deflate (default = Enabled) Compression

IKE Pre-shared Keys

This screen allows you to configure IKE PSKs (pre-shared keys) used to negotiate with the IKE peers with which the switch communicates over IPv6.

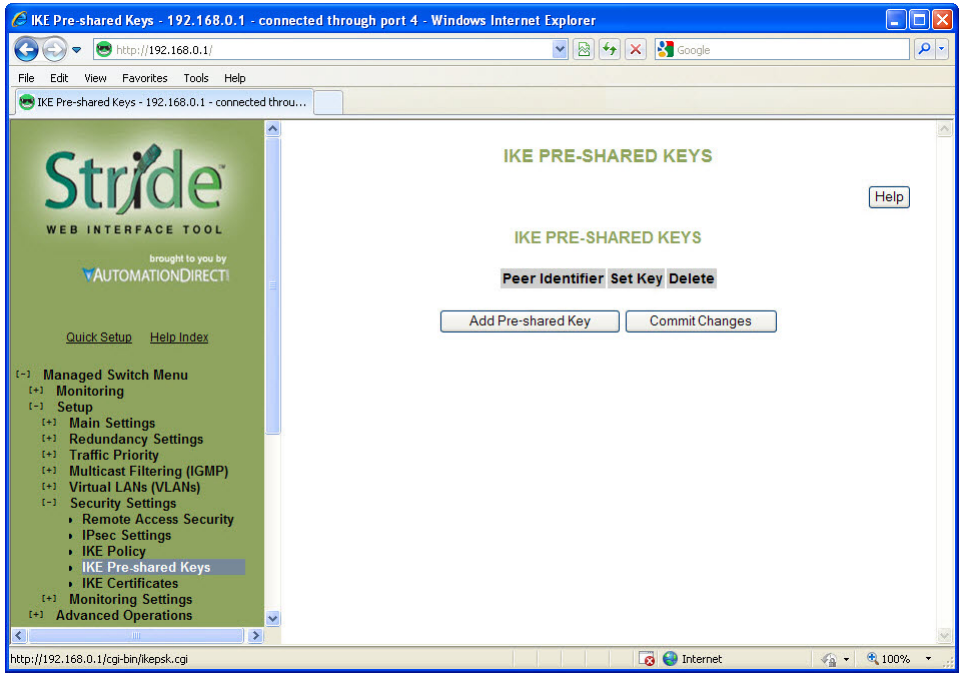

Warning: Misconfiguration on this screen may block network access to the Switch's configuration interface.

The same pre-shared key must be set for both peers. For example, if communicating between two hosts fe80::1 and fe80::2 with a pre-shared key "secret", fe80::1 must have "secret" set as the pre-shared key for peer fe80::2, and fe80::2 must have "secret" set as the pre-shared key for peer fe80::1.

- **Peer Identifier:** The identifier of the peer with which this pre-shared key should be used. Typically this will be the peer's address.
- **Set Key:** The value to set the pre-shared key to. If left blank, the current value will be preserved.
- **Delete:** Mark this pre-shared key for removal when changes are committed.

IKE Certificates

This screen allows you to configure IKE certificates used to identify the switch and IKE peers with which it communicates over IPv6.

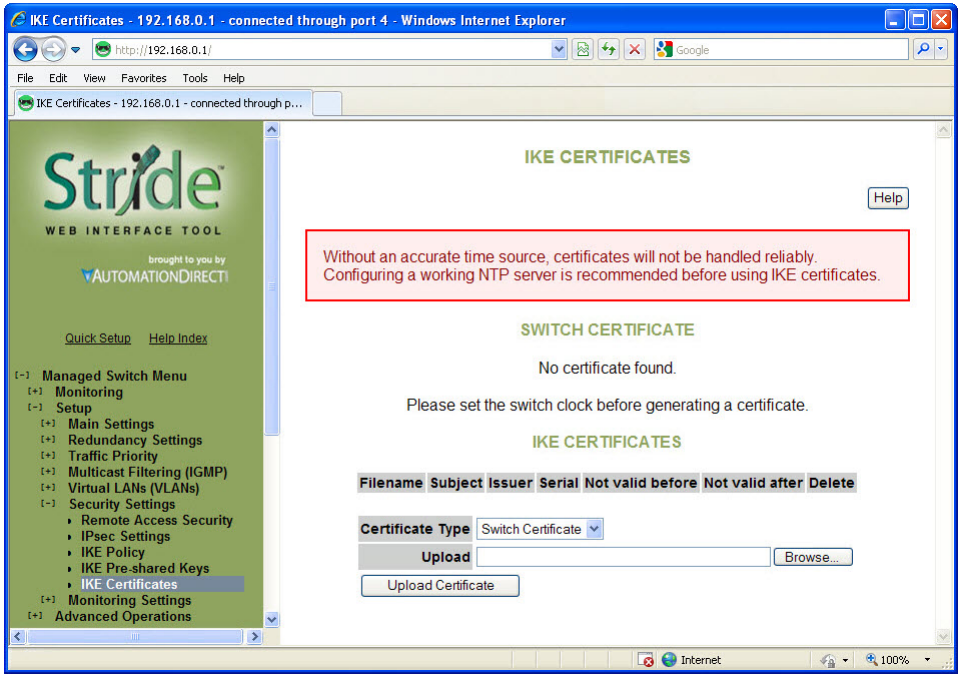

Warning: Misconfiguration on this screen may block network access to the Switch's configuration interface.

Providing a reliable time source, such as NTP, is highly recommended, as IKE will reject certificates which are not valid according to the system time, whether it is before the 'not valid before' time or after the expiration time. If NTP is used, pre-shared keys or hard-wired Security Associations should be used for IPsec communications with the NTP server or updating the clock will fail.

The HTTPS certificate used by the switch's Web interface cannot be changed on this screen.

Switch Certificate: This section may be used to generate or view the details of an X.509 certificate which the switch uses to identify itself via IKE.

A certificate request which can be provided to a third-party Certificate Authority (CA) is also generated. A CA-signed certificate can be uploaded using the form at the bottom of the page and will replace the self-signed certificate used by the switch for IKE. Note that the certificate provided should be generated from the certificate request generated by the switch.

- **Subject:** The DN (distinguished name) identifying the holder of the certificate.
- **Issuer:** The DN (distinguished name) identifying the issuer of the certificate.
- **Serial:** The certificate's serial number.
- **Certificate:** A link which can be used to download the certificate for inspection.
- **Request:** A link which can be used to download a certificate request to be signed by a CA.
- **Not valid before:** The earliest time for which the certificate is valid.
- **Not valid after:** The latest time for which the certificate is valid.
- **Delete:** Pressing this button will delete the certificate and private key, allowing a new one to be generated. This operation cannot be undone.

When no IKE certificate is present on the switch, a certificate and key may be generated. The following options may be set.

- **Common Name:** The CN to use as the subject of the new certificate. This should identify the switch and is typically a hostname or IP address. It defaults to the switch's hostname.
- **Bits:** The size of the private key to create, in bits.
- **Expires:** The number of days the certificate will be valid for, starting from the current day according to the switch's clock. This setting is used only for the self-signed certificate; CAs provides their own expiration dates for certificates they produce.

IKE Certificate: This section is used to add, delete, and view certificates which are trusted by the switch during IKE negotiation.

- **Subject:** The DN (distinguished name) identifying the holder of the certificate.
- **Issuer:** The DN (distinguished name) identifying the issuer of the certificate.
- **Serial:** The certificate's serial number.
- **Not valid before:** The earliest time for which the certificate is valid.
- **Not valid after:** The latest time for which the certificate is valid.
- **Delete:** Pressing this button will delete the certificate.

Certificates can be added to the switch using the upload form.

- **Certificate Type:** Whether the uploaded certificate is to be used as the switch's identity ("Switch Certificate"), or to be added to the certificates trusted by the switch when negotiating with IKE peers ("CA Certificate"). The CA Certificate option may also be used to trust self-signed certificates from peers.
- **Upload:** The certificate to upload.

Monitoring Settings

Alarm (OK) Output

These settings control the events that will trigger the alarm output. The OK discrete output is on during normal conditions and turned off in the event of an alarm condition.

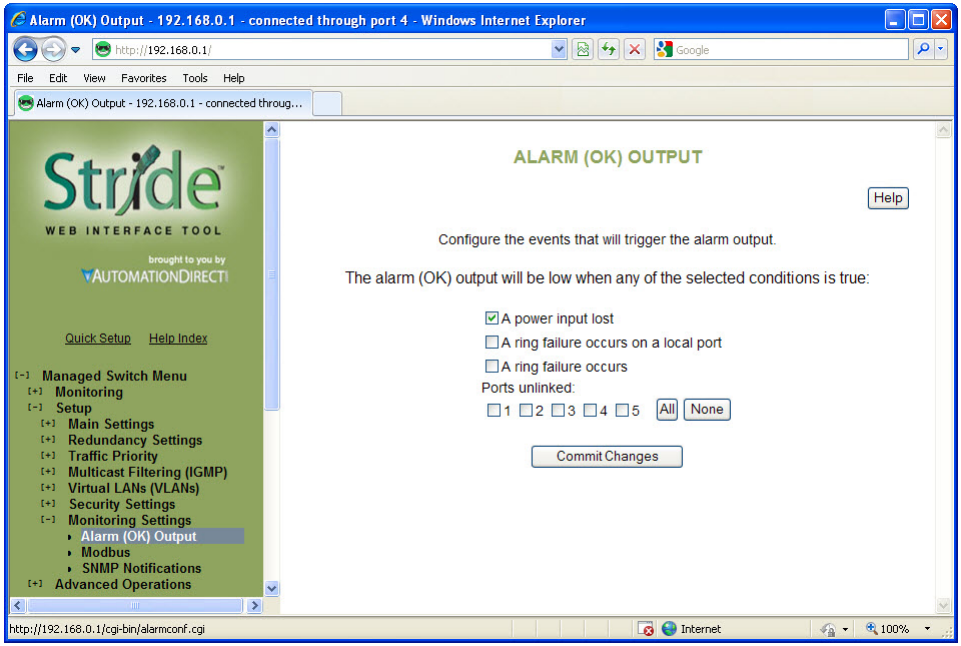

Both Power Inputs On: An alarm condition will be triggered if power is not on for both power inputs.

Ring Failure: An alarm condition will be triggered when a ring failure occurs.

Ring failure on a local port will be triggered when one of this switch's neighbors in the ring goes down; the general ring failure option will be triggered when any switch in the ring goes down.

The general ring failure option implies that local ring port failure is also detected.

Ports Linked: An alarm condition will be triggered whenever any of the selected ports are not linked.

Modbus

These settings control whether and how the switch will respond to Modbus requests. Modbus registers are available for monitoring link status on each Ethernet port, the power and OK status, and the status of each configured Real-Time Ring.

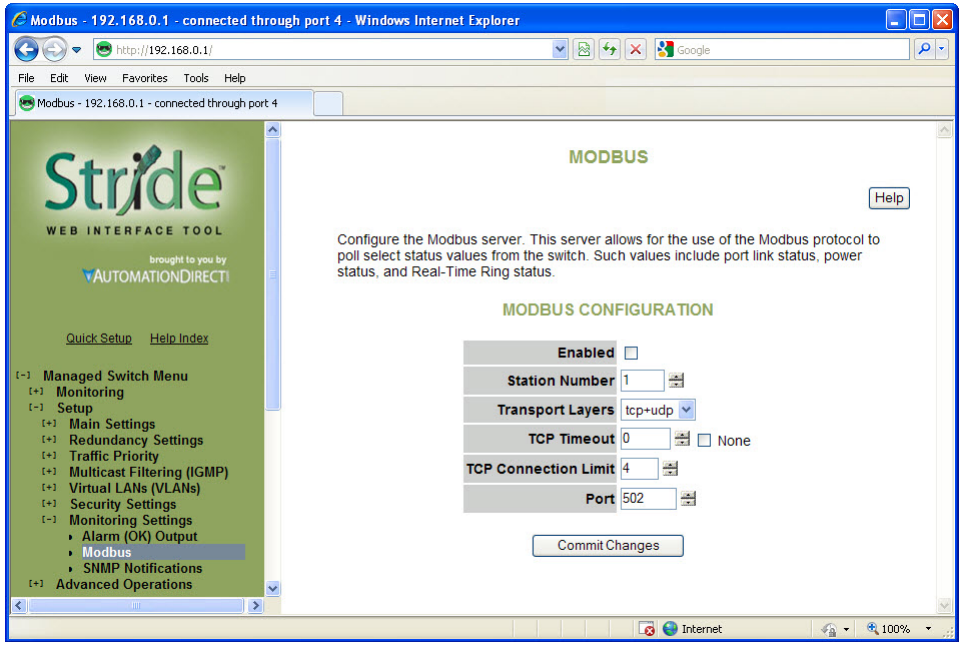

Enabled: If selected, the switch will respond to Modbus requests.

Station Number: The Modbus station number that the switch will respond as.

Transport Layers: The switch will respond to Modbus requests only on the chosen transport layers.

TCP Timeout: If a new TCP connection is received when there are no more free connections (see the TCP Connection Limit), this determines what happens:

0: The least recently active connection will be dropped in favor of the new connection.

>0: The least recently active connection will be dropped in favor of the new connection, but only if the least recently active connection has been inactive for at least this many seconds.

None: The new connection will be dropped immediately after it is accepted.

TCP Connection Limit: The maximum number of active TCP connections that the Modbus server will maintain. Above this limit, the TCP Timeout value will be used to decide how new connections should be handled.

Port: The TCP/UDP port number on which to listen for new connections/requests.

Register Mapping:

The Modbus registers (all discrete inputs) that may be polled for switch status are:

Link Status for Ports 1-16:

10001 Link status of port 1 (1 = link present, 0 = no link present)

10002 Link status of port 2

...10016 Link status of port (register - 10000)

Real-Time Ring Status for Rings 1-4:

10017 Ring 1: Ring is complete (1 = complete, 0 = broken)

10018 Ring 1: First port is passing data (1 = active, 0 = blocked)

10019 Ring 1: Second port is passing data (1 = active, 0 = blocked)

10020 Ring 2: Ring is complete

10021 Ring 2: First port is passing data

10022 Ring 2: Second port is passing data

10023 Ring 3: Ring is complete

10024 Ring 3: First port is passing data

10025 Ring 3: Second port is passing data

10026 Ring 4: Ring is complete

10027 Ring 4: First port is passing data

10028 Ring 4: Second port is passing data

Switch Status:

10030 OK output (1 = on/no alarm, 0 = off/alarm)

10031 First power input active (1 = P1 on, 0 = P1 off)

10032 Second power input active (1 = P2 on, 0 = P2 off)

Extended Link Status for Ports 1-99:

10101 Link status of port 1 (1 = link present, 0 = no link present)

10102 Link status of port 2

10199 Link status of port (register - 10100)

Extended Switch Status:

10300 OK output (1 = on/no alarm, 0 = off/alarm)

10301 First power input active (1 = P1 on, 0 = P1 off)

10302 Second power input active (1 = P2 on, 0 = P2 off)

SNMP Notifications

SNMP (Simple Network Management Protocol) and RMON (Remote Monitoring) provide a means to monitor and manage your network. Each SNMP device maintains Management Information Bases (MIBs) containing information about the operation and configuration of the device.

NOTE: This product uses Net-SNMP (available from www.net-snmp.org) which is subject to the copyrights and license found at: http://www.net-snmp.org/COPYING.txt

The MIBs can be accessed with SNMP tools ranging from simple command-line tools like snmpwalk and snmpget (part of the open source Net-SNMP package available at http://www.net-snmp.org) to commercial network management products from various vendors. Key information from the MIBs is also available via the switch's terminal and web interfaces.

The MIBs are divided into groups of related objects. Objects may be scalar (having only a single value) or tabular (having a list of values varying over time, by port number, etc.).

SNMP Security:

SNMP provides several options for securing access to MIBs. SNMPv1 and SNMPv2 provide only weak authentication. SNMPv3 uses encryption to add stronger authentication as well as privacy. In all versions, you may configure read-only and read/write users.

SNMPv1 and SNMPv2 authenticate users with a "community string" which is sent in clear text (unencrypted) and no password is required. Some measure of security can be achieved by setting long, obscure community strings.

SNMPv3 provides three levels of security and encryption:

- **None:** No password is required to read or write values in the MIB.
- **Authentication:** A password is required and is used to encrypt the user credentials so that security information is not sent in clear text. A variation of MD5 is used for encryption.
- **Privacy:** A password is required and is used to encrypt the user credentials. A second password is used to encrypt the details of the SNMP request using DES encryption.

For SNMPv3 access, the managed switch requires authentication and allows privacy. Only one password is configurable and it is used for both authentication and privacy.

The following examples use snmpget from the Net-SNMP tools to illustrate the use of authentication and privacy when accessing the managed switch.

If SNMPv2 access is enabled, values may be read without a password with a command like:

```
snmpget -v 2c -c public 10.2.0.1 system.sysDescr.0
```

If SNMPv3 access is enabled, values may be read with a command like the following (entered all on one line):

```
snmpget -v 3 -u public -l authNopriv -a MD5 -A publicpwd 10.2.0.1 system.sysDescr.0
```

Finally, if SNMPv3 access is enabled, an authenticated, private request could be made with a command like the following:

```
snmpget -v 3 -u public -l authpriv -a MD5 -A publicpwd -x DES -X publicpwd 10.2.0.1
system.sysDescr.0
```

The switch supports SNMPv1, v2, and v3. SNMPv1 and v2 access are essentially the same from a security standpoint and are enabled and disabled together. SNMPv3 security may be separately controlled. Thus you may prevent unauthenticated access to your switch by disabling SNMPv1/v2 access entirely while retaining password-secured access via SNMPv3.

SNMP Notifications:

Use the SNMP Notifications Menu to enable traps to be sent when the state of the switch changes. Access this menu by selecting Setup from the Main Menu, and then selecting Main Settings.

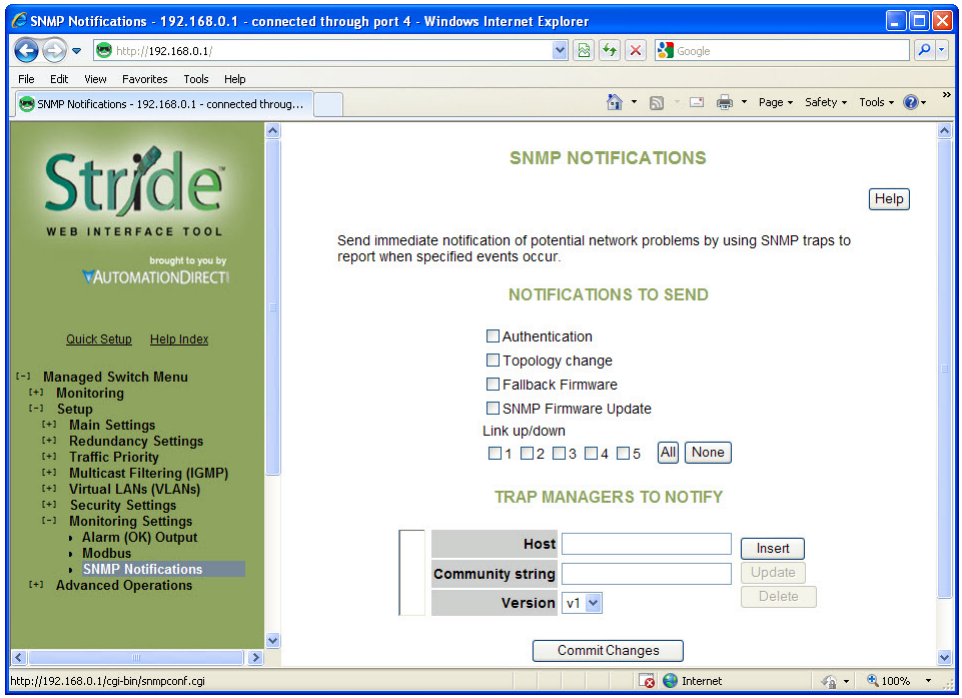

Use the SNMP Notifications Menu to enable traps to be sent when the state of the switch changes. Access this menu by selecting Setup from the Main Menu, and then selecting Main Settings.

- **Authentication:** Traps can be sent when invalid credentials (such as an unrecognized community string) are presented to the SNMP agent. Enable this setting to generate authentication traps.
- **Topology change:** Traps can be sent when the topology of the spanning tree changes. Enable this setting to generate topology change traps.
- **Failback Firmware:** Check this box to send a trap when the switch resets into the non-default firmware image. This can happen if the switch loses power while booting, or if the default firmware image somehow becomes corrupt and is no longer bootable.
- **SNMP Firmware Update:** Check this box to send a trap when the switch has completed an SNMP-initiated firmware update. The trap will trigger regardless of whether the firmware update succeeded. Check the firmware Health entry in the firmware Table over SNMP to determine whether the update was successful. If it lists the non-running image as Healthy (1), then the update succeeded. Otherwise, it failed.
- **Link 1 up/down-Link 18 up/down:** Traps can be sent when a link goes up or down (the same state reflected in the LED for each port). Enable these settings to generate link up/down traps.

Trap Managers to Notify:

Use the Trap Managers Menu to specify where traps will be sent. The Trap Managers Menu can be accessed by selecting Setup from the Main Menu and then selecting Main Settings. Up to five trap managers may be configured. For each one, the following values may be specified.

- **Host:** The IP address of the host where the trap manager is located.
- **Community String:** The community string to use when contacting the trap manager on the host.
- **Version:** The SNMP trap version to send.

NOTE: There are two system traps that cannot be disabled and will be sent to any configured trap managers. A coldStart trap will be sent whenever the SNMP agent starts up (usually, this is only when the Switch is reset). A NotifyRestart trap will be sent whenever the SNMP agent's configuration changes and is reloaded. This will happen, for example, when you commit changes on a configuration menu that includes SNMP settings.