



SonicOS 5.7: Advanced Switching Feature Guide and Screencast Tutorial

This solutions document describes how to configure and manage the Switching feature on a SonicWALL NSA 2400MX running SonicOS 5.7. A screencast tutorial on Port Mirroring is also provided.

This document contains the following sections:


- “Feature Overview” section on page 2
- “Configuring Switching” section on page 5
- “Troubleshooting and Verification” section on page 38
- “Technical FAQ” section on page 44
- “Glossary” section on page 45

Screencast Tutorial - Port Mirroring

Adobe Reader and browser plugin version 9.1 or higher is required to play the embedded Port Mirroring screencast. You can download the latest version of Adobe Reader at <http://get.adobe.com/reader>. The installation will require a restart of your system. When ready, click the image below to play the screencast. Alternatively, you can access the tutorial at: <http://help.mysonicwall.com/enhancedcontent/12mirroring/>.



How to Configure Port Mirroring on an NSA 2400MX Firewall
A SonicWALL Tutorial



Feature Overview

This section provides an introduction to the Switching feature. This section contains the following subsections:

- [“What is Switching on the SonicWALL NSA 2400MX?” section on page 2](#)
- [“Benefits of Switching in SonicOS” section on page 3](#)
- [“How Does Switching Work on the SonicWALL NSA 2400MX?” section on page 4](#)
- [“Supported Platforms” section on page 4](#)

What is Switching on the SonicWALL NSA 2400MX?

SonicOS 5.7 introduces Layer 2 (data link layer) switching functionality on the SonicWALL NSA 2400MX appliance.

The SonicWALL NSA 2400MX appliance is a Unified Threat Management (UTM) security appliance that integrates the WAN flexibility of a router with 24 built-in Ethernet switch ports. The appliance provides two expansion slots to allow modular card flexibility for additional ports. Both 3G wireless cards and V.90 modem cards are supported for WAN access.

The functionality supports the following switching features:

- VLAN Trunking – Provides the ability to trunk different VLANs between multiple switches.
- Rapid Spanning Tree Protocol – Prevents loops from being formed when switches or bridges are interconnected via multiple paths and provides for network convergence after a topology change.
- Layer 2 Network Discovery – Uses IEEE 802.1AB (LLDP) and Microsoft LLTD protocols and switch forwarding table to discover devices visible from a port.
- Link Aggregation – Provides the ability to aggregate ports for increased performance and redundancy.
- Port Mirroring – Allows the administrator to assign a mirror port to mirror ingress, egress or bidirectional packets coming from a group of ports.
- Layer 2 Quality of Service – On a per port basis, allows configuration to trust Cost of Service (CoS) (802.1p) or trust DSCP marking and treat the frames appropriately.
- Rate Control / Flow Control – On a per port basis, the bandwidth of ingress frames can be tuned in four modes by limiting all/flooded unicast/multicast/broadcast frames. Rate limiting for egress frames can be enabled or disabled.
- Port Security – Provides the ability to bind a MAC address or multiple MAC addresses to a specific port interface.

Benefits of Switching in SonicOS

The SonicWALL NSA 2400MX provides a combined security and switching solution with the objective of improved security for all tasks. Layer 2 switching features enhance the deployment and interoperability of SonicWall devices within existing Layer 2 networks.

The SonicWALL NSA 2400MX provides flexible, intelligent switching capabilities with its unique PortShield architecture, increased port density with 26 interfaces, and advanced switching features.

The advanced switching features on a network security appliance provide the following benefits:

- Increased port density – With one appliance providing 26 interfaces, including 24 switch ports, you can decrease the number of devices on your internal network.
- Increased security across multiple switch ports – The PortShield architecture provides the flexibility to configure all 26 LAN switch ports into separate security zones such as LANs, WLANs and DMZs, providing protection not only from the WAN and DMZ, but also between devices inside the LAN. Effectively, each security zone has its own wire-speed 'mini-switch' that benefits from the protection of a dedicated deep packet inspection firewall.
- VLAN Trunking – Simplifies VLAN management and configuration by reducing the need to configure VLAN information on every switch.
- Layer 2 Discovery – Provides Layer 2 network information for all devices attached to the SonicWALL NSA 2400MX.
- Link Aggregation – Aggregated ports provide increased performance through load balancing when connected to a switch that supports aggregation, and provide redundancy when connected to a switch or server that supports aggregation.
- Port Security – Allows administrators to bind a trusted MAC address or multiple MAC addresses to a specific port to decrease unauthorized access on that port.
- Rapid Spanning Tree Protocol – Allows for redundancy in case a connection goes down, while preventing loops from being formed when switches or bridges are interconnected via multiple paths.
- Layer 2 Quality of Service – Allows for traffic prioritization and bandwidth management to minimize network delay using Cost of Service (CoS) classification, and DSCP marking.
- Port Mirroring – Allows the administrator to easily monitor and inspect network traffic on one or more ports.
- Rate Control / Flow Control – Back-pressure flow control on half-duplex ports and pause frame-based flow control on full-duplex ports allow zero packet loss under temporary traffic congestion.
- Port Security – Binding a MAC address or multiple MAC addresses to a specific port interface provides security, as frames whose source addresses are not contained in the table will be dropped.

How Does Switching Work on the SonicWALL NSA 2400MX?

The switching features have their own menu group in the left navigation pane of the SonicOS management interface.

Figure 1 Switching in SonicOS – Navigation Pane



Some switching features operate on PortShield Groups and require preliminary configuration on the Network > PortShield Groups page. Some operate on existing Network > Interface configurations. The Port Security feature uses MAC address objects. For more information about configuring these related features in SonicOS, see the *SonicOS 5.7 Administrator's Guide*.

For details about the operation of each switching feature, see the related section under the [“Configuring Switching”](#) section on page 5.

Supported Platforms

Switching is available on the SonicWALL NSA 2400MX running SonicOS 5.7 and higher. Switching features are only available on ports X2 - X25, not on X0 (LAN) or X1 (WAN).

The hardware design of the SonicWALL NSA 2400MX includes the following elements:

- Dual core 700 MHZ CPU
- 8 Gigabit Ethernet interfaces
- 16 10/100 Megabit Fast Ethernet interfaces
- 1 Gigabit Ethernet WAN port
- 1 Gigabit Ethernet LAN port
- 2 USB extension ports that support external 3G wireless cards or V.90 analog modem cards
- 2 Expansion Slots for future use

Configuring Switching

This section contains the following sections:

- [“Configuring VLAN Trunking” section on page 5](#)
- [“Configuring Rapid Spanning Tree” section on page 13](#)
- [“Configuring Layer 2 Discovery” section on page 18](#)
- [“Configuring Link Aggregation” section on page 21](#)
- [“Configuring Port Mirroring” section on page 25](#)
- [“Configuring Layer 2 Quality of Service” section on page 27](#)
- [“Configuring Rate Control” section on page 32](#)
- [“Configuring Port Security” section on page 35](#)

Configuring VLAN Trunking

VLAN trunking is supported by the IEEE 802.1Q networking standard, also called VLAN Tagging. This standard defines how VLANs operate with regard to Layer 2 (MAC layer) bridging. The use of VLANs and VLAN trunking allows multiple bridged networks to simultaneously share a single physical network while preserving the privacy of information in each (virtual) network. IEEE 802.1Q also refers to the encapsulation protocol used to implement this standard in Ethernet networks. The SonicWALL NSA 2400MX appliance supports 802.1Q encapsulation on its VLAN trunk ports. Encapsulation, in this case, refers to the

For example, a company, university, or other organization can use VLANs to create separate logical (virtual) networks for different departments. Each department is assigned to its own VLAN. The switch ports to which the department computers are connected are configured as members of that VLAN. When network traffic is sent out from a department computer, the switch adds a 32-bit VLAN tag to each data frame before forwarding it via a VLAN trunk port. Each switch in the network examines the VLAN tag, and uses the information to determine that it is a tagged frame, the priority level (defined by IEEE 802.1p), whether it is an Ethernet or a Token Ring frame, and the VLAN to which the frame belongs. The frame makes its way through the physical network until it reaches the last switch before the destination device, at which point the switch removes the VLAN tag and delivers the frame to its destination. This switch only delivers the frame via a port that is configured as a member of the same VLAN, thereby ensuring that the data is not leaked to any other department.

In the above scenario, the switch ports connected to department computers are configured as members of a VLAN. The switch ports that are connected to other switches in the physical network are configured as VLAN trunk ports. This distinction means that only unassigned switch ports on the SonicWALL NSA 2400MX appliance can function as VLAN trunk ports.

You can enable or disable individual VLANs on the trunk ports, allowing the existing VLANs on the SonicWALL NSA 2400MX appliance to be bridged to respective VLANs on another switch connected via the trunk port. A maximum of 32 VLANs can be enabled on each trunk port.

Figure 2 shows the Switching > VLAN Trunking page. The page displays the range of reserved VLANs in the **Reserved VLAN Information** section, details about current VLANs in the **VLAN Table**, and the VLAN trunks configured on the system in the **VLAN Trunks** area.

Figure 2 Switching > VLAN Trunking Page

The screenshot shows the 'Switching / VLAN Trunking' page. It features three main sections: 'Reserved VLAN Information', 'VLAN Table', and 'VLAN Trunks'.

Reserved VLAN Information:

Starting VLAN ID:	3767
Ending VLAN ID:	3791

VLAN Table:

VLAN ID	Interface	Member Ports	Trunked	Configure
26	X0	X2, X3, X4, X5, X6, X7, X8, X9, X10, X11, X12, X13, X18, X19, X0		
3767	X14	X15, X16, X17, X14, X20, X21		
100	X20:V100	X20, X21		
200	X20:V200	X20, X21		
3771	X22	X22		
3772	X23	X23		
3773	X24	X24		
3774	X25	X25		

VLAN Trunks:

Trunk Port	VLAN ID	Configure
<input type="checkbox"/> X20 (7 VLAN entries)		
	3767	
	100	
	200	
<input type="checkbox"/> X21 (7 VLAN entries)		
	3767	
	100	
	200	

At the bottom of the page, there are three buttons: 'Add', 'Delete', and 'Enable VLAN'.

The VLAN trunking feature provides the following functions:

- Change VLAN ID's of existing PortShield groups
- Add/delete VLAN trunk ports
- Enable/disable VLANs on the trunk ports

The allowed VLAN ID range is 1-4094. Some VLAN IDs are reserved for PortShield use and are displayed in the **Reserved VLAN Information** table on the Switching > VLAN Trunking page.

The values displayed on the Switching > VLAN Trunking page are described in [Table 1](#).

Table 1 VLAN Trunking Page Description

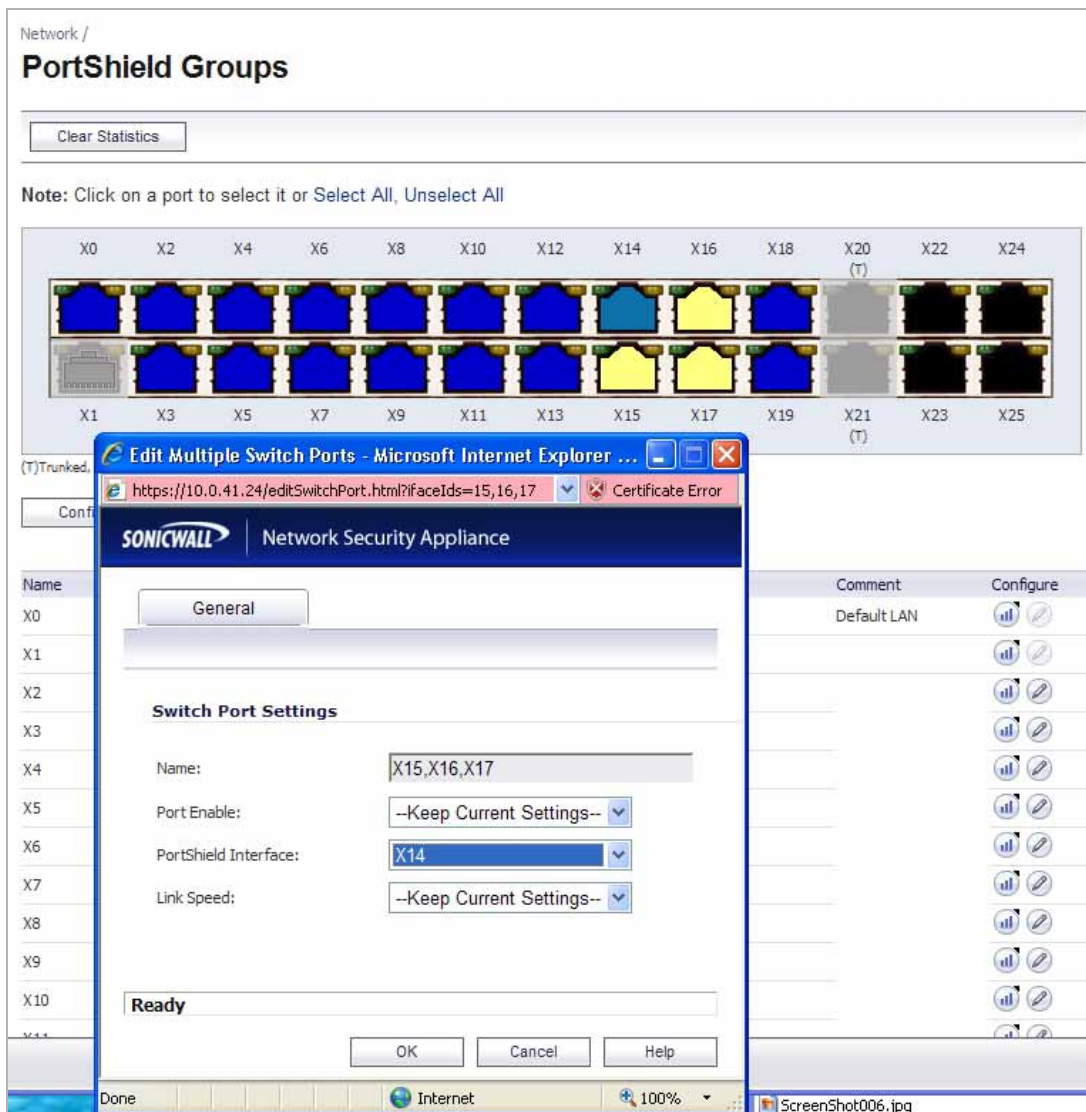
Item	Description
Reserved VLAN Information	
Starting VLAN ID	The lowest ID number in the VLAN range reserved for PortShield use. This VLAN range is reserved for local VLANs associated with a PortShield group.
Ending VLAN ID	The highest ID number in the VLAN range reserved for PortShield use.
VLAN Table	
VLAN ID	The ID number of the VLAN. VLAN ID 26 is the default VLAN that includes all LAN ports on the SonicWALL NSA 2400MX unless configured otherwise.
Interface	<ul style="list-style-type: none"> For the default VLAN that includes all LAN ports unless configured otherwise, the lowest interface, X0, is used. For other VLANs with multiple member ports, the displayed interface is the one configured as the PortShield interface for that PortShield group. For interfaces that are not assigned to a PortShield group, the port number of the interface is used, such as X25. For remote VLANs, the trunked interface and the virtual interface (called the VLAN Trunk Interface) are displayed in the format: [trunked interface]:V[virtual interface number] The virtual VLAN Trunk Interface is automatically created for remote VLANs. When the same remote VLAN is enabled on another trunk port, no new interface is created. All packets with the same VLAN tag ingressing on different trunk ports are handled by the same virtual interface.
Member Ports	<ul style="list-style-type: none"> For PortShield groups, all interfaces in the group are listed as Member Ports. For interfaces that are not assigned to a PortShield group, only the port number of the interface is listed as a Member Port. For remote VLANs, the VLAN trunk ports on which the remote VLAN is enabled are listed.
Trunked	A green check mark is displayed if the VLAN ID has been configured as trunked, and is enabled for trunking on all VLAN trunk ports. A VLAN can be enabled for trunking on an individual trunk port or a subset of all trunk ports, in which case the green check mark does not appear. Enabling trunking allows traffic for this VLAN to be sent to remote members of the VLAN who are connected to a different switch in the network. To enable trunking for this VLAN on all trunk ports, see the “Editing VLANs” section on page 12 .
Configure	The Configure icon is enabled for rows that contain PortShield groups. When the Configure icon is disabled, you can edit the settings for that row on the Network > PortShield Groups page.
VLAN Trunks	
Trunk Port	The interface name is displayed in the Trunk Port column. Also, the number of VLAN entries that are enabled on this trunk port is given in parentheses.

Item	Description
VLAN ID	The VLAN ID of each VLAN enabled on the trunk port is displayed when the arrow next to the interface name is pointing downward. Click the right arrow to expand the list.
Configure	The Configure column shows a delete icon if the entry on the row can be deleted. A row containing a VLAN ID that is marked as Trunked in the VLAN Table will not display a delete icon.

You can mark certain PortShield groups as “Trunked”. For information about how to do this, see the “Editing VLANs” section on page 12. Once the PortShield group is dismantled, the associated VLAN is automatically disabled on the trunk ports.

VLANs can exist locally in the form of PortShield groups or can be totally remote VLANs. In Figure 3, the Network > PortShield page shows a PortShield group with X14 as the PortShield interface and X15, X16, and X17 as members of the PortShield group. X20 and X21 are VLAN trunk ports.

Figure 3 Switch Port Settings on Network > PortShield Groups Page



You can change the VLAN ID of PortShield groups on the SonicWALL NSA 2400MX appliance. This allows easy integration with existing VLAN numbering.

Unlike traditional Layer 2 switches, the SonicWALL NSA 2400MX appliance does not allow changing port VLAN membership in an ad-hoc manner. VLAN membership of a port must be configured via PortShield configuration in the SonicOS management interface.

For more information about configuring PortShield groups, see the “Configuring PortShield Interfaces” chapter in the *SonicOS 5.7 Administrator's Guide*.

A virtual interface (called the VLAN Trunk Interface) is automatically created for remote VLANs. When the same remote VLAN is enabled on another trunk port, no new interface is created. All packets with the same VLAN tag ingressing on different trunk ports are handled by the same virtual interface. This is a key difference between VLAN sub-interfaces and VLAN trunk interfaces.

As shown in [Figure 4](#), the **Name** column on the Network > Interfaces page displays the VLAN Trunk Interfaces for the VLAN trunks on which VLAN IDs 100 and 200 are enabled.

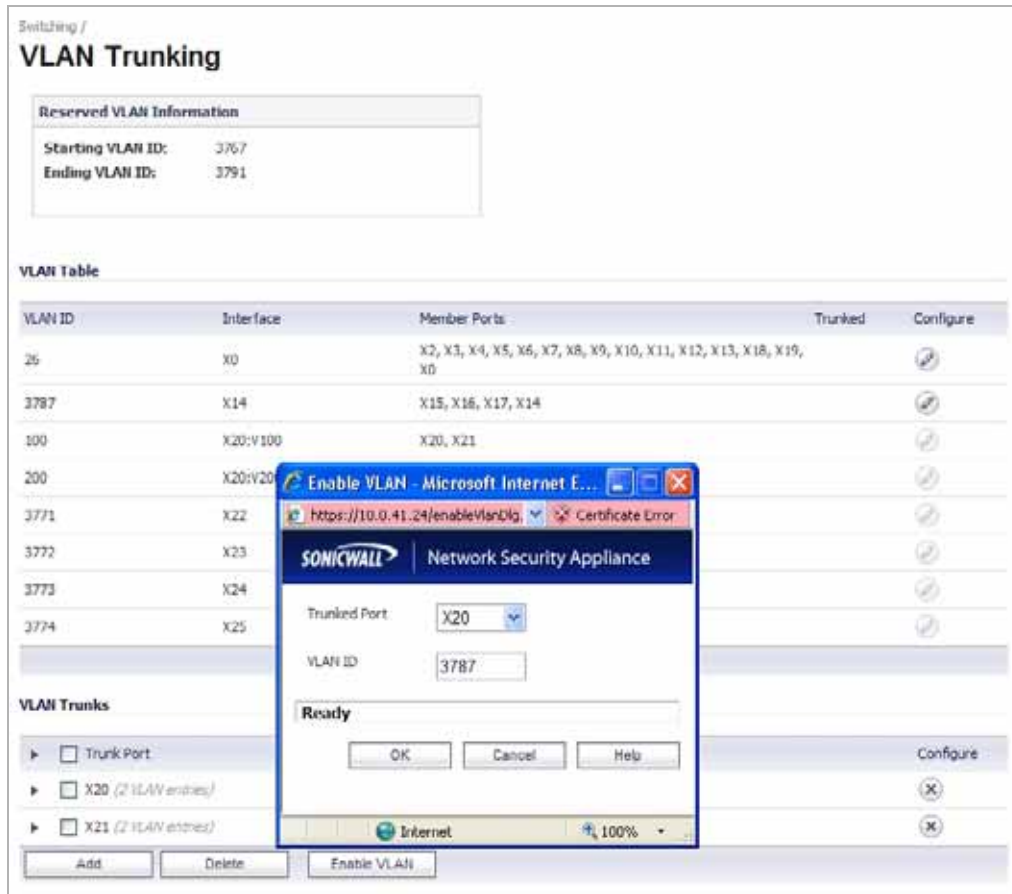
Figure 4 VLAN Trunk Interfaces on Network > Interfaces Page

Name	Zone	Group	IP Address	Subnet Mask	IP Assignment	Status	Comment	Configure
X0	LAN		192.168.168.168	255.255.255.0	Static	No link	Default LAN	
X1	WAN	Default LB Group	10.0.41.24	255.255.0.0	Static	100 Mbps full-duplex	Default WAN	
X20:V100	LAN		192.144.144.10	255.255.255.0	Static	Trunk-VLAN I/F	Sales	
X20:V200	LAN		192.145.145.10	255.255.255.0	Static	Trunk-VLAN I/F	Engineering	
X22	Unassigned		0.0.0.0	0.0.0.0	N/A	No link		
X23	Unassigned		0.0.0.0	0.0.0.0	N/A	No link		
X24	Unassigned		0.0.0.0	0.0.0.0	N/A	No link		
X25	Unassigned		0.0.0.0	0.0.0.0	N/A	No link		

You can enable any VLAN, local or remote, on a VLAN trunk to allow bridging to respective VLANs on another switch. For example, local VLAN 3787, created from a PortShield group, can be enabled on the VLAN trunk for port X20, which also has two remote VLANs enabled on it.

Figure 5 shows the user interface while enabling the local VLAN 3787 on the trunk port, X20.

Figure 5 Enabling a Local VLAN on a VLAN Trunk



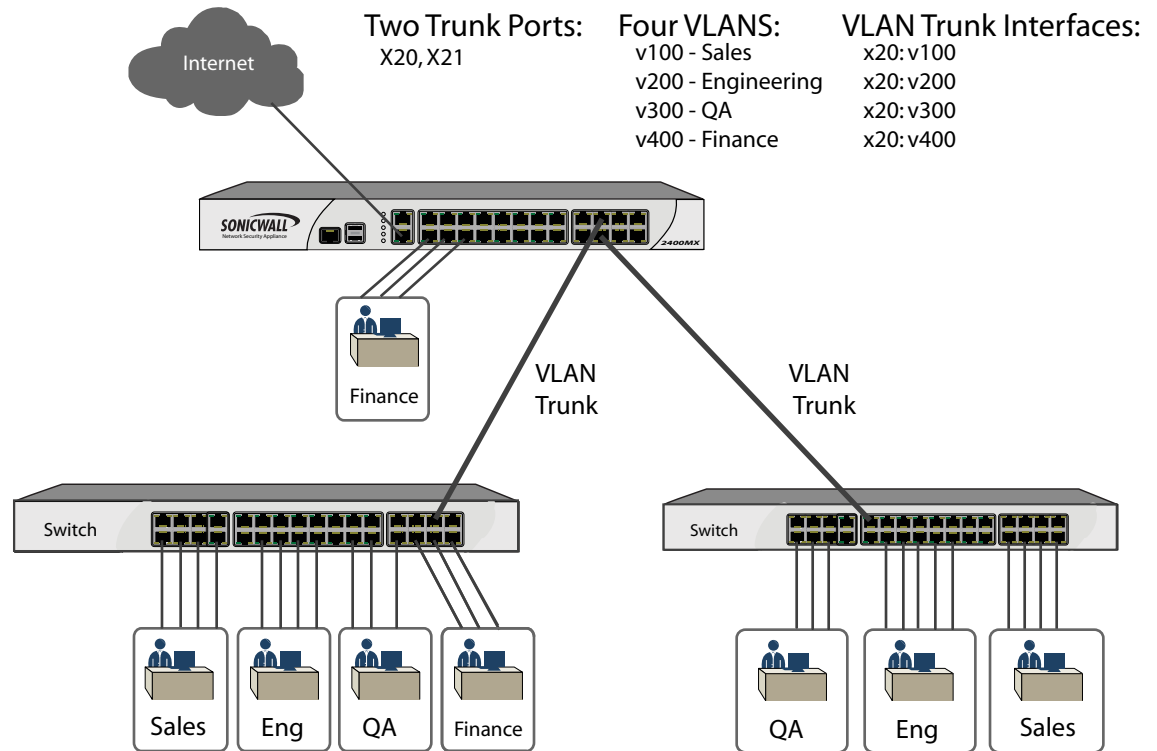
In Figure 6, the **VLAN Table** on the Switching > VLAN Trunking page displays the trunk port, X20, as a member of local VLAN 3787 after the VLAN is enabled on the VLAN trunk.

Figure 6 VLAN Table

VLAN ID	Interface	Member Ports	Trunked	Configure
26	X0	X2, X3, X4, X5, X6, X7, X8, X9, X10, X11, X12, X13, X18, X19, X0		
3787	X14	X15, X16, X17, X14, X20		
100	X20:V100	X20, X21		
200	X20:V200	X20, X21		

Figure 7 illustrates a VLAN trunk with two trunk ports, bridging the Sales, Engineering, QA, and Finance VLANs through the SonicWALL NSA 2400MX. Each remote VLAN is initially enabled on VLAN trunk port X20, causing the creation of four virtual VLAN trunk interfaces. When these VLANs are also enabled on trunk port X21, no new virtual interfaces are created.

Figure 7 VLAN Trunk Bridging Four VLANs Through the SonicWALL NSA 2400MX



VLAN trunking interoperates with Rapid Spanning Tree Protocol (RSTP), Link Aggregation and Port Mirroring features. A VLAN trunk port can be mirrored, but cannot act as a mirror port itself. VLAN trunk ports are used to pass traffic to other networking devices. By comparison, traffic arriving on a mirror port has already been handled or sent to its destination, and the mirror port does not forward it again.

You cannot enable static port security on the VLAN trunk port. When using static port security, MAC address objects for the trusted MAC addresses are bound to a port, and frames from other source addresses are dropped. This is not supported on VLAN trunk ports in this release.

Ports configured as VLAN trunks cannot be used for any other function and are reserved for use in Layer 2 only. For example, you cannot configure an IP Address for the trunk ports.


When a Trunk VLAN interface has been configured on a particular trunk port, that trunk port cannot be deleted until the VLAN interface is removed, even though the VLAN is enabled on multiple trunk ports.

See the following procedures:

- [“Editing VLANs” on page 12](#)
- [“Adding a VLAN Trunk Port” on page 12](#)
- [“Enabling a VLAN on a Specific Trunk Port” on page 13](#)
- [“Deleting VLAN Trunk Ports” on page 13](#)

Editing VLANs

To edit a VLAN, perform the following steps:

- Step 1** On the Switching > VLAN Trunking page, click the Configure icon  in the VLAN Table row for the VLAN ID you want to edit.
- Step 2** In the **Edit Vlan for PortShield** window, do one of the following:
- Type a different VLAN ID into the **Vlan ID** field. You can enter any VLAN ID except the original system-specified VLAN ID or any others in the Reserved VLAN IDs.
 - Use the VLAN ID number in the **Vlan ID** field, which matches the one for which you clicked the Configure icon.



- Step 3** To enable trunking for this VLAN, select the **Trunked** checkbox. To disable trunking for this VLAN, clear the checkbox.

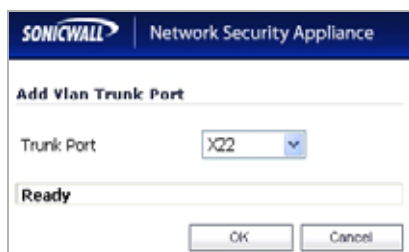
Enabling trunking allows traffic for this VLAN to be sent to remote members of the VLAN who are connected to a different switch in the network. The VLAN ID is automatically added to the list of VLAN entries for each trunk port, indicating that the VLAN is enabled on the trunk ports. Also, a green check mark is displayed for this VLAN in the VLAN Table on the Switching > VLAN Trunking page.

- Step 4** Click **OK**.

Adding a VLAN Trunk Port

To add a VLAN trunk port, perform the following steps:

- Step 1** On the Switching > VLAN Trunking page under **VLAN Trunks**, click the **Add** button.
- Step 2** In the Add VLAN Trunk Port window, select the port to add from the **Trunk Port** drop-down list.




- Step 3** Click **OK**.

Enabling a VLAN on a Specific Trunk Port

Using this method rather than the method described in “Editing VLANs” on page 12, you can specify a single trunk port to be used for a particular VLAN ID.

To enable a custom VLAN ID on a specific trunk port, perform the following steps:

- Step 1** On the Switching > VLAN Trunking page under **VLAN Trunks**, click the **Enable VLAN** button.
- Step 2** In the **Enable VLAN** window, select a trunked port from the **Trunked Port** drop-down list. This is the port that you want to use to trunk the VLAN ID indicated in the next field.




- Step 3** In the **VLAN ID** field, type in the VLAN ID to be trunked. This can be a VLAN ID on another switch.
- Step 4** Click **OK**.

Deleting VLAN Trunk Ports

To delete one or more VLAN trunk ports, perform the following steps:

- Step 1** On the Switching > VLAN Trunking page under **VLAN Trunks**, select one or more checkboxes for the VLAN trunk ports you want to delete.



- Step 2** Click the **Delete** button.
- Step 3** Click **OK** in the confirmation dialog box.

Configuring Rapid Spanning Tree

The Rapid Spanning Tree Protocol (RSTP) is implemented to support Layer 2 network designs with redundant paths.

The first Spanning Tree Protocol was defined in by IEEE 802.1D in 1990, updated in 1998, and replaced in 2004. The 802.1D-2004 standard removed STP and replaced it with RSTP, also adding other extensions 802.11c and 802.1t. RSTP was originally defined in 802.1w, introduced in 1998 and published in 2001. RSTP

allows faster spanning tree convergence after a topology change, typically within 3 times the Hello interval, a total of 6 seconds in the default configuration. The original STP can take 30 to 50 seconds for convergence after a topology change.

SonicWALL's RSTP implementation conforms to the IEEE 802.1D-2004 specification. The 802.1D specification is VLAN unaware and creates a common spanning tree (CST) that is applied to all VLANs present in the network. The RSTP implementation is backward compatible with the original 802.1D standard (STP).

The Switching > Rapid Spanning Tree page, shown in Figure 8, displays the root bridge ID and other information about the root bridge. It also provides a way to configure bridge settings and port settings for the SonicWALL NSA 2400MX.

See the following sections for bridge and port settings configuration information:

- “Configuring Bridge Settings” on page 17
- “Configuring Port Settings” on page 17

Figure 8 Switching > Rapid Spanning Tree Page

The screenshot displays the 'Rapid Spanning Tree' configuration page. At the top, there are 'Apply' and 'Cancel' buttons. Below this is the 'Bridge Information' section, which lists the following details:

- Root Bridge ID: SonicWALL:3c:d0:7d
- Root Bridge: Yes
- Root Priority: 32768
- Root Path Cost: 0
- Root Port: 0
- Root Age Time (sec): 0
- Root Max Age (sec): 20
- Root Forward Delay (sec): 15
- Root Hello Time (sec): 3

The 'Bridge Settings' section includes the following fields:

- Force Version: RSTP Operation (dropdown menu)
- Bridge Priority: 32768 (text input)
- Hello Time (secs): 3 (text input)
- Forward Delay (secs): 15 (text input)

The 'Port Settings' section contains a table with the following data:

Name	Type	Cost	Priority	State	Role	Enable	Configure
X20	No link	0(Auto)	16	Discarding	Disabled	<input checked="" type="checkbox"/>	
X21	No link	0(Auto)	21	N/A	N/A	<input type="checkbox"/>	

The settings displayed in the Bridge Information section of the Switching > Rapid Spanning Tree page are described in [Table 2](#).

Table 2 RSTP Configurable Objects

Item	Description
Root Bridge ID	The root bridge ID is an 8-byte value with 2 bytes for the bridge priority and 6 bytes for the MAC address. The root bridge has the lowest value for priority among all switches in the network.
Root Bridge	Indicates whether or not the SonicWALL NSA 2400MX is the root bridge in the bridged LAN. The root bridge is chosen by an election process among all switches in the network, based on bridge priority and bridge MAC address.
Root Priority	Bridge priority is configurable in multiples of 4,096 with a default value of 32,768 and maximum of 61,440. Lower numbers indicate higher priority. Bridge priority is the key factor in determining the root bridge. It also determines the designated bridge for each LAN segment when multiple bridges have the same path cost to the root bridge.
Root Path Cost	The root path cost is based on the speed of the interface that connects to the root bridge, and is used by RSTP to calculate the shortest path to the elected root bridge.
Root Port	The root port is the interface on a bridge that provides the shortest path to the root bridge.
Root Age Time (sec)	The root age time is the number of seconds since the last hello packet arrived from the root bridge.
Root Max Age (sec)	The root max age is the time interval without sending a hello packet after which a switch is assumed to be unreachable and the Spanning Tree network topology is updated.
Root Forward Delay (sec)	The root forward delay is the time allowed for the listening and learning state. It is also the time that it takes to convert an interface from a blocking state to a forwarding state. The default is 15 seconds.
Root Hello Time	The root hello time is the time interval between hello packets sent to the root bridge.

Auto detection of non-edge ports is not supported. A non-edge port is one that is connected directly to an end-user computer such as a PC or laptop.

You can enable/disable RSTP on VLAN trunk ports only. By default, RSTP is disabled on trunk ports.




Note

You should enable the RSTP before performing physical network connectivity between the SonicWALL NSA 2400MX and another switch.

When the SonicWALL NSA 2400MX is booting up, ports are disabled until Spanning Tree configuration is applied. The NSA 2400MX automatically soft-bridges the STP Bridge Protocol Data Units (BPDUs) between the ports to prevent loops when ports in the same VLAN (PortShield group or L2 Bridge mode) are connected to another switch. This allows the remote switch to detect that its ports are connected to another switch and it can automatically block certain ports.

Figure 9 illustrates the possible loop that is logically blocked when ports in the same VLAN exist on multiple connected switches in a network.

Figure 9 RSTP on VLAN Trunk Ports Detects and Blocks Loops Between Switches



You can view the port settings for RSTP-eligible interfaces on the Switching > Rapid Spanning Tree page. See Table 3 for an explanation of the port settings.

Table 3 Switching > Rapid Spanning Tree Port Settings

Item	Description	Configurable from this Page
Name	Indicates the interface number, such as X20.	No
Type	Indicates link status and speed.	No
Cost	Displays the port cost. This can be a configured value or an automatically calculated value based on link speed.	Yes
Priority	Displays the port priority. This defaults to the interface number unless configured otherwise.	Yes
State	Indicates whether the port is currently forwarding, discarding, or blocking packets.	No
Role	Indicates the role and status of the port, such as root, designated, alternate, backup, or disabled.	No
Enabled	Indicates whether or not RSTP (or STP) is enabled.	Yes

Configuring Bridge Settings

To configure the Bridge Settings on the Switching > Rapid Spanning Tree page, perform the following steps:

-
- Step 1** To specify the spanning tree protocol version to use, select one of the following from the **Force Version** drop-down list:
- **RSTP Operation** – Use Rapid Spanning Tree Protocol.
 - **STP Only** – Use the original Spanning Tree Protocol.
- Step 2** To specify the priority of the root bridge, type the desired priority into the **Bridge Priority** field. Bridge priority is configurable in multiples of 4,096 with a default value of 32,768 and maximum of 61,440. Lower numbers indicate higher priority. Bridge priority is the key factor in determining the root bridge. It also determines the designated bridge for each LAN segment when multiple bridges have the same path cost to the root bridge.
- Step 3** To specify the Hello time, type the desired number of seconds to allow into the **Hello Time (secs)** field. The Hello time is the time interval between transmission of BPDUs by the root bridge and other bridges/switches in the network. The default is 3 and the range is 1 to 10 seconds. The Hello time is communicated to other switches by including it in the BPDU along with other topology information for the bridged LAN.
- Step 4** To specify the forward delay, type the desired number of seconds into the **Forward Delay (secs)** field. The forward delay is the time allowed for the listening and learning state. It is also the time that it takes to convert an interface from a blocking state to a forwarding state. The default is 15 and the range is 4 to 30 seconds. The forward delay setting is communicated to other switches by including it in the BPDU.
- Step 5** When finished, click **Apply**.

Configuring Port Settings

You can enable or disable RSTP and configure path cost and priority for each VLAN trunk port on your SonicWALL NSA 2400MX.

When port settings have been specified for an interface, the Port Settings table on the Switching > Rapid Spanning Tree page contains a row for that interface. A Configure icon is enabled for it unless Link Aggregation is enabled for the interface.



Note If you need to enable RSTP on interfaces in a link aggregation group, first enable RSTP on the individual ports and then enable link aggregation.

To configure the Port Settings on the Switching > Rapid Spanning Tree page, perform the following steps:

-
- Step 1** Under Port Settings, click the Configure icon in the row for the interface you want to edit.

- Step 2** In the **Edit RSTP Settings** window, select the **Enable RSTP** checkbox to enable Rapid Spanning Tree Protocol for this interface. Clear the checkbox to disable RSTP on this interface.

- Step 3** To allow the path cost for the port to be automatically calculated by SonicOS, select the **Auto** checkbox. The **Auto** option is enabled by default. If left in auto-mode, the port cost is determined based on link speed.
- Step 4** To specify the path cost for the port, type the desired cost value into the **Port Path Cost** field. You can assign an arbitrary cost value or base the cost on guidelines provided by the RSTP or STP specification. The cost is higher for lower bandwidth connections. According to guidelines, the cost of a 1 Gbps bandwidth connection would be 2, compared to the cost of 19 for a 100 Mbps Fast Ethernet connection or 100 for a 10 Mbps connection.
- Step 5** To specify the port priority, type the desired priority into the **Port Priority** field. The port priority defaults to the interface number unless configured otherwise. The range is 0 to 15. A lower number indicates higher priority. Port priority is important when multiple ports are connected to the same switch and there is a possible loop, or in the case where the cost of the path to the root bridge is the same as it is for another port. The port with the lower number for port priority is used to forward traffic. The port with the lower priority (indicated by a higher number) is blocked.

**Note**

The range for input into the Port Priority field is 0 to 15. The number you enter is automatically multiplied by 16 when applied.

Configuring Layer 2 Discovery

The Switching > Layer 2 Discovery page accesses information about switches and other devices in the network, using the Link Layer Discovery Protocol (LLDP). LLDP is a non-proprietary protocol used by network devices in the LAN to advertise their identity, capabilities, and interconnections. The LLDP protocol is defined by the IEEE 802.1AB standard, which is titled “Station and Media Access Control Connectivity Discovery.”

In addition to LLDP, the SonicWALL NSA 2400MX uses Microsoft Link Layer Topology Discovery (LLTD) protocol to discover nodes visible from a port. Link Layer Topology Discovery (LLTD) is a Microsoft proprietary protocol with functionality similar to LLDP. It operates on wired or wireless networks (Ethernet 802.3 or wireless 802.11). LLTD is included on Windows Vista and Windows 7, and can be installed on Windows XP.

**Note**

Windows XP users need to download, install, and enable the LLTD responder driver from Microsoft.

Both LLDP and LLTD are Layer 2 protocols and do not cross a broadcast domain. A switch forwarding table is also used during discovery, and an ARP table is used to connect MAC addresses to IP addresses.

On many switches and network devices, the LLDP information is stored as a management information database (MIB). Simple Network Management Protocol (SNMP) is used to query the MIB for device information, including system name, port name, VLAN name, IP address, system capabilities (such as switching or routing), MAC address, and link aggregation settings. The topology of a network can be discovered by crawling the hosts and querying the MIB database on each.


The LLDP transmitter is not implemented in SonicOS 5.7.0.0. The Switching > L2 Discovery page displays Layer 2 information obtained via LLDP from other, LLDP-enabled, switches and devices in the network. [Figure 10](#) shows information obtained via discovery on the X1 (WAN) interface.

Figure 10 Switching > L2 Discovery Page

Switching /





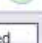
L2 Discovery

<input type="checkbox"/> Interface	MAC Address	IP Address	System Name	Description
<input type="checkbox"/> X0 (0 entries)				
<input checked="" type="checkbox"/> X1 (10 entries)	SonicWALL:65:29:8b	10.0.81.140		
	SonicWALL:0f:4c:94	10.48.254.254		
	SonicWALL:0d:d2:29	10.0.79.5		
	SonicWALL:6a:00:00	10.48.254.254		
	SonicWALL:15:22:7d	10.0.84.77		
	SonicWALL:6a:00:01	10.48.254.253		
	SonicWALL:3a:1a:89	10.0.11.131		
	SonicWALL:16:66:41	10.0.32.24		
	SonicWALL:0f:5b:de	100.0.0.254		
	SonicWALL:10:68:d5	10.0.37.98		
<input type="checkbox"/> X2 (0 entries)				
<input type="checkbox"/> X3 (0 entries)				
<input type="checkbox"/> X4 (0 entries)				

The Switching > L2 Discovery feature does not proactively manage the discovery. Discovery is active when the system boots up and then does not restart unless you click the L2 Discovery refresh button  in the SonicOS management interface.

To restart Layer 2 discovery on multiple interfaces, you can select the checkbox next to the desired interfaces and then click the **Refresh Selected** button at the bottom of the page, as shown in [Figure 11](#).

Figure 11 Restarting Layer 2 Discovery on Selected Interfaces

<input type="checkbox"/> X23 (0 entries)	
<input checked="" type="checkbox"/> X24 (0 entries)	
<input checked="" type="checkbox"/> X25 (0 entries)	
<input type="checkbox"/> U0 (0 entries)	
<input type="checkbox"/> U1 (0 entries)	

Viewing Device Information in the Layer 2 Discovery Page

To view the LLDP/LLTD discovery results for your network, perform the following steps:


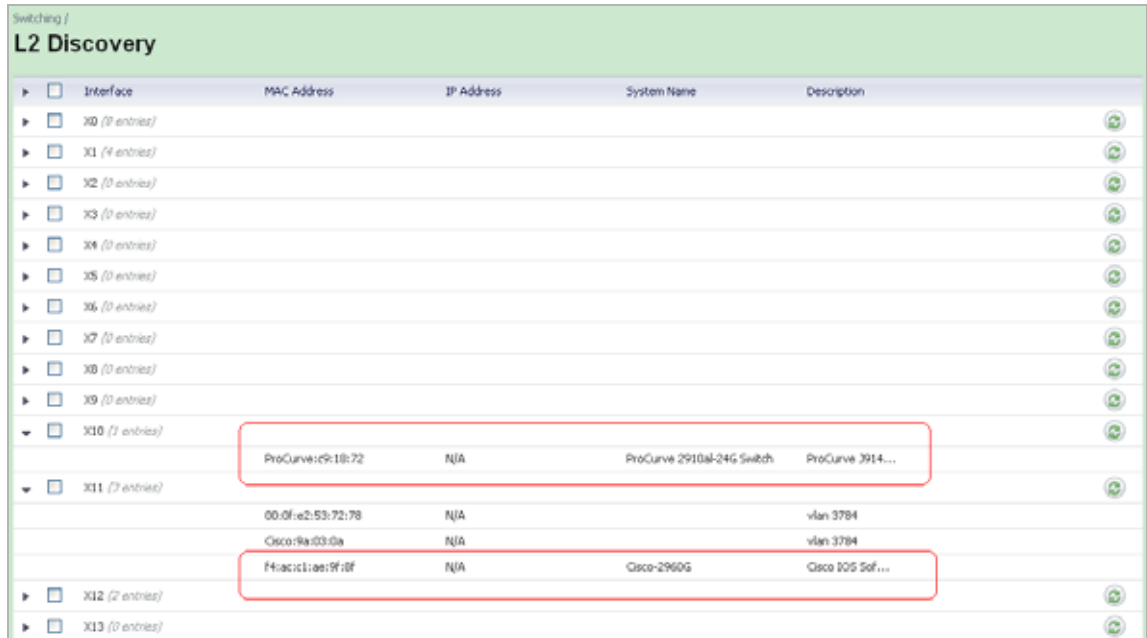
- Step 1** Enable LLDP on any switches or other network devices in your network, using a command such as “lldp run”. LLDP is usually not enabled by default.
- Step 2** To get LLTD results from Windows XP machines in your network, download, install, and enable the LLTD responder driver from Microsoft on those machines. LLTD is installed and enabled by default on Windows Vista and Windows 7 machines.
- Step 3** In the SonicOS management interface, navigate to the Switching > L2 Discovery page.
- Step 4** To view the LLDP/LLTD results for a single interface, click the Refresh  button in the same row as the interface.
- Step 5** To view the LLDP/LLTD results for multiple interfaces, select the checkboxes for those interfaces and then click the **Refresh Selected** button at the bottom of the page. You can select all interfaces by selecting the checkbox next to **Interface** in the table heading.

Figure 12 shows the LLDP discovery results for an HP ProCurve switch and a Cisco switch (both circled in red). The other details are the LLTD results.

Figure 12 Switch Information in the Switching > Layer 2 Discovery Page



Interface	MAC Address	IP Address	System Name	Description
X10 (1 entries)	ProCurve:c9:1b:72	N/A	ProCurve 2910al-24G Switch	ProCurve 3914...
X11 (7 entries)	00:0f:e2:53:72:78	N/A		vlan 3784
	Cisco:9a:03:0a	N/A		vlan 3784
	F4:ac:c1:a:e:9f:bf	N/A	Cisco-2960G	Cisco IOS Sof...

Configuring Link Aggregation

SonicOS 5.7 supports the IEEE 802.1AX-2008 Link Aggregation Control Protocol (LACP). LACP is used when multiple network ports are connected in parallel between two switches or between a switch and a server. Link aggregation makes it possible to increase the bandwidth beyond the limits of a single connection, and to provide seamless, higher availability by creating a redundant link.

Link aggregation in SonicOS 5.7 allows port redundancy and load balancing in Layer 2 networks. Load balancing is controlled by the hardware, based on source and destination MAC address pairs. The Switching > Link Aggregation page provides information and statistics, and allows configuration of interfaces for aggregation.

Figure 13 shows the Switching > Link Aggregation page in the SonicOS user interface.

Figure 13 Switching > Link Aggregation Page

Port	LAG ID	Key	Aggregator	LACP Enable	Status	Partner	Action
X20	0	99	✓	✓	down	00:00:00:00:00:00	
X21	0	99		✓	down	00:00:00:00:00:00	

Add...


Static and Dynamic Link Aggregation are supported. Dynamic Link Aggregation is supported with the use of LACP (IEEE 802.1AX). Ports that are in the same VLAN (same PortShield Group) or are VLAN trunk ports are eligible for link aggregation. Up to four ports can be aggregated in a logical group called a Logical Link, and there can be four Logical Links configured.

Two main types of usage are enabled by this feature:

- SonicWALL NSA 2400MX to Server – This is implemented by enabling link aggregation on ports within the same VLAN (same PortShield Group). This configuration allows port redundancy, but does not support load balancing in the NSA 2400MX-to-Server direction due to a hardware limitation on the NSA 2400MX.
- SonicWALL NSA 2400MX to Switch – This is allowed by enabling link aggregation on VLAN trunk ports. Load balancing is automatically performed by the hardware. The NSA 2400MX supports one load balancing algorithm based on source and destination MAC address pairs.

Figure 14 illustrates the two types of link aggregation.

Figure 14 Two Types of Link Aggregation: NSA to Server and NSA to Switch



Similarly to PortShield configuration, you select an interface that represents the aggregated group. This port is called an aggregator. The aggregator port must be assigned a unique key. By default, the aggregator port key is the same as its interface number. Non-aggregator ports can be optionally configured with a key, which can help prevent an erroneous Logical Link if the switch connections are wired incorrectly.

Ports bond together if connected to the same link partner and their keys match. If there is no key configured for a port (if the port is in auto mode), it will bond with an aggregator that is connected to the same link partner. The link partner is discovered via LACP messages. A link partner cannot be discovered for Static link aggregation. In this case, ports aggregate based on keys alone.

Like a PortShield host, the aggregator port cannot be removed from the Logical Link since it represents the Logical Link in the system.



Note

Once link aggregation has been enabled on VLAN trunk ports, additional VLANs cannot be added or deleted on the Logical Link.



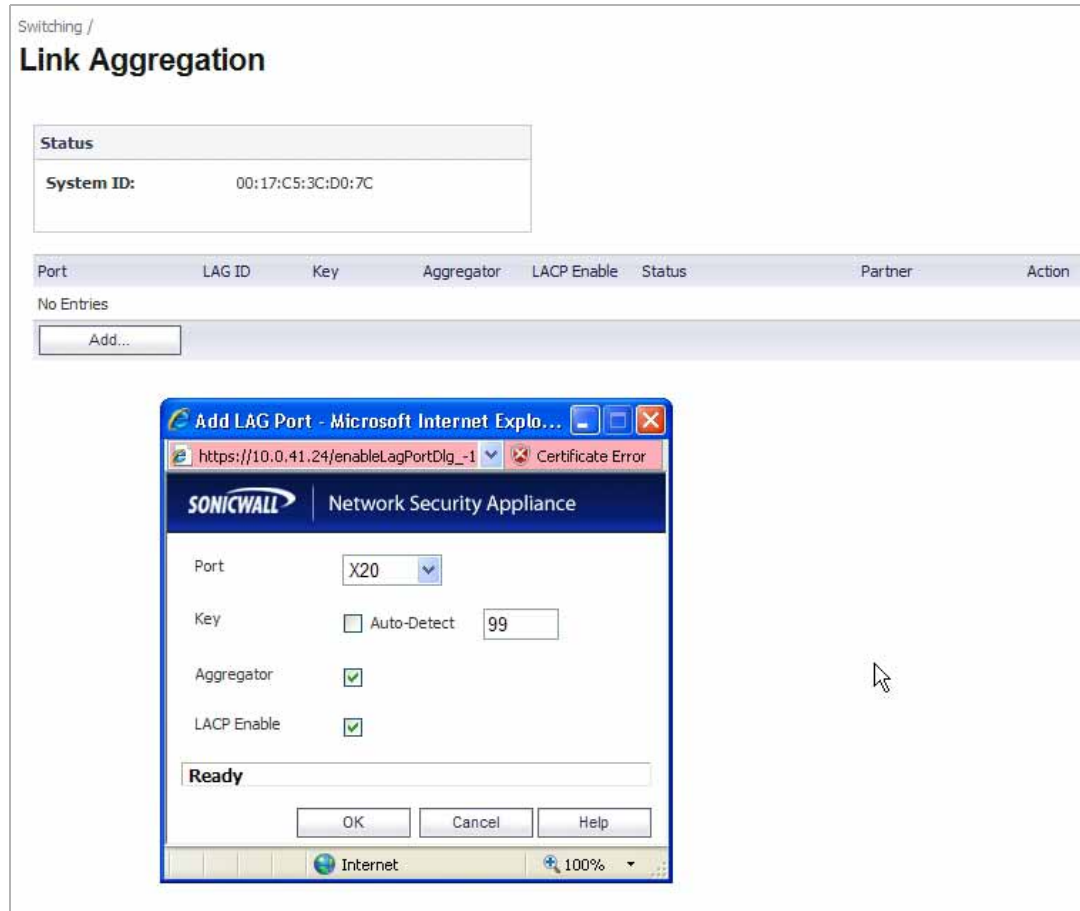
Note

If you need to enable RSTP on the Logical Link, first enable RSTP on the individual members and then enable link aggregation.

Creating a Logical Link

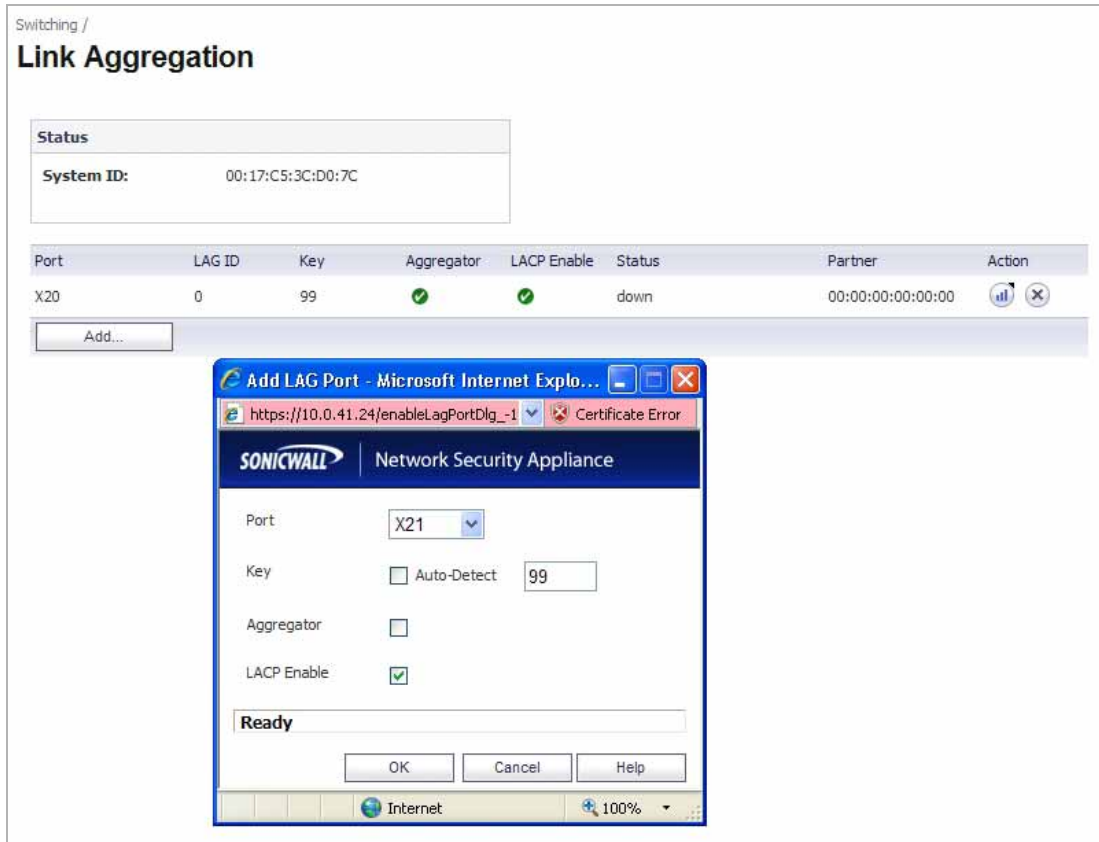
To create a Logical Link, perform the following steps:

- Step 1** On the Switching > Link Aggregation page, click the **Add** button.
- Step 2** In the **Add LAG Port** window, select the interface from the **Port** drop-down list.



- Step 3** To specify a key, clear the **Auto-Detect** checkbox and type the desired key into the **Key** field.
- Step 4** If this interface will be the aggregator for the Logical Link, select the **Aggregator** checkbox. Only one interface can be an aggregator for a Logical Link.
- Step 5** To enable LACP, select the **LACP Enable** checkbox. Dynamic Link Aggregation is supported with the use of LACP. The link partner is discovered via LACP messages.
- Step 6** Click **OK**.
- Step 7** On the Switching > Link Aggregation page, click the **Add** button again.

Step 8 In the **Add LAG Port** window, select the interface for the link partner from the **Port** drop-down list.



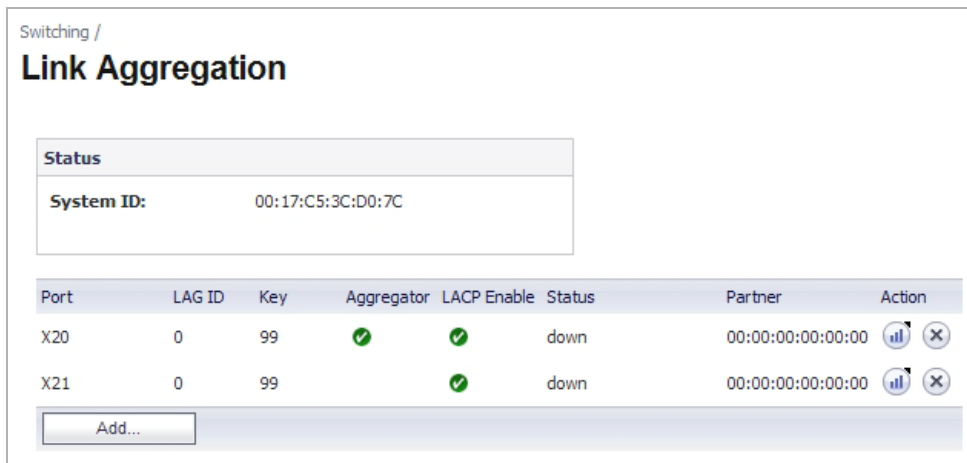
Step 9 If you specified a key for the first interface (the aggregator), clear the **Auto-Detect** checkbox and type the same key into the **Key** field. If **Auto-Detect** was left enabled for the first interface, leave it enabled for this one as well.

Step 10 Clear the **Aggregator** checkbox. Only one interface can be an aggregator for a Logical Link.

Step 11 Select the **LACP Enable** checkbox. This is necessary to create the Logical Link.

Step 12 Click **OK**.

The Switching > Link Aggregation page displays the Logical Link. The **Partner** column will display the MAC addresses of the link partners after they are physically connected.



Configuring Port Mirroring

You can configure Port Mirroring on the SonicWALL NSA 2400MX to send a copy of network packets seen on one or more switch ports (or on a VLAN) to another switch port called the mirror port. By connecting to the mirror port, you can monitor the traffic passing through the mirrored port(s).

Figure 15 shows the Switching > Port Mirroring page with one mirror group configured. Traffic on interface X20 is being mirrored to interface X22, the designated mirror port.

Figure 15 Switching > Port Mirroring Page

Group Name	Mirror Port	Direction	Ingress	Egress	Enable	Configure
Mirror22	X22	both	0	0	<input checked="" type="checkbox"/>	
X20			0	0		

A VLAN trunk port can be mirrored, but cannot act as a mirror port itself. VLAN trunk ports are used to pass traffic to other networking devices. By comparison, traffic arriving on a mirror port has already been handled or sent to its destination, and the mirror port does not forward it again. Typically, the traffic is passed from the mirror port to a computer where the administrator can use an application, such as Wireshark, to view the traffic content.

The Switching > Port Mirroring page allows the administrator to assign mirror ports to mirror ingress, egress or bidirectional packets coming from a group of ports.

See the following procedures:

- [“Configuring a Port Mirroring Group” on page 25](#)
- [“Deleting a Port Mirroring Group” on page 26](#)

Configuring a Port Mirroring Group

To create a new port mirroring group, perform the following steps:

-
- Step 1** On the Switching > Port Mirroring page, click the **New Group** button.

- Step 2** In the **Edit Mirror Group** window, type a descriptive name for the group into the **Interface Group Name** field.

- Step 3** For the **Direction**, select one of the following:
- **ingress** – Select **ingress** to monitor traffic arriving on the mirrored port(s).
 - **egress** – Select **egress** to monitor traffic being sent out on the mirrored port(s).
 - **both** – Select **both** to monitor traffic in both directions on the mirrored port(s).
- Step 4** In the **All Interfaces** list, select the port to mirror the traffic to and click the top right-arrow button to move it to the **Mirror Port** field. You must use an unassigned port as the mirror port.
- Step 5** In the **All Interfaces** list, select one or more ports to be monitored, and click the lower right-arrow button to move it/them to the **Mirrored Ports** field. You will be able to monitor traffic on the mirrored port(s) by connecting to the mirror port.
- Step 6** Click **OK**.
- Step 7** To begin copying network packets from the Mirrored Ports to the Mirror Port, select the **Enable** checkbox in the Switching > Port Mirroring page.
- Step 8** Click **OK**.

Deleting a Port Mirroring Group

To remove a port mirroring group, perform the following steps:

- Step 1** On the Switching > Port Mirroring page, select the checkbox next to the port mirroring group that you want to delete.
- Step 2** Click the **Ungroup** button.
- Step 3** Click **OK** in the confirmation dialog box.

Configuring Layer 2 Quality of Service

Quality of service (QoS) refers to a method of resource control that provides different priority to different types of applications, data, or users. QoS can also be used to guarantee a certain bit rate, delay, jitter, or error rate to a type of network traffic. When network capacity is not large enough to accommodate all traffic at full speed, QoS performance guarantees are essential in the delivery of delay sensitive applications such as Voice over IP (VoIP), online gaming, and Internet TV.

Layer 2, or Ethernet, QoS, provided by SonicOS 5.7 on the SonicWALL NSA 2400MX, is appropriate for real-time streaming multimedia applications such as those mentioned above. At Layer 2, VLANs may also be used to separate traffic of different QoS levels (VLAN tagging is defined in IEEE 802.1Q).

SonicOS Layer 2 QoS supports Class of Service (CoS) as specified in IEEE 802.1p. CoS uses a 3 bit field within the Ethernet frame header. It provides 8 levels of priority for use with QoS algorithms to handle different types of traffic.

SonicOS also supports Differentiated Services Code Point (DSCP). Also known as Differentiated Services or DiffServ, DSCP uses a 6-bit field in the header of IP packets (Layer 3) for packet classification, and provides a simple method of providing QoS guaranteed service to voice or video while using best-effort for traffic that is not delay sensitive. Best-effort service is the default for most Internet traffic and does not provide any guarantees.

The SonicWALL NSA 2400MX appliance can be configured to trust Class of Service (CoS) (IEEE 802.1p) and/or trust Differentiated Services Code Point (DSCP) per port and treat the frames appropriately.

The Switching > Layer 2 QoS page allows the administrator to configure QoS settings per interface. [Figure 16](#) shows the Switching > Layer 2 QoS page.

Figure 16 Switching > Layer 2 QoS Page

Switching /

Layer 2 QoS

Apply Cancel

Settings

Output Scheduling Mechanism: Weighted Round-Robin

DSCP Remap Table [Hide/Show](#)

Value	Priority	Value	Priority	Value	Priority	Value	Priority
Click "Hide/Show" to view the DSCP Remap Table							

[Reset DSCP Remap...](#)

CoS Remap Table [Hide/Show](#)

Value	Priority	Value	Priority	Value	Priority	Value	Priority
Click "Hide/Show" to view the CoS Remap Table							

QoS Settings

<input type="checkbox"/>	Name	Mode	Configure
<input type="checkbox"/>	X0	Both CoS and DSCP (Prefer CoS)	

In SonicOS, four queues with different priority levels (low, normal, high, highest) are supported. These are mapped to the eight levels defined in IEEE 802.1p (CoS) and cannot be changed. [Table 4](#) shows the mapping between the CoS priority levels and the four supported queue priority levels.

Table 4 802.1p Priority Levels Mapped to Four Queues

CoS Priority	Traffic Type	Queue Priority
0	Best Effort	Normal
1	Background	Low
2	Spare	Low
3	Excellent Effort (Business critical)	Normal
4	Controlled Load (Streaming multimedia)	High
5	Video (Interactive Media) [Less than 100ms latency and jitter]	High
6	Voice [Less than 10ms latency and jitter]	Highest
7	Network Control [Lowest latency and jitter]	Highest

On the Switching > Layer 2 QoS page, you can configure the mapping between each value in the DSCP range of 0-63 and the four queues (Low, Normal, High, Highest) supported for Layer 2 QoS, or reset the entire map to the default settings (Normal queue for all DSCP values). Also, on the Firewall > QoS Mapping page, you can configure the mapping between the DSCP range and the 8 levels defined by CoS.

Frames received on ports configured to trust CoS or DSCP are queued appropriately according to the mapping table. An option is provided to select the field to use when both the 802.1p tag field and the DSCP field are present in ingressing frames.

For QoS settings, ports can be assigned a default priority. The default priority is used when Trust CoS or Trust DSCP is enabled, but the information is absent. When Fixed Priority is enabled, the 802.1p tag field and DSCP field are ignored and the default priority is used.

See the following procedures:

- [“Configuring the Scheduling Mechanism” on page 28](#)
- [“Configuring DSCP Mapping” on page 29](#)
- [“Showing the CoS Remap Table” on page 29](#)
- [“Configuring QoS Settings” on page 30](#)

Configuring the Scheduling Mechanism

To configure Weighted Round-Robin or Strict Priority Queue as the output scheduling mechanism, perform the following steps:

-
- Step 1** On the Switching > Layer 2 QoS page, select one of the following from the **Output Scheduling Mechanism** drop-down list:
- **Weighted Round-Robin** – When Weighted Round-Robin is selected, the weighting factors are 8:4:2:1.

- **Strict Priority Queue** – When Strict Priority Queue is selected, packets containing an 802.1p tag or DSCP marking with a priority level matching the Highest queue priority are forwarded or received. Packets matching High, Normal or Low priority may be dropped.

Step 2 Click the **Apply** button.

Configuring DSCP Mapping

You can configure the DSCP mapping by setting the priority levels for DSCP values 0 through 63. The Switching > Layer 2 QoS page also provides a **Reset DSCP Remap** button to reset the priority levels back to the default, which is “Normal.”

To configure DSCP mapping, perform the following steps:

Step 1 To show the DSCP Remap table, click **Hide/Show** next to the **DSCP Remap Table** heading. The priority settings for all DSCP values, 0 - 63, are displayed.

DSCP Remap Table Hide/Show							
Value	Priority	Value	Priority	Value	Priority	Value	Priority
0	Normal	1	Normal	2	Normal	3	Normal
4	Normal	5	Normal	6	Normal	7	Normal
8	Normal	9	Normal	10	Normal	11	Normal

Step 2 For each DSCP value (**0 - 63**) that you want to change, select one of the following from the **Priority** drop-down list:

- **Low**
- **Normal**
- **High**
- **Highest**

Step 3 Click the **Apply** button. The DSCP Remap table is hidden, but if you show it again you will see the updated priority settings.

Step 4 To reset all DSCP mapping back to the default, Normal, click the **Reset DSCP Remap** button and then click **OK** in the confirmation dialog box.

Showing the CoS Remap Table

To show the CoS Remap table, click **Hide/Show** next to the **CoS Remap Table** heading. The priority levels cannot be configured. The CoS Remap table is shown in [Figure 17](#).

Figure 17 Showing the CoS Remap Table on the Switching > Layer 2 QoS Page




CoS Remap Table Hide/Show							
Value	Priority	Value	Priority	Value	Priority	Value	Priority
0	Normal	1	Low	2	Low	3	Normal
4	High	5	High	6	Highest	7	Highest

To hide the CoS Remap table, click **Hide/Show** next to the **CoS Remap Table** heading again.

Configuring QoS Settings

The QoS Settings table on the Switching > Layer 2 QoS page lists all interfaces on the SonicWALL NSA 2400MX. You can configure the QoS settings for each interface individually or for multiple interfaces at the same time. A portion of the QoS Settings table is shown in [Figure 18](#).

Figure 18 QoS Settings Table on the Switching > Layer 2 QoS Page

QoS Settings		
<input type="checkbox"/> Name	Mode	Configure
<input type="checkbox"/> X0	Both CoS and DSCP (Prefer CoS)	
<input type="checkbox"/> X2	Both CoS and DSCP (Prefer CoS)	
<input type="checkbox"/> X3	Both CoS and DSCP (Prefer CoS)	

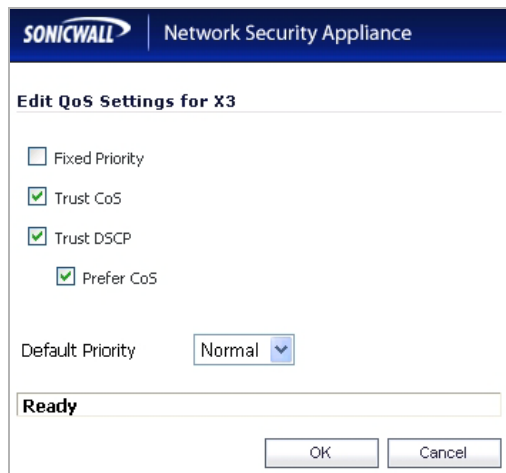
See the following procedures:

- [“Configuring QoS Settings for an Individual Interface” on page 30](#)
- [“Configuring QoS Settings for Multiple Interfaces” on page 31](#)

Configuring QoS Settings for an Individual Interface

To configure QoS settings for frames received on an individual interface, perform the following steps:

- Step 1** On the Switching > Layer 2 QoS page under **QoS Settings**, click the Configure icon in the row for the interface you want to configure. The Edit QoS Settings window opens.



- Step 2** In the Edit QoS Settings window, to enable fixed priority for frames arriving on this interface, select the **Fixed Priority** checkbox. When **Fixed Priority** is selected, the remaining checkboxes are cleared and disabled (greyed out). The CoS 802.1p tag field and DSCP field are ignored and the ingress port's default priority is always used.
- Step 3** To enable the use of the CoS 802.1p tag field settings for Quality of Service on this interface, select the **Trust CoS** checkbox. The **Fixed Priority** checkbox must be cleared before you can select any other checkbox.
- Step 4** To enable the use of the DSCP field settings for Quality of Service on this interface, select the **Trust DSCP** checkbox. The **Fixed Priority** checkbox must be cleared before you can select any other checkbox.

- Step 5** If both **Trust CoS** and **Trust DSCP** are selected, do one of the following:
- Select the **Prefer CoS** checkbox to give preference to the CoS 802.1p tag field settings when both the 802.1p tag field and the DSCP field are present in ingressing frames.
 - Clear the **Prefer CoS** checkbox to give preference to the DSCP field settings when both the 802.1p tag field and the DSCP field are present in ingressing frames.
- Step 6** Select one of the following priority levels from the **Default Priority** drop-down list:
- **Low**
 - **Normal**
 - **High**
 - **Highest**
- If ingressing frames do not contain either a CoS 802.1p tag field or a DSCP field, the default priority is used.
- Step 7** Click **OK**.

Configuring QoS Settings for Multiple Interfaces

To configure QoS settings for frames received on any of several interfaces, perform the following steps:

- Step 1** On the Switching > Layer 2 QoS page under **QoS Settings**, select the checkboxes next to the interfaces you want to configure, and then click the **Configure** button at the bottom of the page. The Edit QoS Settings window opens.

- Step 2** The **Keep original QoS mode of each port** checkbox is selected by default. When this checkbox is selected, each individual port's QoS mode remains unchanged, and only the **Default Priority** setting is changed to the configured value ([Step 7](#)) for each port being configured.
- To activate the other checkboxes in this window and make changes to the QoS settings of the selected interfaces, clear the **Keep original QoS mode of each port** checkbox.
- Step 3** To enable fixed priority for frames arriving on these interfaces, select the **Fixed Priority** checkbox. When **Fixed Priority** is selected, the subsequent checkboxes are cleared and disabled (greyed out). The CoS 802.1p tag field and DSCP field are ignored and the ingress port's default priority is always used.
- Step 4** To enable the use of the CoS 802.1p tag field settings for Quality of Service on these interfaces, select the **Trust CoS** checkbox. The **Fixed Priority** checkbox must be cleared before you can select this checkbox.

- Step 5** To enable the use of the DSCP field settings for Quality of Service on these interfaces, select the **Trust DSCP** checkbox. The **Fixed Priority** checkbox must be cleared before you can select this checkbox.
- Step 6** If both **Trust CoS** and **Trust DSCP** are selected, do one of the following:
- Select the **Prefer CoS** checkbox to give preference to the CoS 802.1p tag field settings when both the 802.1p tag field and the DSCP field are present in ingressing frames.
 - Clear the **Prefer CoS** checkbox to give preference to the DSCP field settings when both the 802.1p tag field and the DSCP field are present in ingressing frames.
- Step 7** Select one of the following priority levels from the **Default Priority** drop-down list:
- **Keep Original Settings** – Choose this setting to allow each interface to default to its original individual QoS settings.
 - **Low**
 - **Normal**
 - **High**
 - **Highest**
- If ingressing frames do not contain either a CoS 802.1p tag field or a DSCP field, the default priority is used.
- Step 8** Click **OK**.

Configuring Rate Control

SonicOS supports per-interface rate limiting and flow control on the Switching > Rate Control page.

Rate limiting provides a way to control the rate of traffic sent or received on a network interface. Traffic is sent or received while its rate is less than or equal to the specified rate limit, while traffic that exceeds the rate is dropped or delayed.

Flow control allows you to manage the rate of data transmission between two devices to prevent a fast sender from getting too far ahead of a slow receiver. Flow control provides a mechanism for the receiver to control the transmission speed to avoid being overwhelmed with data from the sender. Flow control is important in cases where the receiver has a heavy load to process compared to the sender, or if the receiver is a slower computer than the sender.

The Switching > Rate Control page, shown in [Figure 19](#), provides information and configuration for per-interface rate limiting and flow control. Both the rate limiting and flow control features are configured on a per port basis.

Figure 19 Switching > Rate Control Page

Name	Ingress Limit Mode	Ingress Rate (kbits/s)	Egress Rate (kbits/s)	Flow Control	Configure
X0	Limit Broadcast, Multicast and Flooded Unicast	256	0		
X2	Limit Broadcast, Multicast and Flooded Unicast	256	0		
X3	Limit Broadcast, Multicast and Flooded Unicast	256	0		
X4	Limit Broadcast, Multicast and Flooded Unicast	256	0		
X5	Limit Broadcast, Multicast and Flooded Unicast	256	0		
X6	Limit Broadcast, Multicast and Flooded Unicast	256	0		
X7	Limit Broadcast, Multicast and Flooded Unicast	256	0		
X8	Limit Broadcast, Multicast and Flooded Unicast	256	0		
X9	Limit Broadcast, Multicast and Flooded Unicast	256	0		

Egress Traffic Rate Limiting

In SonicOS, the rate limiting for egress frames can only be enabled or disabled, no mode can be selected.

Ingress Traffic Rate Limiting

In SonicOS, the bandwidth of ingress frames can be rate-limited in four modes:

- Limit All Frames – Apply rate limiting to all types of network traffic arriving on the interface
- Limit just broadcast, multicast and flooded unicast frames – Apply rate limiting to broadcast, multicast and flooded unicast traffic arriving on the interface, where:
 - Broadcast traffic is traffic that is simultaneously sent to all computers in the network, typically to announce shared services and state information.
 - Multicast traffic is traffic that is simultaneously sent to a group of computers in the network. Specific multicast addresses are used for different purposes. For example, Spanning Tree Protocol 802.1D for bridges uses the well-known multicast address 01-80-C2-00-00-00, and Spanning Tree Protocol 802.1AD for provider bridges uses the well-known multicast address 01-80-C2-00-00-08.
 - Flooded unicast traffic is traffic that is sent by a switch to all the interfaces in the VLAN. This can occur if the destination MAC address for a frame has not yet been learned by the switch, making it an unknown unicast.
- Limit just broadcast and multicast frames – Apply rate limiting to broadcast and multicast traffic arriving on the interface
- Limit just broadcast frames – Apply rate limiting to broadcast traffic arriving on the interface

The ingress rate limit is rounded to the nearest increment, depending on the the granularity available for that rate. The granularities are different depending on the range of rates:

- 128kbps ~ 1Mbps – increments of 64kbps
- 1Mbps ~ 100Mbps – increments of 1Mbps
- 100Mbps ~ 1000Mbps – increments of 10Mbps (for gigabit ports)

Flow Control

In SonicOS, back-pressure flow control on half-duplex ports and pause frame-based flow control on full-duplex ports are provided to support zero packet loss under temporary traffic congestion.

- Full-duplex flow control requires support from the peer end station. Full-duplex flow control works as follows: when a port's free buffer space is almost empty, the devices send out a PAUSE frame with the maximum pause time to stop the remote node from sending more frames into the switch. The devices also respond to the pause command. Once the PAUSE frame is detected, the port will stop transmission of new data for the amount of time defined in the pause time field of the received PAUSE frame.
- Half-duplex flow control is used to throttle the throughput rate of an end station to avoid dropping packets during network congestion.

Configuring Rate Control Settings for an Interface

To configure rate control settings or to enable flow control, perform the following steps:

- Step 1** On the Switching > Rate Control page, click the Configure icon in the row for the interface you want to configure. The Edit Rate Control Settings window opens.

- Step 2** To enable flow control on this interface, select the **Enable Flow Control** checkbox.
- Step 3** To set the mode for limiting the bandwidth of ingressing frames, select one of the following from the **Ingress Mode** drop-down list:
- **Limit All** – Apply rate limiting to all types of network traffic arriving on the interface
 - **Limit Broadcast, Multicast and Flooded Unicast** – Apply rate limiting to broadcast, multicast and flooded unicast traffic arriving on the interface
 - **Limit Broadcast and Multicast** – Apply rate limiting to broadcast and multicast traffic arriving on the interface, where:
 - Broadcast traffic is traffic that is simultaneously sent to all computers in the network, typically to announce shared services and state information.
 - Multicast traffic is traffic that is simultaneously sent to a group of computers in the network. Specific multicast addresses are used for different purposes. For example, Spanning Tree Protocol 802.1D for bridges uses the well-known multicast address 01-80-C2-00-00-00, and Spanning Tree Protocol 802.1AD for provider bridges uses the well-known multicast address 01-80-C2-00-00-08.
 - Flooded unicast traffic is traffic that is sent by a switch to all the interfaces in the VLAN. This can occur if the destination MAC address for a frame has not yet been learned by the switch, making it an unknown unicast.
 - **Limit Only Broadcast** – Apply rate limiting to broadcast arriving on the interface

- Step 4** Type the desired ingress rate limit in kilobits per second into the **Ingress Rate** field. To turn off the ingress rate limit and allow unlimited traffic, type **0** (zero). The value you type will be rounded to the nearest increment, depending on the the granularity available for that rate. The granularities are different depending on the range of rates:
- 128kbps ~ 1Mbps – increments of 64kbps
 - 1Mbps ~ 100Mbps – increments of 1Mbps
 - 100Mbps ~ 1000Mbps – increments of 10Mbps (for gigabit ports)
- Step 5** Type the desired egress rate limit in kilobits per second into the **Egress Rate** field. To turn off the egress rate limit and allow unlimited traffic, type **0** (zero). The value you type will be rounded to the nearest increment, depending on the the granularity available for that rate. The granularities are the same as for the ingress rate.
- Step 6** Click **OK**.

Configuring Port Security

On the Switching > Port Security page, each port can be configured to enable or disable the **Discard Tagged** option. When it is enabled, all frames with a 802.3ac tag (or “Q-tag”) are discarded. IEEE 802.3ac specifies an extension of 4 bytes to the Ethernet frame size, allowing 1522 bytes per frame. The additional 4 bytes are for the “Q-tag”, which includes 802.1Q VLAN information and 802.1p priority information.

A secure port is meant to receive untagged frames. If a frame has a tag, even when its Security Association (SA) is trusted, it will be discarded.

Only static port security is supported. This means that the SonicWALL NSA 2400MX administrator must create MAC address objects for the trusted MAC addresses and bind the MAC address objects to specific ports. Frames whose source addresses are not contained in the table will be dropped.

MAC address objects are one type of address object in SonicOS. Address objects allow for entities to be defined one time and then re-used in multiple referential instances throughout the SonicOS interface. Address objects can be selected from a drop-down menu in many configuration screens throughout the user interface.

A VLAN trunk port or a port currently configured for link aggregation as part of a Logical Link cannot be a secure port at the same time. This prevents a non-trunk port from connecting to a trunk port.

[Figure 20](#) shows part of the Switching > Port Security page, with one secure port configured.

Figure 20 Switching > Port Security Page

Switching /

Port Security

Static MAC Address

Port	MAC Address Object	Discard Tagged	Configure
<input type="checkbox"/> X0		<input type="checkbox"/>	
<input type="checkbox"/> X2		<input checked="" type="checkbox"/>	
<input type="checkbox"/>	pubs_MAC		
<input type="checkbox"/> X3		<input type="checkbox"/>	
<input type="checkbox"/> X4		<input type="checkbox"/>	

See the following procedures for information about configuring port security:

- “Creating a Secure Port by Adding a MAC Address Object” on page 36
- “Editing MAC Address Objects in Port Security Settings” on page 37
- “Deleting MAC Address Objects from Port Security Settings” on page 37

Creating a Secure Port by Adding a MAC Address Object

To configure port security, you must use an address object to bind MAC address(es) to an interface. You can create an address object from within the procedure described below, or use an existing one. For more information about address objects, see the *SonicOS 5.7 Administrator's Guide*.

To create a secure port by adding a MAC address object to an interface, perform the following steps:

- Step 1** On the Switching > Port Security page, click the **Add** button at the bottom of the page. The Add Static MAC Address window opens.

- Step 2** Select the desired interface from the **Port** drop-down list. This will be the secure port.
- Step 3** If the address object that contains the desired MAC addresses already exists, select it from the **MAC Address** drop-down list and skip to [Step 10](#).
- Step 4** To create a MAC address object, select **Create new address object** from the **MAC Address** drop-down list. The Add Address Object window opens.



Note

Turn off the pop-up blocker in your browser before selecting **Create new address object**.

- Step 5** Type a descriptive name for the address object into the **Name** field.

- Step 6** Select the zone from the **Zone Assignment** drop-down list. This is the zone for the computer with this MAC address. You can select any zone that exists on the SonicWALL NSA 2400MX, including custom zones and the SonicOS default zones, which are LAN, WAN, DMZ, VPN, SSLVPN, MULTICAST, and WLAN.
- Step 7** The only available selection for **Type** is **MAC**, indicating that you are creating a MAC Address Object.
- Step 8** If the device with this MAC address can have multiple IP addresses, select the **Multi-homed host** checkbox. Otherwise, clear this checkbox.
- Step 9** Click **OK** in the Add Address Object window. The new address object appears in the **MAC Address** field of the Add Static MAC Address window.
- Step 10** Click **OK** in the Add Static MAC Address window to complete the secure port configuration using these settings.

Editing MAC Address Objects in Port Security Settings

To edit a MAC address object for a secure port on the Switching > Port Security page, perform the following steps:

- Step 1** Click the Configure icon in the row for the MAC address object you want to edit. The Edit Static MAC Address window opens.

- Step 2** Select a different address object or select **Create new address object** from the **MAC Address** drop-down list and follow the steps provided in [Step 4](#) through [Step 9](#) of the “[Creating a Secure Port by Adding a MAC Address Object](#)” section on page 36.
- Step 3** When finished, click **OK**.

Deleting MAC Address Objects from Port Security Settings

To delete one or more MAC address objects from your secure port settings on the Switching > Port Security page, perform the following steps:

- Step 1** To delete a single MAC address object, click the Delete icon (X) in the Configure column for the row with the MAC address object you want to delete.
- Step 2** To delete multiple MAC address objects, select the checkboxes next to the MAC address objects you want to delete and then click the **Delete Selected** button at the bottom of the page.
- Step 3** Click **OK** in the confirmation dialog box.

Troubleshooting and Verification

This section provides methods you can use to verify and troubleshoot the behavior of your Switching configuration.

See the following sections:

- [“Using Port Mirroring with Wireshark” on page 38](#)
- [“Viewing Log Event Messages for Switching” on page 42](#)

Using Port Mirroring with Wireshark

Because switched frames cannot be viewed with SonicOS Packet Monitoring, configuring Port Mirroring and viewing the mirrored traffic with a network analyzer is the preferred method for examining traffic as it passes through the appliance interfaces. When using Packet Monitoring, the packets are displayed in the SonicOS user interface. However, with Port Mirroring, the mirrored traffic is sent to a port rather than to the user interface. A computer running a network analyzer, such as Wireshark, is connected to the mirror port and the network analyzer is used to display the mirrored traffic.

This section describes an example situation where two VLAN Trunk ports are mirrored to a port that is connected to a computer running Wireshark. The two trunk ports are enabled for Rapid Spanning Tree and Link Aggregation, and the RSTP and LACP Layer 2 frames can be seen in Wireshark.

Configuring Ports for Mirroring

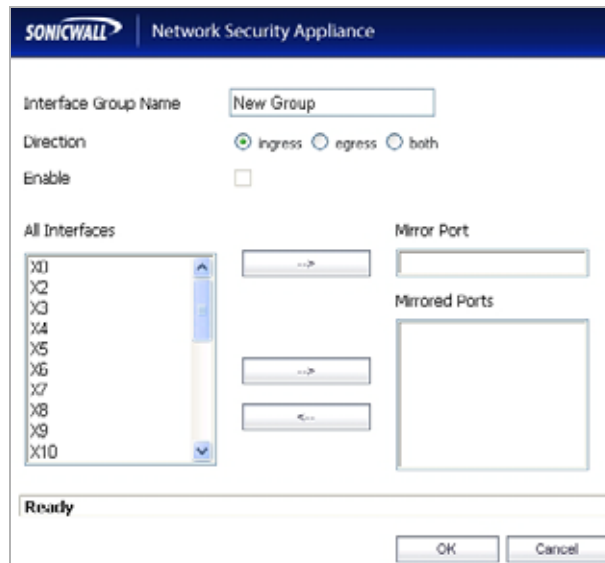
In this use case, the VLAN Trunk ports that we want to mirror are gigabit Ethernet ports. Therefore, we select a gigabit port as the Mirror Port.

To configure Port Mirroring for this use case, perform the following steps:

- Step 1** On the Switching > Port Mirroring page, click the **New Group** button to open the Edit Mirror Group window.



Step 2 In the Edit Mirror Group window, type a name for the Mirror Group.



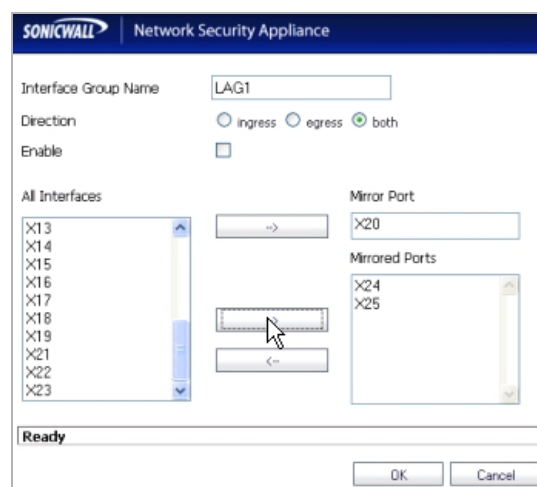
Step 3 For **Direction**, select **both**. This allows mirroring of traffic arriving on, and being transmitted from, the mirrored ports.

Step 4 In the **All Interfaces** box, scroll down and select a gigabit interface for the Mirror Port. The selected interface must have an **Unassigned** zone. You can verify this on the Network > Interfaces page.

Step 5 Click the right arrow button to move the selected interface to the **Mirror Port** field. In this case, we select X20 as the Mirror Port.

Step 6 In the **All Interfaces** box, scroll down and select the interfaces that you want to mirror. We select X24 and X25. These gigabit interfaces are enabled for RSTP and are configured as a Logical Link with X24 as the aggregator.

Step 7 Click the right arrow button to move the selected interfaces to the **Mirrored Ports** field.



Step 8 Click **OK**.

The X20, X24, and X25 interfaces are shown below on the Network > Interfaces page. You can see that all three are gigabit Ethernet ports, that the Zone for X20 is Unassigned and it is configured as a Mirror Port, and that X24 and X25 are configured as VLAN Trunk ports. X25 is marked as a member of a Logical Link.

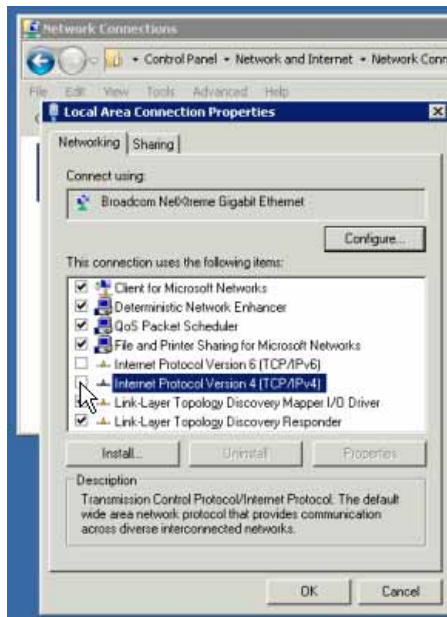
▼ X20	Unassigned			Mirror Port	1000 Mbps full-duplex
▼ X21	Unassigned	0.0.0.0	0.0.0.0	N/A	No link
▼ X22	Unassigned			VLAN Trunk	1000 Mbps full-duplex
X22:V550	E5500_TSL_X2:V550_WAN	3.3.3.1	255.255.255.0	Static	Trunk-VLAN I/F
▼ X23	Unassigned			VLAN Trunk	1000 Mbps full-duplex
▼ X24	Unassigned			VLAN Trunk	1000 Mbps full-duplex
▼ X25	Unassigned			VLAN Trunk,LAG	1000 Mbps full-duplex

Using Wireshark

Wireshark is a popular, open source network analysis tool that runs on Windows or Mac OS X computers. It allows you to analyze network traffic at the frame or packet level. You can download Wireshark for free from <http://www.wireshark.org/>.

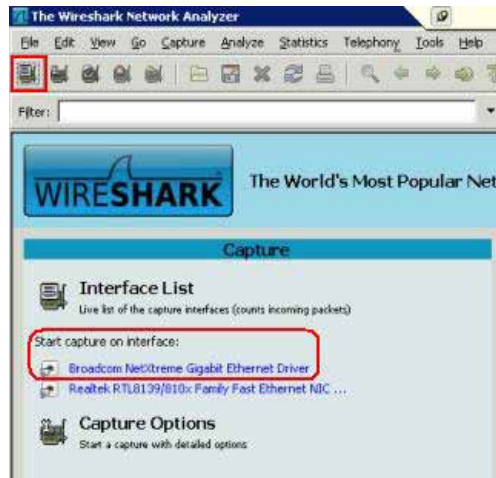
To use Wireshark to view the mirrored packets from the SonicWALL NSA 2400MX:

- Step 1** Connect the SonicWALL NSA 2400MX port X20 to a gigabit Ethernet interface on your Wireshark computer.
- Step 2** To view only Layer 2 traffic in Wireshark, open **Network Connections** on the computer and right-click the gigabit interface you are using.
- Step 3** Select **Properties** in the drop-down list.
- Step 4** Clear the checkboxes for **IPv4** and **IPv6** traffic.



- Step 5** Click **OK**.


Step 6 Launch **Wireshark**.



Step 7 Do one of the following:

- Click the **Start Capture** button in the upper left corner, and then, in the dialog box, select the **Start** checkbox for the gigabit interface connected to the Mirror Port.
- Under **Start capture on interface**, click the link for the interface connected to the Mirror Port.

Step 8 View the frames in the Wireshark main window.



Step 9 When finished, click the **Stop Capture** button.

Step 10 Optionally save the output as a **.cap** or **.pcap** file, by selecting **File > Save as**, choosing a location, and entering a file name.

Viewing Log Event Messages for Switching

A new log event, **logstrAdvSwitch**, is introduced in SonicOS 5.7 to address SonicOS Switching activities. It falls under a new category, Advanced Switching, which can only be seen on devices with the switching hardware, such as the SonicWALL NSA 2400MX. Other SonicWALL appliances will not show the new category, as it is not applicable to the hardware.

You can filter the log for the Advanced Switching category to display only the log events related to switching. For instructions, see the [“Filtering the Log for Switching Events” section on page 43](#).

All messages are informational, and no actions are required by the administrator when these log messages occur.

[Table 5](#) describes the various instances of the **logstrAdvSwitch** log event specific to Switching in SonicOS 5.7.

Table 5 Switching Log Event Messages

Log Event Message	Comments
User configuration of Flow Control on an interface can cause the following log events:	
Flow control on X%d is enabled	<ul style="list-style-type: none"> Where %d is a decimal number in the interface name. For example, X%d could be X3, X22, etc.
Flow control on X%d is disabled	
User configuration of Rate Limiting on an interface can cause the following log events:	
Ingress rate on X%d is limited to x kbps	<ul style="list-style-type: none"> Where %d is a decimal number in the interface name. For example, X%d could be X3, X22, etc. Where x kbps is the number of kilobits per second for the limit.
Ingress rate on X%d is not limited	
Egress rate on X%d is limited to x kbps	
Egress rate on X%d is not limited	
Ingress Rate limiting mode on X%d : Limit All	
Ingress Rate limiting mode on X%d : Limit Broadcast, Multicast and Flooded Unicast	
Ingress Rate limiting mode on X%d : Limit Broadcast and Multicast	
Ingress Rate limiting mode on X%d : Limit Only Broadcast	
User configuration of QoS Priority or Rules on an interface can cause the following log events:	
Default QoS Priority on X%d : x	<ul style="list-style-type: none"> Where x is the number for the priority level, with possible values of 0-7. See Table 4 on page 28 for a mapping of 802.1p priority levels to the four priority queues supported by the SonicOS Switching feature. CoS is Class of Service, IEEE 802.1p DSCP is Differentiated Services Code Point
Desired QoS Rule on X%d : Trust CoS	
Desired QoS Rule on X%d : Trust DSCP	
Desired QoS Rule on X%d : Both CoS and DSCP (Prefer DSCP)	
Desired QoS Rule on X%d : Both CoS and DSCP (Prefer CoS)	
The following message is generated when the hardware determines that both ends of an Ethernet link support flow control and flow control will take effect as needed. the link Xi (e.g X0, X1, X2):	
Flow control on X%d is activated	<ul style="list-style-type: none"> Where X%d indicates the link interface, such as X3.

Filtering the Log for Switching Events

To display only the log events related to switching, perform the following steps:

- Step 1** Navigate to the Log > View page in the SonicOS management interface.
- Step 2** In the Log View Settings section, select **Advanced Switching** from the **Category** drop-down list.

Filter	Value	Group Filters
Priority:	All	<input type="checkbox"/>
Category:	Advanced Switching	<input type="checkbox"/>
Source (IP, Interface):	All Interfaces	<input type="checkbox"/>
Destination (IP, Interface):	All Interfaces	<input type="checkbox"/>

Filter Logic: Priority && Category && Source && Destination

- Step 3** Optionally select specific interfaces from the **Source** and/or **Destination** drop-down lists. The fields for which you select values into are combined into a search string with a logical AND. For example, if you select an interface for **Source** and for **Destination**, the search string will look for log events matching:

Source interface AND Destination interface



- Note** Do not enter IP addresses for Source or Destination, as these are not used in the Layer 2 Advanced Switching log event messages.

- Step 4** Select the **Group Filters** checkbox next to any two or more criteria to combine them with a logical OR.
- Step 5** For example, if you select interfaces for **Source** and **Destination**, and select **Advanced Switching** for **Category**, and then select the **Group Filters** checkboxes next to **Source** and **Destination**, the search string will look for log event messages matching:

(Source interface OR Destination interface) AND Advanced Switching category

- Step 6** Click **Apply Filters** to apply the filter immediately to the Log View table of event messages. Click **Reset Filters** to clear the filter and display the unfiltered results again.

Technical FAQ

How do I view the CAM table on the SonicWALL NSA 2400MX?

The SonicOS 5.7.0.0 user interface or CLI does not provide a way to display the CAM, or MAC Address, table directly, but provides the same information in the ARP table and on the Switching > L2 Discovery page.

A Content Addressable Memory (CAM) table is a dynamic, internal, purely Layer 2 mapping between switch ports and the MAC addresses that are bound to them. The CAM table information is also referred to as the MAC address table, switching cache, or forwarding data. The CAM table is used to quickly dereference MAC addresses to the switch ports where they are connected, allowing the speedy switching of traffic out the port to the destination. The CAM table is populated when the switch receives a data frame on one of its ports and updates the table with the frame's source MAC address and the port on which it was received.

In SonicOS 5.7.0.0, the information displayed on the the Switching > L2 Discovery page is derived from three sources:

- MAC address table, internal to the switch (SonicWALL NSA 2400MX)
- ARP table maintained by the gateway
- Layer 2 Discovery Protocol exchanges

To illustrate the difference between the MAC address table and the ARP table, consider a situation where you have two computers that use static IP addresses and communicate with each other within the same VLAN. The traffic between them never reaches the IP layer (the traffic is never forwarded, always switched).

These machines will only show up in the MAC address table of the switch. The Switching > L2 Discovery page will display the MAC addresses and VLAN for these computers, but nothing else (assuming there is no discovery protocol agent running on these machines).

If the machines stop communicating for awhile, the switch ages out the MAC address table and the entries will be gone. If you refresh the Switching > L2 Discovery page, you will no longer see these entries.

On the other hand, if the machines connect to the Internet or to another VLAN, the traffic will be forwarded and the gateway ARP table is populated with entries for these computers. It is possible for entries to exist only in the gateway ARP table, but not in the switch MAC address table.

The Switching > L2 Discovery page consolidates entries from the MAC address table and the ARP table, and displays one entry per machine.

Many switches, such as the HP ProCurve, Dell PowerConnect, or Cisco switches, provide a command to display the CAM or MAC Address table. For example, the following output is from a Cisco switch running IOS:

```
Cisco_L3# show mac-address-table dynamic
          Mac Address Table
-----
Vlan      Mac Address           Type           Ports
----      -
1         0017.c52e.59ba        DYNAMIC        Fa0/3
1         0017.c52e.5aa4        DYNAMIC        Fa0/4
1         0017.c53c.d425        DYNAMIC        Po1
Total Mac Addresses for this criterion: 3
Cisco_L3#
```

The display shows two dynamic entries for SonicPoint-Ns, connected to switch ports 3 and 4 of the Cisco switch, and one entry for the LACP Link Aggregation Group, which is connected to a SonicWALL NSA 2400MX and is not blocked by RSTP.

Glossary

BPDU	Bridge Protocol Data Unit – Used in RSTP, BPDUs are special data frames used to exchange information about bridge IDs and root path costs. BPDUs are exchanged every few seconds to allow switches to keep track of network topology and start or stop port forwarding.
bridge	A bridge is a data communications device that connects two Ethernet segments of a network together. A bridge operates by forwarding packets according to the destination Ethernet, or MAC, address, rather than by IP address. Because a bridge uses Layer 2, the data link layer, all protocols can be handled. Switches operate like bridges, but provide more ports for LAN connectivity and offer features that reduce collisions and latency on the network.
CoS	Class Of Service – Cos (IEEE 802.1p) defines eight different classes of service that are indicated in a 3-bit user_priority field in an IEEE 802.1Q header added to an Ethernet frame when using tagged frames on an 802.1 network.
DSCP	Differentiated Services Code Point – Also known as DiffServ, DSCP is a networking architecture that defines a simple, coarse-grained, class-based mechanism for classifying and managing network traffic and providing Quality of Service (QoS) guarantees on IP networks. RFC 2475, published in 1998 by the IETF, defines DSCP. DSCP operates by marking an 8-bit field in the IP packet header.
IETF	Internet Engineering Task Force – The IETF is an open standards organization that develops and promotes Internet standards.
L2	OSI Layer 2 (Ethernet) – Layer 2 of the seven layer OSI model is the Data Link Layer, on which the Ethernet protocol runs. Layer 2 is used to transfer data among network entities.
LACP	Link Aggregation Control Protocol – LACP is an IEEE specification that provides a way to combine multiple physical ports together to form a single logical channel. LACP allows load balancing by the connected devices.
LLDP	Link Layer Discovery Protocol (IEEE 802.1AB) – LLDP is a Layer 2 protocol used by network devices to communicate their identity, capabilities, and interconnections. This information is stored in a MIB database on each host, which can be queried with SNMP to determine the network topology. The information includes system name, port name, VLAN name, IP address, system capabilities (switching, routing), MAC address, link aggregation, and more.
LLTD	Link Layer Topology Discovery (Microsoft Standard) – LLTD is a Microsoft proprietary protocol with functionality similar to LLDP. It operates on wired or wireless networks (Ethernet 802.3 or wireless 802.11). LLTD is included on Windows Vista and Windows 7, and can be installed on Windows XP.
PDU	Protocol Data Unit – In the context of the Switching feature, the Layer 2 PDU is the frame. It contains the link layer header followed by the packet.
RSTP	Rapid Spanning Tree Protocol (IEEE 802.1D-2004) – RSTP was defined in 1998 as an improvement to Spanning Tree Protocol. It provides faster spanning tree convergence after a topology change.

Solution Document Version History

Version Number	Date	Notes
1	3/30/2010	This document was created by Susan Weigand
2	5/19/2010	Added conceptual information, more details about configuration, Troubleshooting and Technical FAQ sections. Embedded Port Mirroring screencast tutorial.
3	6/29/2010	Replaced embedded Port Mirroring screencast tutorial with final version. Added direct link to it as well.