



Switch 3812 and Switch 3824 Implementation Guide

3C17401, 3C17400

<http://www.3com.com/>

Part No. DUA1740-0BAA01
Published May 2003



3Com Corporation
5500 Great America
Parkway, Santa Clara,
California 95052-8145

Copyright © 2003, 3Com Corporation. All rights reserved. No part of this documentation may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without written permission from 3Com Corporation.

3Com Corporation reserves the right to revise this documentation and to make changes in content from time to time without obligation on the part of 3Com Corporation to provide notification of such revision or change.

3Com Corporation provides this documentation without warranty, term, or condition of any kind, either implied or expressed, including, but not limited to, the implied warranties, terms or conditions of merchantability, satisfactory quality, and fitness for a particular purpose. 3Com may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

If there is any software on removable media described in this documentation, it is furnished under a license agreement included with the product as a separate document, in the hard copy documentation, or on the removable media in a directory file named LICENSE.TXT or !LICENSE.TXT. If you are unable to locate a copy, please contact 3Com and a copy will be provided to you.

UNITED STATES GOVERNMENT LEGEND

If you are a United States government agency, then this documentation and the software described herein are provided to you subject to the following:

All technical data and computer software are commercial in nature and developed solely at private expense. Software is delivered as "Commercial Computer Software" as defined in DFARS 252.227-7014 (June 1995) or as a "commercial item" as defined in FAR 2.101(a) and as such is provided with only such rights as are provided in 3Com's standard commercial license for the Software. Technical data is provided with limited rights only as provided in DFAR 252.227-7015 (Nov 1995) or FAR 52.227-14 (June 1987), whichever is applicable. You agree not to remove or deface any portion of any legend provided on any licensed program or documentation contained in, or delivered to you in conjunction with, this User Guide.

Unless otherwise indicated, 3Com registered trademarks are registered in the United States and may or may not be registered in other countries.

3Com, the 3Com logo and SuperStack are registered trademarks of 3Com Corporation.

Intel and Pentium are registered trademarks of Intel Corporation. Microsoft, MS-DOS, Windows, and Windows NT are registered trademarks of Microsoft Corporation. Novell and NetWare are registered trademarks of Novell, Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd. Solaris is a registered trademark of Sun Microsystems.

IEEE and 802 are registered trademarks of the Institute of Electrical and Electronics Engineers, Inc.

All other company and product names may be trademarks of the respective companies with which they are associated.

ENVIRONMENTAL STATEMENT

It is the policy of 3Com Corporation to be environmentally-friendly in all operations. To uphold our policy, we are committed to:

Establishing environmental performance standards that comply with national legislation and regulations.

Conserving energy, materials and natural resources in all operations.

Reducing the waste generated by all operations. Ensuring that all waste conforms to recognized environmental standards. Maximizing the recyclable and reusable content of all products.

Ensuring that all products can be recycled, reused and disposed of safely.

Ensuring that all products are labelled according to recognized environmental standards.

Improving our environmental record on a continual basis.

End of Life Statement

3Com processes allow for the recovery, reclamation and safe disposal of all end-of-life electronic components.

Regulated Materials Statement

3Com products do not contain any hazardous or ozone-depleting material.

Environmental Statement about the Documentation

The documentation for this product is printed on paper that comes from sustainable, managed forests; it is fully biodegradable and recyclable, and is completely chlorine-free. The varnish is environmentally-friendly, and the inks are vegetable-based with a low heavy-metal content.

CONTENTS

ABOUT THIS GUIDE

Conventions	8
Related Documentation	9
Documentation Comments	9
Product Registration	10

1 SWITCH FEATURES OVERVIEW

What is Management Software?	11
Switch Features Explained	11
Automatic IP Configuration	12
Port Security	12
Aggregated Links	12
Auto-negotiation	12
Multicast Filtering	13
Spanning Tree Protocol and Rapid Spanning Tree Protocol	13
Switch Database	14
Traffic Prioritization	14
RMON	14
Broadcast Storm Control	15
VLANs	15
Configuration Save and Restore	15

2 OPTIMIZING BANDWIDTH

Port Features	17
Duplex	17
Flow Control	18
Auto-negotiation	18
Aggregated Links	19
How 802.3ad Link Aggregation Operates	19
Implementing 802.3ad Aggregated Links	20

Aggregated Links and Your Switch	21
Aggregated Link — Manual Configuration Example	24

3 USING MULTICAST FILTERING

What is an IP Multicast?	27
Benefits of Multicast	28
Multicast Filtering	28
Multicast Filtering and Your Switch	29
IGMP Multicast Filtering	30

4 USING RESILIENCE FEATURES

Resilience Feature Overview	34
Spanning Tree Protocol (STP)	34
Rapid Spanning Tree Protocol (RSTP)	35
What is STP?	35
How STP Works	37
STP Requirements	37
STP Calculation	38
STP Configuration	38
STP Reconfiguration	39
How RSTP Differs to STP	39
STP Example	39
STP Configurations	41
Using STP on a Network with Multiple VLANs	43

5 USING THE SWITCH DATABASE

What is the Switch Database?	45
How Switch Database Entries Get Added	45
Switch Database Entry States	46

6 USING TRAFFIC PRIORITIZATION

What is Traffic Prioritization?	48
How Traffic Prioritization Works	48
Traffic Classification	49
Traffic Marking	50

Traffic Re-Marking	52
Traffic Prioritization	52
Traffic Queues	56
Important QoS Considerations	56
Default QoS Configurations	58

7 STATUS MONITORING AND STATISTICS

RMON	59
What is RMON?	59
The RMON Groups	59
Benefits of RMON	61
RMON and the Switch	61
Alarm Events	62

8 SETTING UP VIRTUAL LANs

What are VLANs?	63
Benefits of VLANs	64
VLANs and Your Switch	65
The Default VLAN	65
Communication Between VLANs	65
Creating New VLANs	66
VLANs: Tagged and Untagged Membership	66
VLAN Configuration Examples	67
Using Untagged Connections	67
Using 802.1Q Tagged Connections	68

9 USING AUTOMATIC IP CONFIGURATION

How Your Switch Obtains IP Information	72
How Automatic IP Configuration Works	72
Automatic Process	72
Important Considerations	72
Server Support	73
Event Log Entries and Traps	73

A CONFIGURATION RULES

- Configuration Rules for Gigabit Ethernet 75
- Configuration Rules for Fast Ethernet 76
 - Configuration Rules with Full Duplex 77

B NETWORK CONFIGURATION EXAMPLES

- Simple Network Configuration Examples 80
 - Desktop Switch Example 80
- Advanced Network Configuration Examples 81
 - Improving the Resilience of Your Network 81
 - Enhancing the Performance of Your Network 82

C IP ADDRESSING

- IP Addresses 83
 - Simple Overview 83
 - Advanced Overview 84
- Subnets and Subnet Masks 86
- Default Gateways 88

GLOSSARY

INDEX

ABOUT THIS GUIDE

This guide describes the features of the 3Com Switch 3812 (12-port, Managed Gigabit) and 3Com Switch 3824 (24-port, Managed Gigabit). It outlines how to use these features to optimize the performance of your network.

The term Switch 3812 and Switch 3824 is used when referring to the 3Com Switch 3812 (12-port, Managed Gigabit) and 3824 (24-port, Managed Gigabit).

Refer to the Management Quick Reference Guide that accompanies your Switch for details of the specific features your Switch supports.

This guide is intended for the system or network administrator who is responsible for configuring, using, and managing the Switches. It assumes a working knowledge of local area network (LAN) operations and familiarity with communication protocols that are used to interconnect LANs.



For detailed descriptions of the Web interface operations and the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch or on the 3Com Web site.



If release notes are shipped with your product and the information there differs from the information in this guide, follow the instructions in the release notes.

Most user guides and release notes are available in Adobe Acrobat Reader Portable Document Format (PDF) or HTML on the 3Com World Wide Web site:

<http://www.3com.com/>

Conventions

[Table 1](#) and [Table 2](#) list conventions that are used throughout this guide.

Table 1 Notice Icons

Icon	Notice Type	Description
	Information note	Information that describes important features or instructions
	Caution	Information that alerts you to potential loss of data or potential damage to an application, system, or device
	Warning	Information that alerts you to potential personal injury

Table 2 Text Conventions

Convention	Description
Screen displays	This typeface represents information as it appears on the screen.
Syntax	<p>The word “syntax” means that you must evaluate the syntax provided and then supply the appropriate values for the placeholders that appear in angle brackets. Example:</p> <p>To change your password, use the following syntax:</p> <pre>system password <password></pre> <p>In this example, you must supply a password for <password>.</p>
Commands	<p>The word “command” means that you must enter the command exactly as shown and then press Return or Enter. Commands appear in bold. Example:</p> <p>To display port information, enter the following command:</p> <pre>bridge port detail</pre>
The words “enter” and “type”	When you see the word “enter” in this guide, you must type something, and then press Return or Enter. Do not press Return or Enter when an instruction simply says “type.”
Keyboard key names	<p>If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example:</p> <pre>Press Ctrl+Alt+Del</pre>
Words in <i>italics</i>	<p>Italics are used to:</p> <ul style="list-style-type: none"> ■ Emphasize a point. ■ Denote a new term at the place where it is defined in the text. ■ Identify menu names, menu commands, and software button names. Examples: <p>From the <i>Help</i> menu, select <i>Contents</i>.</p> <p>Click <i>OK</i>.</p>

Related Documentation

In addition to this guide, each Switch documentation set includes the following:

- *Switch 3812 and Switch 3824 Getting Started Guide*

This guide contains:

- all the information you need to install and set up the Switch in its default state
- information on how to access the management software to begin managing the Switch.

- *Switch 3812 and Switch 3824 Management Interface Reference Guide*

This guide provides detailed information about the Web interface and Command Line Interface that enable you to manage the Switch. It is supplied in HTML format on the CD-ROM that accompanies the Switch.

- *Switch 3812 and Switch 3824 Management Quick Reference Guide*

This guide contains:

- a list of the features supported by the Switch.
- a summary of the Web interface and Command Line Interface commands for the Switch.

- *Release Notes*

These notes provide information about the current software release, including new features, modifications, and known problems.

There are other publications you may find useful, such as:

- Documentation accompanying 3Com Network Supervisor. This is supplied on the CD-ROM that accompanies the Switch.

Documentation Comments

Your suggestions are very important to us. They will help make our documentation more useful to you. Please e-mail comments about this document to 3Com at:

pddtechpubs_comments@3com.com

Please include the following information when contacting us:

- Document title

- Document part number (on the title page)
- Page number (if appropriate)

Example:

- Switch 3812 and Switch 3824 Implementation Guide
- Part number: DUA1740-0BAA01
- Page 25



Please note that we can only respond to comments and questions about 3Com product documentation at this e-mail address. Questions related to technical support or sales should be directed in the first instance to your network supplier.

Product Registration

You can now register your Switch on the 3Com Web site:

<http://www.3com.com/register/>

You will need your product part number (3Cxxxxx), product serial number and date and place of purchase to register your 3Com product.

Registering your product enables you to: process Repair Requests on-line, check the status of your requests at anytime, provides you with important warranty information as well as activating your entitlement to additional service benefits and receive up-to-date information on your product.

1

SWITCH FEATURES OVERVIEW

This chapter contains introductory information about the Switch 3812 and Switch 3824 management software and supported features. It covers the following topics:

- [What is Management Software?](#)
- [Switch Features Explained](#)



For detailed descriptions of the Web interface operations and the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch.

What is Management Software?

Your Switch can operate in its default state. However, to make full use of the features offered by the Switch, and to change and monitor the way it works, you have to access the management software that resides on the Switch. This is known as managing the Switch.

Managing the Switch can help you to improve its efficiency and therefore the overall performance of your network.

There are several different methods of accessing the management software to manage the Switch. These methods are explained in Chapter 3 of the Getting Started Guide that accompanies your Switch.

Switch Features Explained

The management software provides you with the capability to change the default state of some of the Switch features. This section provides a brief overview of these features — their applications are explained in more detail later in this guide.



For a list of the features supported by your Switch, please refer to the Management Quick Reference Guide that accompanies your Switch.

Automatic IP Configuration

Your Switch can have its IP information automatically configured using a DHCP server. Alternatively, you can manually configure the IP information.



For more information about how the automatic IP configuration feature works, see [Chapter 9 “Using Automatic IP Configuration”](#).

Port Security

Your Switch supports the following port security modes, which you can set for an individual port or a range of ports:

- **No Security**

Port security is disabled and all network traffic is forwarded through the port without any restrictions.

- **Secure**

All currently learnt addresses on the port are made permanent. Any packets containing a source address not learnt on the port will be dropped.



The maximum number of permanent addresses on the Switch is 1000.

Aggregated Links

Aggregated links are connections that allow devices to communicate using up to eight links in parallel. Aggregated links provide two benefits:

- They can potentially increase the bandwidth of a connection.
- They can provide redundancy — if one link is broken, the other links share the traffic for that link.



For more information about aggregated links, see [Chapter 2 “Optimizing Bandwidth”](#).

Auto-negotiation

Auto-negotiation allows ports to auto-negotiate port speed, duplex-mode (only at 10 Mbps and 100 Mbps) and flow control. When auto-negotiation is enabled (default), a port “advertises” its maximum capabilities — these capabilities are by default the parameters that provide the highest performance supported by the port.



SFP ports do not support auto-negotiation of port speed.



Ports operating at 1000 Mbps only support full duplex mode.



For details of the auto-negotiation features supported by your Switch, please refer to the Management Quick Reference Guide that accompanies your Switch.

Duplex

Full duplex mode allows packets to be transmitted and received simultaneously and, in effect, doubles the potential throughput of a link.

Flow Control

All Switch ports support flow control, which is a mechanism that minimizes packet loss during periods of congestion on the network.

Flow control is supported on ports operating in half duplex mode, and is implemented using the IEEE Std 802.3-2002 (incorporating 802.3x) on ports operating in full duplex mode.



For more information about auto-negotiation and port capabilities, see [Chapter 2 “Optimizing Bandwidth”](#).

Multicast Filtering

Multicast filtering allows the Switch to forward multicast traffic to only the endstations that are part of a predefined multicast group, rather than broadcasting the traffic to the whole network.

The multicast filtering system supported by your Switch uses IGMP (Internet Group Management Protocol) snooping to detect the endstations in each multicast group to which multicast traffic should be forwarded.



For more information about multicast filtering, see [Chapter 3 “Using Multicast Filtering”](#).

Spanning Tree Protocol and Rapid Spanning Tree Protocol

Spanning Tree Protocol (STP) and Rapid Spanning Tree Protocol (RSTP) are bridge-based systems that make your network more resilient to link failure and also provide protection from network loops — one of the major causes of broadcast storms.

STP allows you to implement alternative paths for network traffic in the event of path failure and uses a loop-detection process to:

- Discover the efficiency of each path.
- Enable the most efficient path.

- Disable the less efficient paths.
- Enable one of the less efficient paths if the most efficient path fails.

RSTP is an enhanced version of the STP feature and is enabled by default. RSTP can restore a network connection quicker than the legacy STP feature. RSTP can detect if it is connected to a legacy device that only supports IEEE 802.1D STP and will automatically downgrade to STP on that particular port.

STP conforms to the IEEE Std 802.1D, 1998 Edition and RSTP conforms to the IEEE Std 802.1w-2001.



For more information about STP and RSTP, see [Chapter 4 “Using Resilience Features”](#).

Switch Database

The Switch Database is an integral part of the Switch and is used by the Switch to determine if a packet should be forwarded, and which port should transmit the packet if it is to be forwarded.



For more information about the Switch Database, see [Chapter 5 “Using the Switch Database”](#).

Traffic Prioritization

Using the traffic prioritization capabilities of your Switch provides Quality of Service (QoS) to your network through increased reliability of data delivery. You can prioritize traffic on your network to ensure that high priority data is transmitted with minimum delay.



For more information about traffic prioritization, see [Chapter 6 “Using Traffic Prioritization”](#).

RMON

Remote Monitoring (RMON) is an industry standard feature for traffic monitoring and collecting network statistics. The Switch software continually collects statistics about the LAN segments connected to the Switch. If you have a management workstation with an RMON management application, the Switch can transfer these statistics to your workstation on request or when a pre-defined threshold is exceeded.



For more information about RMON and Event Notification, see [Chapter 7 “Status Monitoring and Statistics”](#).

Broadcast Storm Control

Broadcast Storm Control is a system that monitors the level of broadcast traffic on that port. If the broadcast traffic level rises to a pre-defined number of frames per second (threshold), the broadcast traffic on the port is blocked until the broadcast traffic level drops below the threshold. This system prevents the overwhelming broadcast traffic that can result from network equipment which is faulty or configured incorrectly.

VLANs

A Virtual LAN (VLAN) is a flexible group of devices that can be located anywhere in a network, but which communicate as if they are on the same physical segment. With VLANs, you can segment your network without being restricted by physical connections — a limitation of traditional network design. As an example, with VLANs you can segment your network according to:

- Departmental groups
- Hierarchical groups
- Usage groups



For more information about VLANs, see [Chapter 8 “Setting Up Virtual LANs”](#).

Configuration Save and Restore

The Configuration Save and Restore feature allows the configuration of your Switch to be saved as a file on a remote server, or to be restored onto the Switch from a remote file. The configuration information is stored in a readable ASCII text file.

All configuration information that can be set using the Switch's Command Line Interface is saved and restored.

You must have *security* management access level to be able to save and restore the Switch configuration.



The password is stored in the backup file. If you want to restore the file but don't have the password you can delete the password from the file.

Important Considerations

- 3Com recommends the Switch unit is reset to its factory default settings before you restore a configuration onto it. You can reset the Switch using the **system control initialize** CLI command or the *System > Control > Initialize* Web interface operation.

- The configuration can only be restored onto a device which has the same physical connections and configuration, as when the configuration was initially saved. The restore operation will be unsuccessful if the physical configuration of the device is different.
- The configuration of the Switch must only be restored or saved by a single user at a time.
- When using the Configuration Save and Restore feature, 3Com recommends that aggregated links are configured as either:
 - Manual aggregations with Link Aggregation Configuration Protocol (LACP) disabled on the ports that are to be manually placed in the aggregated link.

or

- LACP automatic aggregations — that is, LACP enabled on all ports and the aggregated links created automatically. The aggregated link should be enabled and Spanning Tree Protocol enabled.

Parameters such as VLANs and Fast Start may be set up as required.

Other combinations of port settings, however, are not recommended as Configuration Restore will only perform a “best effort” restore of the configuration. For example, LACP automatic aggregations with manually defined ports are restored as manual aggregations with manual ports. LACP automatic aggregations with automatic ports where the aggregated link is disabled and Spanning Tree Protocol is disabled are restored as manual aggregations with the aggregated link disabled.



For further information about LACP, see [Chapter 2 “Optimizing Bandwidth”](#).

- When restoring a configuration onto a unit over an aggregated link, communication with that unit may be lost because the restore operation disables the aggregated link ports. Communication over the aggregated links is re-established when the restore operation has been completed.



For detailed descriptions of the Configuration Save and Restore Web interface operations and Command Line Interface (CLI) commands, please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch.

2

OPTIMIZING BANDWIDTH

There are many ways you can optimize the bandwidth on your network and improve network performance. If you utilize certain Switch features you can provide the following benefits to your network and end users:

- Increased bandwidth
- Quicker connections
- Faster transfer of data
- Minimized data errors
- Reduced network downtime



For detailed descriptions of the Web interface operations and the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch.

Port Features

The default state for all the features detailed below provides the best configuration for most users. *In normal operation, you do not need to alter the Switch from its default state.* However, under certain conditions you may wish to alter the default state of these ports, for example, if you are connecting to old equipment that does not comply with the IEEE 802.3x standard.

Duplex Full duplex allows packets to be transmitted and received simultaneously and, in effect, doubles the potential throughput of a link. Half duplex only allows packets to be transmitted or received at any one time.

To communicate effectively, both devices at either end of a link *must* use the same duplex mode. If the devices at either end of a link support auto-negotiation, this is done automatically. If the devices at either end of

a link do not support auto-negotiation, both ends must be manually set to full duplex or half duplex accordingly.



Ports operating at 1000 Mbps support full duplex mode only.

Flow Control

All Switch ports support flow control, which is a mechanism that prevents packet loss during periods of congestion on the network. Packet loss is caused by one or more devices sending traffic to an already overloaded port on the Switch. Flow control prevents packet loss by inhibiting the transmitting port from generating more packets until the period of congestion ends.

Flow control is implemented using the IEEE Std 802.3-2002 (incorporating 802.3x) for ports operating in full duplex mode, and Intelligent Flow Management (IFM) for ports operating in half duplex mode.

Auto-negotiation

Auto-negotiation allows ports to automatically determine the best port speed, duplex-mode (only at 10 Mbps and 100 Mbps) and flow control. When auto-negotiation is enabled (default), a port “advertises” its maximum capabilities — these capabilities are by default the parameters that provide the highest performance supported by the port.

You can modify the capabilities that a port “advertises” on a per port basis, dependant on the type of port.

You can disable auto-negotiation for the whole Switch, or per port. You can also modify the capabilities that a port “advertises” on a per port basis, dependant on the type of port.



SFP ports do not support auto-negotiation of port speed.



Ports operating at 1000 Mbps support full duplex mode only.



If auto-negotiation is disabled, the auto-MDIX feature does not operate on the ports. Therefore the correct cables, that is, cross-over or straight-through need to be used. For more information, see the Getting Started Guide that accompanies your Switch.



Ports at both ends of the link should be set to auto-negotiate.

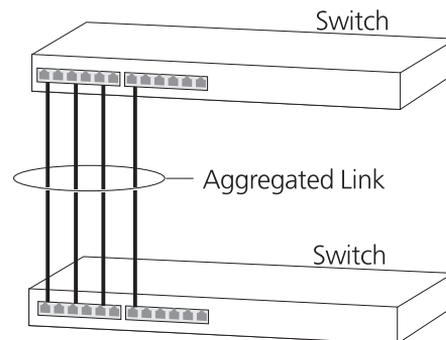
Aggregated Links

Aggregated links are connections that allow devices to communicate using up to eight member links in parallel. Aggregated links provide the following benefits:

- They can potentially increase the bandwidth of a connection. The capacity of the multiple links is combined into one logical link.
- They can provide redundancy — if one link is broken, the other links share the traffic for that link.

[Figure 1](#) shows two Switches connected using an aggregated link containing four member links. If all ports on both Switch units are configured as 1000BASE-TX and they are operating in full duplex, the potential maximum bandwidth of the connection is 8 Gbps.

Figure 1 Switch units connected using an aggregated link



How 802.3ad Link Aggregation Operates

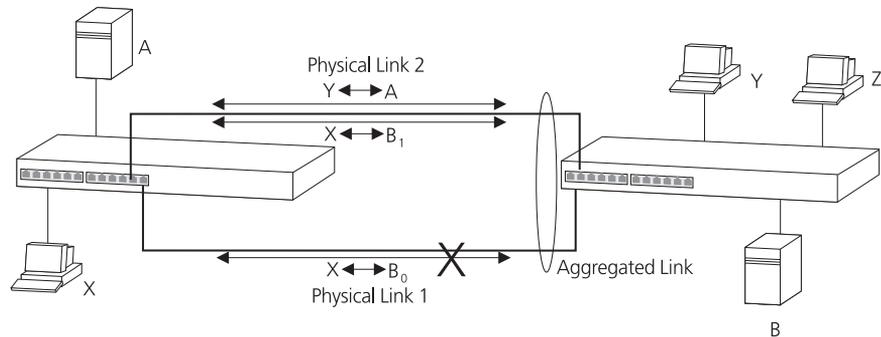
Your Switch supports IEEE Std 802.3-2002 (incorporating 802.3ad) aggregated links which use the Link Aggregation Control Protocol (LACP). LACP provides automatic, point-to-point redundancy between two devices (switch-to-switch or switch-to-server) that have full duplex connections operating at the same speed.

By default, LACP is disabled on all Switch ports.

If a member link in an aggregated link fails, the traffic using that link is dynamically reassigned to the remaining member links in the aggregated link. [Figure 2](#) shows the simplest case: two member links, that is the physical links, form an aggregated link. In this example, if link 1 fails, the data flow between X and B is remapped to physical link 2. The re-mapping occurs as soon as the Switch detects that a member link has

failed — almost instantaneously. As a result, aggregated link configurations are extremely resilient and fault-tolerant.

Figure 2 Dynamic Reassignment of Traffic Flows



The key benefits of 802.3ad link aggregation are:

- Automatic configuration — network management does not need to be used to manually aggregate links.
- Rapid configuration and reconfiguration — approximately one to three seconds.
- Compatibility — non-802.3ad devices can interoperate with 802.3ad enabled devices. However, you will need to manually configure the aggregated links as LACP will not be able to automatically detect and form an aggregation with a non-802.3ad device.
- The operation of 802.3ad can be configured and managed via network management.

Implementing 802.3ad Aggregated Links

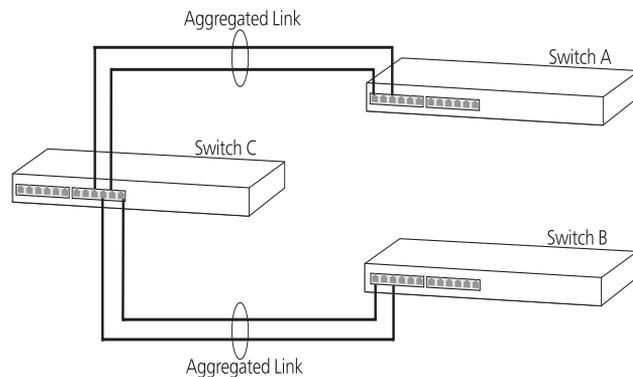
LACP can be enabled or disabled on a per port basis. You can implement 802.3ad aggregated links in two ways:

- Manual Aggregations — You can manually add and remove ports to and from an aggregated link via Web commands. However, if a port has LACP enabled, and if a more appropriate or correct automatic membership is detected by LACP, it will override the manual configuration.

For example, in [Figure 3](#), if a port on Switch C is physically connected to Switch B, but you manually configure the port on Switch C to be a

member of an aggregated link for Switch A in error, LACP (if it is enabled) will detect this and place the port in the aggregated link for Switch B, thus overriding the manual configuration.

Figure 3 Aggregated Link — Example



- **LACP Automatic Aggregations** — If LACP detects at least two active ports sharing the same partner device, and if no matching pre-configured aggregated links exist, LACP will automatically assign a free un-configured aggregated link to form an aggregated link with the partner device.

If you have an existing single port connection between two devices, this automatic behavior allows quick and easy addition of extra bandwidth by simply adding an extra physical link between the units.

The Spanning Tree costs for a port running LACP is the cost assigned for an aggregated link running at that speed. As required by the IEEE Std 802.3-2002 (incorporating 802.3ad), no changes in cost are made according to the number of member links in the aggregated link.

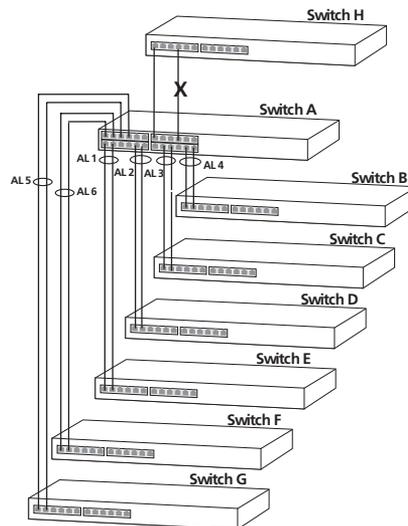
Aggregated Links and Your Switch

- When any port is assigned to an aggregated link (either manually or via LACP) it will adopt the configuration settings of the aggregated link. When a port leaves an aggregated link its original configuration settings are restored.
- A maximum of six active aggregations can be created. A maximum of up to eight ports may be added manually to any individual aggregation, or via LACP. There are however a few points to consider:

- The Switch only supports a maximum of eight active ports in any individual aggregation.
- If multiple links are connected between a unit and more than six other devices as shown in [Figure 4](#), only six of the devices will be assigned to aggregated links. The remaining devices will each only have one link made *active*, that is, passing data. All other links will be made *inactive* to prevent loops occurring.

LACP detects if one of the existing six aggregated links is removed and will then automatically assign one of the remaining devices to the aggregated link that has become free.

Figure 4 How LACP works on a Switch with multiple aggregated links



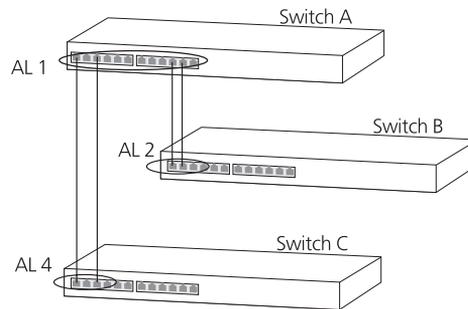
- When multiple links of different speed connect two devices only the highest speed links will be aggregated.
- A LinkUp / LinkDown trap will only be sent for individual links. The Traps will not be sent for an aggregation.

When setting up an aggregated link, note that:

- The ports at both ends of a member link must be configured as members of an aggregated link, if you are manually configuring aggregated links.
- A member link port can only belong to one aggregated link.

- The member link ports can be mixed media, that is fiber and/or twisted pair ports within the same aggregated link.
- The member link ports must have the same configuration.
- Member links must retain the same groupings at both ends of an aggregated link. For example, the configuration in [Figure 5](#) will not work as Switch A has one aggregated link defined whose member links are then split between two aggregated links defined on Switches B and C. Note that this illegal configuration could not occur if LACP is enabled.

Figure 5 An illegal aggregated link configuration



To make this configuration work you need to have two aggregated links defined on Switch A, one containing the member links for Switch B and one containing the member links for Switch C.

When using an aggregated link, note that:

- To gather statistics about an aggregated link, you must add together the statistics for each port in the aggregated link.
- If you wish to disable a single member link of an aggregated link, you must first physically remove the connection to ensure that you do not lose any traffic, before you disable both ends of the member link separately. If you do this, the traffic destined for that link is distributed to the other links in the aggregated link.

If you do not remove the connection and only disable one end of the member link port, traffic is still forwarded to that port by the aggregated link port at the other end. This means that a significant amount of traffic may be lost.

- Before removing an entire aggregated link, you must disable all the aggregated link ports or disconnect all the links, except one — if you do not, a loop may be created.
- When manually creating an aggregated link between two devices, the ports in the aggregated link must not be physically connected together until the aggregated link has been correctly configured at both ends of the link. Failure to configure the aggregated link at both ends before physically connecting the ports can result in a number of serious network issues such as lost packets and network loops.

Traffic Distribution and Link Failure on Aggregated Links

To maximize throughput, all traffic is distributed across the individual links that make up an aggregated link. Therefore, when a packet is made available for transmission down an aggregated link, a hardware-based traffic distribution mechanism determines which particular port in the link should be used. The traffic is distributed among the member links as efficiently as possible.

To avoid the potential problem of out-of-sequence packets (or “packet re-ordering”), the Switch ensures that all the conversations between a given pair of endstations will pass through the same port in the aggregated link. Single-to-multiple endstation conversations, on the other hand, may still take place over different ports.

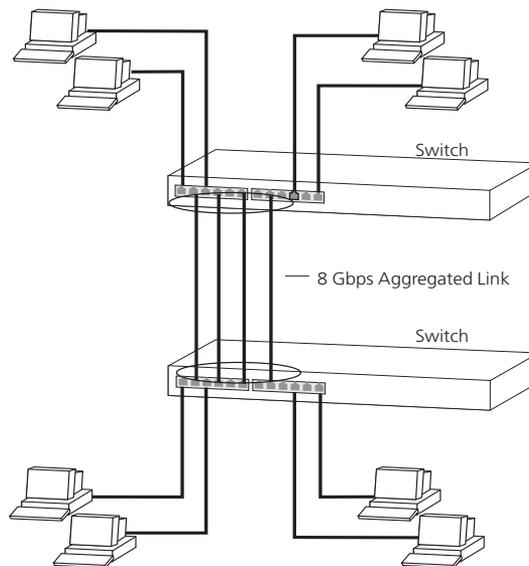
If the link state on any of the ports in an aggregated link becomes inactive due to link failure, then the Switch will automatically redirect the aggregated link traffic to the remaining ports. Aggregated links therefore provide built-in resilience for your network.

The Switch also has a mechanism to prevent the possible occurrence of packet re-ordering when a link recovers too soon after a failure.

Aggregated Link — Manual Configuration Example

The example shown in [Figure 6](#) illustrates an 8 Gbps aggregated link between two Switch units, (that is, each port is operating at 1000 Mbps, full duplex).

Figure 6 An 8 Gbps aggregated link between two Switch units



To manually set up this configuration:

- 1** Prepare ports 2, 4, 6 and 8 on the upper Switch for aggregated links. To do this:
 - a** Check that the ports have an identical configuration using your preferred management interface.
 - b** Add the ports 2, 4, 6 and 8 on the specified unit to the aggregated link.
- 2** Prepare ports 2, 4, 6 and 8 on the lower Switch for aggregated links. To do this:
 - a** Check that the ports have an identical configuration using your preferred management interface.
 - b** Add the ports 2, 4, 6 and 8 on the specified unit to the aggregated link.
- 3** Connect port 2 on the upper Switch to port 2 on the lower Switch.
- 4** Connect port 4 on the upper Switch to port 4 on the lower Switch.
- 5** Connect port 6 on the upper Switch to port 6 on the lower Switch.
- 6** Connect port 8 on the upper Switch to port 8 on the lower Switch.

3

USING MULTICAST FILTERING

Multicast filtering improves the performance of networks that carry multicast traffic.

This chapter explains multicasts, multicast filtering, and how multicast filtering can be implemented on your Switch. It covers the following topics:

- [What is an IP Multicast?](#)
- [Multicast Filtering](#)
- [IGMP Multicast Filtering](#)



For detailed descriptions of the Web interface operations and the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch.

What is an IP Multicast?

A *multicast* is a packet that is intended for “one-to-many” and “many-to-many” communication. Users explicitly request to participate in the communication by joining an endstation to a specific multicast group. If the network is set up correctly, a multicast can only be sent to an endstation or a subset of endstations in a LAN, or VLAN, that belong to the relevant multicast group.

Multicast group members can be distributed across multiple subnetworks; thus, multicast transmissions can occur within a campus LAN or over a WAN. In addition, networks that support IP multicast send only *one* copy of the desired information across the network until the delivery path that reaches group members diverges. It is only at these points that multicast packets are replicated and forwarded, which makes efficient use of network bandwidth.

A multicast packet is identified by the presence of a multicast group address in the destination address field of the packet's IP header.

Benefits of Multicast

The benefits of using IP multicast are that it:

- Enables the simultaneous delivery of information to many receivers in the most efficient, logical way.
- Reduces the load on the source (for example, a server) because it does not have to produce multiple copies of the same data.
- Makes efficient use of network bandwidth and scales well as the number of participants or collaborators expands.
- Works with other IP protocols and services, such as Quality of Service (QoS).

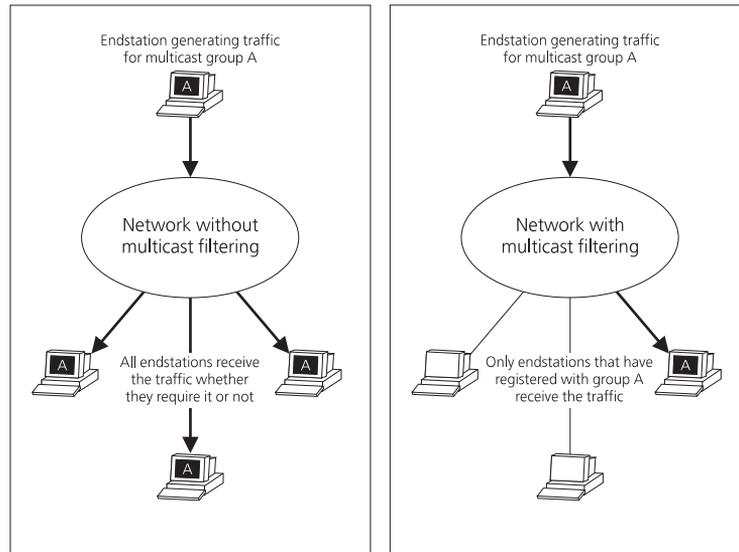
There are situations where a multicast approach is more logical and efficient than a unicast approach. Application examples include distance learning, transmitting stock quotes to brokers, and collaborative computing.

A typical use of multicasts is in video-conferencing, where high volumes of traffic need to be sent to several endstations simultaneously, but where broadcasting that traffic to all endstations would seriously reduce network performance.

Multicast Filtering

Multicast filtering is the process that ensures that endstations only receive multicast traffic if they register to join specific multicast groups. With multicast filtering, network devices only forward multicast traffic to the ports that are connected to registered endstations.

[Figure 7](#) shows how a network behaves without multicast filtering and with multicast filtering.

Figure 7 The effect of multicast filtering

Multicast Filtering and Your Switch

Your Switch provides automatic multicast filtering support using IGMP (Internet Group Management Protocol) Snooping. It also supports IGMP query mode.

Snooping Mode

Snooping Mode allows your Switch to forward multicast packets only to the appropriate ports. The Switch “snoops” on exchanges between endstations and an IGMP device, typically a router, to find out the ports that wish to join a multicast group and then sets its filters accordingly.

Query Mode

Query mode allows the Switch to function as the Querier if it has the lowest IP address in the subnetwork to which it belongs.

IGMP querying is disabled by default on the Switch 3812 and Switch 3824. This helps prevent interoperability issues with core products that may not follow the lowest IP address election method.

You can enable or disable IGMP query mode for the Switch using the *Bridge > Multicast filter > IGMP > Querymode* operation on the Web Interface.

You would enable query mode if you wish to run multicast sessions in a network that does not contain any IGMP routers (or queriers). This command will configure the Switch to automatically negotiate with compatible devices on VLAN 1 to become the querier.



*The Switch 3812 and Switch 3824 are compatible with any device that conforms to the IGMP v2 protocol. The Switch does not support IGMP v3. If you have an IGMP v3 network, you should disable IGMP snooping for the Switch using the **snoopMode** command on the Web Interface .*

IGMP Multicast Filtering

IGMP is the system that all IP-supporting network devices use to register endstations with multicast groups. It can be used on all LANs and VLANs that contain a multicast capable IP router and on other network devices that support IP.

IGMP multicast filtering works as follows:

- 1 The IP router (or querier) periodically sends *query* packets to all the endstations in the LANs or VLANs that are connected to it.
If your network has more than one IP router, then the one with the lowest IP address becomes the querier. The Switch can be the IGMP querier and will become so if its own IP address is lower than that of any other IGMP queriers connected to the LAN or VLAN. However, as the Switch only has an IP address on its default VLAN, the Switch will only ever query on the default VLAN (VLAN1). Therefore, if there are no other queriers on other VLANs, the IP multicast traffic will not be forwarded on them.
- 2 When an IP endstation receives a query packet, it sends a *report* packet back that identifies the multicast group that the endstation would like to join.
- 3 When the report packet arrives at a port on a Switch with *IGMP multicast learning* enabled, the Switch learns that the port is to forward traffic for the multicast group and then forwards the packet to the router.
- 4 When the router receives the report packet, it registers that the LAN or VLAN requires traffic for the multicast groups.
- 5 When the router forwards traffic for the multicast group to the LAN or VLAN, the Switch units only forward the traffic to ports that received a report packet.

Enabling IGMP Multicast Learning

You can enable or disable multicast learning and IGMP querying using the `snoopMode` command on the Web interface. For more information about enabling IGMP multicast learning, please refer to the Management Interface Reference Guide supplied on your Switch CD-ROM.

If IGMP multicast learning is not enabled then IP multicast traffic is always forwarded, that is, it floods the network.



For information about configuring IGMP functionality on an endstation, refer to the user documentation supplied with your endstation or the endstation's Network Interface Card (NIC).

4

USING RESILIENCE FEATURES

Setting up resilience on your network helps protect critical links against failure, protects against network loops, and reduces network downtime to a minimum.

This chapter explains the features supported by the Switch that provide resilience for your network. It covers the following topics:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP) — an enhanced version of the STP feature.



For detailed descriptions of the Web interface operations and the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch.

Resilience Feature Overview

Table 3 lists the key differences between each feature, so you can evaluate the benefits of each to determine which feature is most suitable for your network.

Table 3 Spanning Tree Protocols — Key Differences

Spanning Tree Protocol	Rapid Spanning Tree Protocol
STP is disabled by default. User enables STP on each Switch.	RSTP is enabled by default.
Automatic configuration.	Automatic configuration.
Up to 30 second delay on link failure to restoring a network connection.	Within 5 seconds restores a network connection.



3Com recommends that you use the Rapid Spanning Tree Protocol feature (default enabled) to provide optimum performance for your network and ease of use.

The Switch also supports aggregated links which increase bandwidth and also provide resilience against individual link failure. Aggregated links will operate with STP enabled. For more information, see [Aggregated Links](#) on [page 19](#).

Spanning Tree Protocol (STP)

The Spanning Tree Protocol (STP) makes your network more resilient to link failure and also provides a protection from loops — one of the major causes of broadcast storms. STP is enabled by default on your Switch.



To be fully effective, STP must be enabled on all Switches in your network.



RSTP provides the same functionality as STP. For details on how the two systems differ, see [“How RSTP Differs to STP”](#) on [page 39](#).

The following sections explain more about STP and the protocol features supported by your Switch. They cover the following topics:

- [What is STP?](#)
- [How STP Works](#)
- [Using STP on a Network with Multiple VLANs](#)



The protocol is a part of the IEEE Std 802.1D, 1998 Edition bridge specification. To explain STP more effectively, your Switch will be referred to as a bridge.

Rapid Spanning Tree Protocol (RSTP)

The Rapid Spanning Tree (RSTP) is an enhanced Spanning Tree feature. RSTP implements the Spanning Tree Algorithm and Protocol, as defined in the IEEE Std 802.1w-2001.

Some of the benefits of RSTP are:

- Faster determination of the Active Spanning Tree topology throughout a bridged network.
- Support for bridges with more than 256 ports.
- Support for the Fast-Forwarding configuration of edge ports provided by the 'Fast Start' feature. Fast Start allows a port that is connected to an endstation to begin forwarding traffic after only 4 seconds. During this 4 seconds RSTP (or STP) will detect any misconfiguration that may cause a temporary loop and react accordingly.
- Easy deployment throughout a legacy network, through backward compatibility:
 - it will default to sending 802.1D style BPDU's on a port if it receives packets of this format.
 - it is possible for some ports on a Switch to operate in RSTP (802.1w) mode, and other ports, for example those connected to a legacy Switch, to operate in STP (802.1D) mode.
 - you have an option to force your Switch to use the legacy 802.1D version of Spanning Tree, if required.

What is STP?

STP (802.1D) is a bridge-based system that allows you to implement parallel paths for network traffic and uses a loop-detection process to:

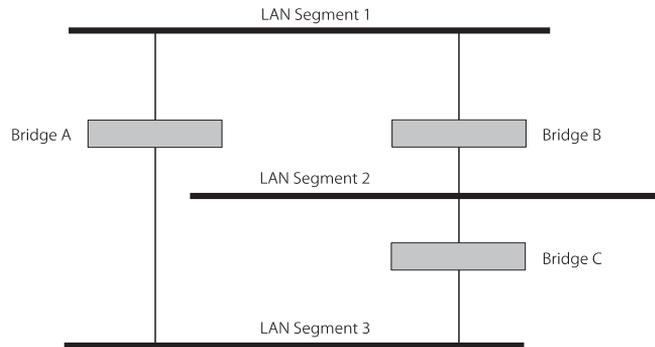
- Find and disable the less efficient paths (that is, the paths that have a lower bandwidth).
- Enable one of the less efficient paths if the most efficient path fails.



RSTP provides the same functionality as STP. For details on how the two systems differ, see ["How RSTP Differs to STP"](#) on [page 39](#).

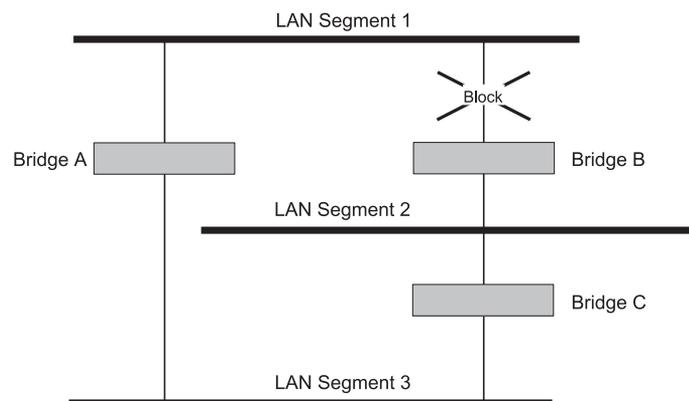
As an example, [Figure 8](#) shows a network containing three LAN segments separated by three bridges. With this configuration, each segment can communicate with the others using two paths. Without STP enabled, this configuration creates loops that cause the network to overload.

Figure 8 A network configuration that creates loops



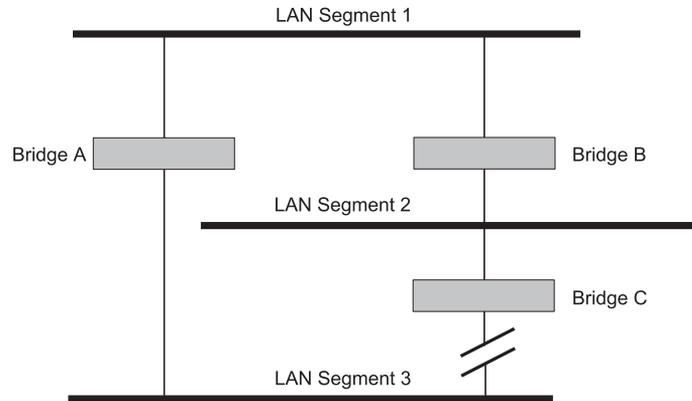
[Figure 9](#) shows the result of enabling STP on the bridges in the configuration. STP detects the duplicate paths and prevents, or *blocks*, one of them from forwarding traffic, so this configuration will work satisfactorily. STP has determined that traffic from LAN segment 2 to LAN segment 1 can only flow through Bridges C and A, because, for example, this path has a greater bandwidth and is therefore more efficient.

Figure 9 Traffic flowing through Bridges C and A



If a link failure is detected, as shown in [Figure 10](#), the STP process reconfigures the network so that traffic from LAN segment 2 flows through Bridge B.

Figure 10 Traffic flowing through Bridge B



STP determines which is the most efficient path between each bridged segment and a specifically assigned reference point on the network. Once the most efficient path has been determined, all other paths are blocked. Therefore, in [Figure 8](#), [Figure 9](#), and [Figure 10](#), STP initially determined that the path through Bridge C was the most efficient, and so blocked the path through Bridge B. After the failure of Bridge C, STP re-evaluated the situation and opened the path through Bridge B.

How STP Works

When enabled, STP determines the most appropriate path for traffic through a network. It does this as outlined in the sections below.

STP Requirements

Before it can configure the network, the STP system requires:

- Communication between all the bridges. This communication is carried out using Bridge Protocol Data Units (BPDUs), which are transmitted in packets with a known multicast address.
- Each bridge to have a Bridge Identifier. This specifies which bridge acts as the central reference point, or Root Bridge, for the STP system — the lower the Bridge Identifier, the more likely the bridge is to become the Root Bridge. The Bridge Identifier is calculated using the MAC address of the bridge and a priority defined for the bridge. The default priority of your Switch is 32768.

- Each port to have a cost. This specifies the efficiency of each link, usually determined by the bandwidth of the link — the higher the cost, the less efficient the link. [Table 4](#) shows the default port costs for a Switch.

Table 4 Default port costs

Port Speed	Link Type	Path Cost 802.1D, 1998 Edition	Path Cost 802.1w-2001
10 Mbps	Half Duplex	100	2,000,000
	Full Duplex	95	1,999,999
	Aggregated Link	90	1,000,000
100 Mbps	Half Duplex	19	200,000
	Full Duplex	18	199,999
	Aggregated Link	15	100,000
1000 Mbps	Full Duplex	4	20,000
	Aggregated Link	3	10,000

STP Calculation

The first stage in the STP process is the calculation stage. During this stage, each bridge on the network transmits BPDUs that allow the system to work out:

- The identity of the bridge that is to be the Root Bridge. The Root Bridge is the central reference point from which the network is configured.
- The Root Path Costs for each bridge — that is, the cost of the paths from each bridge to the Root Bridge.
- The identity of the port on each bridge that is to be the Root Port. The Root Port is the one that is connected to the Root Bridge using the most efficient path, that is, the one that has the lowest Root Path Cost. Note that the Root Bridge does not have a Root Port.
- The identity of the bridge that is to be the Designated Bridge of each LAN segment. The Designated Bridge is the one that has the lowest Root Path Cost from that segment. Note that if several bridges have the same Root Path Cost, the one with the lowest Bridge Identifier becomes the Designated Bridge.

All traffic destined to pass in the direction of the Root Bridge flows through the Designated Bridge. The port on this bridge that connects to the segment is called the Designated Bridge Port.

STP Configuration

After all the bridges on the network have agreed on the identity of the Root Bridge, and have established the other relevant parameters, each

bridge is configured to forward traffic only between its Root Port and the Designated Bridge Ports for the respective network segments. All other ports are blocked, which means that they are prevented from receiving or forwarding traffic.

STP Reconfiguration

Once the network topology is stable, all the bridges listen for Hello BPDUs transmitted from the Root Bridge at regular intervals. If a bridge does not receive a Hello BPDU after a certain interval (the Max Age time), the bridge assumes that the Root Bridge, or a link between itself and the Root Bridge, has gone down. The bridge then reconfigures the network to cater for the change. If you have configured an SNMP trap destination, when the topology of your network changes, the first bridge to detect the change sends out an SNMP trap.



CAUTION: *Network loops can occur if aggregated links are manually configured incorrectly, that is, the physical connections do not match the assignment of ports to an aggregated link. RSTP and STP may not detect these loops. So that RSTP and STP can detect all network loops you must ensure that all aggregated links are configured correctly.*

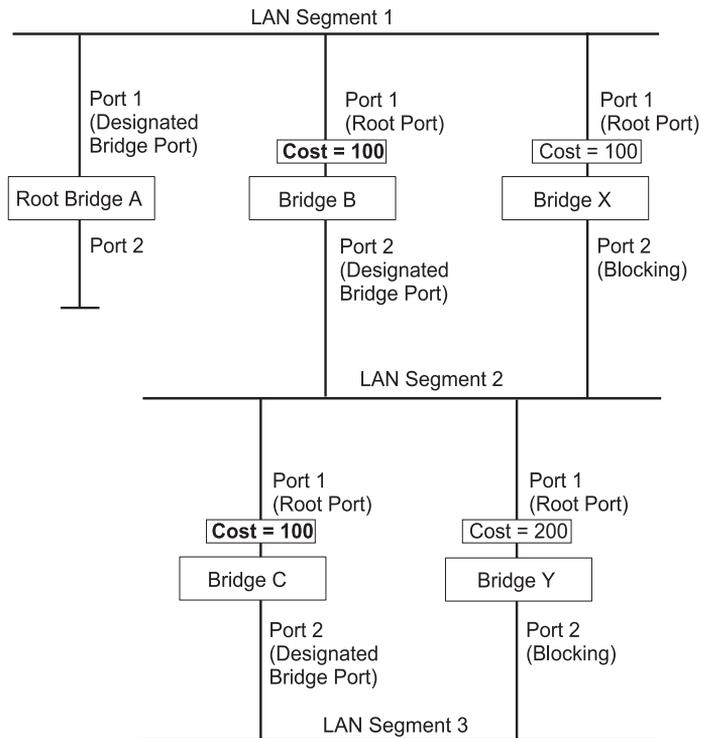
How RSTP Differs to STP

RSTP works in a similar way to STP, but it includes additional information in the BPDUs. This information allows each bridge to confirm that it has taken action to prevent loops from forming when it wants to enable a link to a neighbouring bridge. This allows adjacent bridges connected via point-to-point links to enable a link without having to wait to ensure all other bridges in the network have had time to react to the change.

So the main benefit of RSTP is that the configuration decision is made locally rather than network-wide which is why RSTP can carry out automatic configuration and restore a link faster than STP.

STP Example

[Figure 11](#) shows a LAN that has STP enabled. The LAN has three segments, and each segment is connected using two possible links.

Figure 11 Port costs in a network

- Bridge A has the lowest Bridge Identifier in the network, and has therefore been selected as the Root Bridge.
- Because Bridge A is the Root Bridge, it is also the Designated Bridge for LAN segment 1. Port 1 on Bridge A is therefore selected as the Designated Bridge Port for LAN Segment 1.
- Port 1 of Bridges B, C, X and Y have been defined as Root Ports because they are the nearest to the Root Bridge and therefore have the most efficient path.
- Bridges B and X offer the same Root Path Cost for LAN segment 2, however, Bridge B has been selected as the Designated Bridge for the segment because it has a lower Bridge Identifier. Port 2 on Bridge B is therefore selected as the Designated Bridge Port for LAN Segment 2.

- Bridge C has been selected as the Designated Bridge for LAN segment 3, because it offers the lowest Root Path Cost for LAN Segment 3:
 - the route through Bridges C and B costs 200 (C to B=100, B to A=100)
 - the route through Bridges Y and B costs 300 (Y to B=200, B to A=100).

Port 2 on Bridge C is therefore selected as the Designated Bridge Port for LAN Segment 3.

STP Configurations [Figure 12](#) shows three possible STP configurations using SuperStack 3 Switch units.

- **Configuration 1 — Redundancy for Backbone Link**

In this configuration, the Switches both have STP enabled and are connected by two links. STP discovers a duplicate path and blocks one of the links. If the enabled link breaks, the disabled link becomes re-enabled, therefore maintaining connectivity.

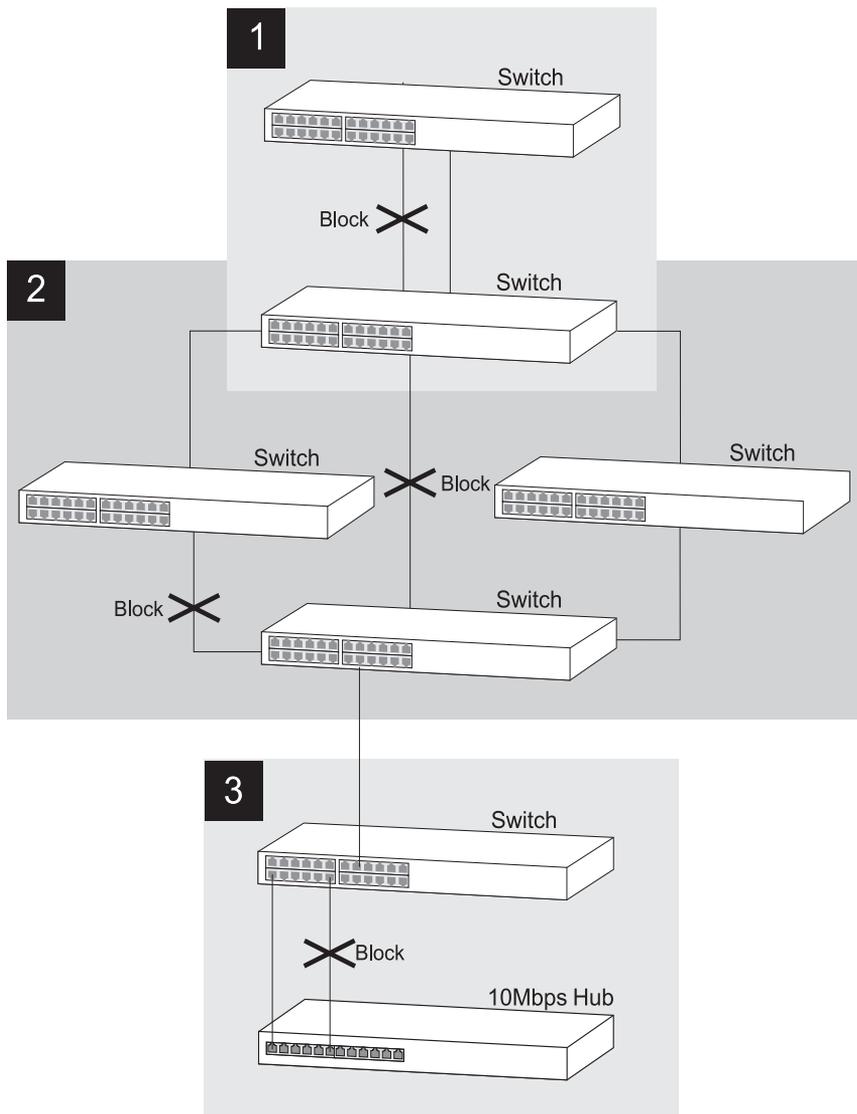
- **Configuration 2 — Redundancy through Meshed Backbone**

In this configuration, four Switch units are connected in a way that creates multiple paths between each one. STP discovers the duplicate paths and blocks two of the links. If an enabled link breaks, one of the disabled links becomes re-enabled, therefore maintaining connectivity.

- **Configuration 3 — Redundancy for Cabling Error**

In this configuration, a Switch has STP enabled and is accidentally connected to a hub using two links. STP discovers a duplicate path and blocks one of the links, therefore avoiding a loop.

Figure 12 STP configurations

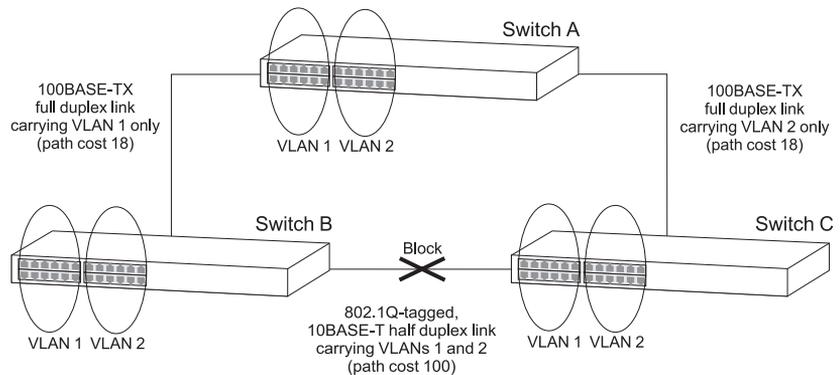


Using STP on a Network with Multiple VLANs

The IEEE Std 802.1D, 1998 Edition does not take into account VLANs when it calculates STP information — the calculations are only performed on the basis of physical connections. For this reason, some network configurations can result in VLANs being subdivided into a number of isolated sections by the STP system. Therefore, you must ensure that any VLAN configuration on your network takes into account the expected STP topology and alternative topologies that may result from link failures.

For example, [Figure 13](#) shows a network containing VLANs 1 and 2. They are connected using the 802.1Q-tagged link between Switch B and Switch C. By default, this link has a path cost of 100 and is automatically blocked because the other Switch-to-Switch connections have a path cost of 36 (18+18). This means that both VLANs are now subdivided — VLAN 1 on Switch units A and B cannot communicate with VLAN 1 on Switch C, and VLAN 2 on Switch units A and C cannot communicate with VLAN 2 on Switch B.

Figure 13 Configuration that separates VLANs



To avoid any VLAN subdivision, it is recommended that all inter-Switch connections are made members of all available 802.1Q VLANs to ensure connectivity at all times. For example, the connections between Switches A and B, and between Switches A and C should be 802.1Q tagged and carrying VLANs 1 and 2 to ensure connectivity.



For more information about VLAN Tagging, see [Chapter 8 “Setting Up Virtual LANs”](#).

5

USING THE SWITCH DATABASE

What is the Switch Database?

The Switch Database is used by the Switch to determine where a packet should be forwarded to, and which port should transmit the packet if it is to be forwarded.

The database contains a list of entries — each entry contains three items:

- MAC (Ethernet) address information of the endstation that sends packets to the Switch.
- Port identifier, that is the port attached to the endstation that is sending the packet.
- VLAN ID of the VLAN to which the endstation belongs.



For details of the number of addresses supported by your Switch database, please refer to Chapter 1 of the Getting Started Guide that accompanies your Switch.



For detailed descriptions of the Web interface operations and the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch.

How Switch Database Entries Get Added

Entries are added to the Switch Database in one of two ways:

- The Switch can learn entries. The Switch updates its database with the source MAC address of the endstation that sent the packet, the VLAN ID, and the port identifier on which the packet is received.
- You can enter and update entries using the management interface via the *Bridge > Address Database* Web interface operation, or an SNMP Network Manager.

Switch Database Entry States

Databases entries can have three states:

- *Learned* — The Switch has placed the entry into the Switch Database when a packet was received from an endstation. Note that:
 - Learned entries are removed (aged out) from the Switch Database if the Switch does not receive further packets from that endstation within a certain period of time (the *aging time*). This prevents the Switch Database from becoming full with obsolete entries by ensuring that when an endstation is removed from the network, its entry is also removed from the database.
 - Learned entries are removed from the Switch Database if the Switch is reset or powered-down.
- *Non-aging learned* — If the aging time is set to 0 seconds, all learned entries in the Switch Database become non-aging learned entries. This means that they are not aged out, but they are still removed from the database if the Switch is reset or powered-down.
- *Permanent* — The entry has been placed into the Switch Database using the management interface. Permanent entries are not removed from the Switch Database unless they are removed using the Switch management interface via the `bridge > addressDatabase > remove` Web operation or the Switch is initialized.

6

USING TRAFFIC PRIORITIZATION

Using the traffic prioritization capabilities of your Switch provides Quality of Service (QoS) to your network through increased reliability of data delivery. You can prioritize traffic on your network to ensure that high priority data is transmitted with minimum delay.

Traffic can be controlled by a set of rules to obtain the required Quality of Service for your network. The rules define different types of traffic and specify how each type should be treated as it passes through the Switch, for example, prioritized or discarded. Being able to define exactly how you want your Switch to treat selected applications, devices, users and types of traffic allows you to have more control over your network.

There are two different categories of rules:

- **Application-based rules** — describe how to deal with traffic for a specific application, for example, Netmeeting or Lotus Notes.
- **Device-based rules** — describe how to deal with traffic that flows to and from specific devices, for example, servers or server farms.

This chapter explains more about traffic prioritization.

- [What is Traffic Prioritization?](#)
- [How Traffic Prioritization Works](#)
- [Important QoS Considerations](#)
- [Default QoS Configurations](#)



Basic traffic prioritization is the default level of QoS supported by the Switch 3812 and Switch 3824.



For a list of the features supported by your Switch, please refer to the Management Quick Reference Guide that accompanies your Switch.



For detailed descriptions of the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch.

What is Traffic Prioritization?

Today's application traffic consists of three common types of data:

- Time critical data such as video and voice.
- Business critical data such as database transactions and online transactions.
- Opportunistic data such as web browsing, email and file transfers.

When these different types of data compete for the same bandwidth, a network can quickly become overloaded, resulting in slow response times (long latency), and application time-outs. Traffic prioritization is a mechanism that allows you to prioritize data so that time-sensitive and system-critical data can be transferred smoothly and with minimal delay over a network.

The benefits of using traffic prioritization are:

- You can control a wide variety of traffic and manage congestion on your network, therefore improving performance.
- You can assign priorities to traffic, for example, set higher priorities to time-critical or business-critical applications.
- You can provide predictable throughput for multimedia applications such as video conferencing or voice over IP platforms like the 3Com NBX, as well as minimizing traffic delay and jitter.
- You can improve network performance as the amount of traffic grows, which also reduces the need to constantly add bandwidth to the network, therefore saving cost.

How Traffic Prioritization Works

Traffic prioritization uses the eight traffic queues that are present in your Switch to ensure that high priority traffic is forwarded on a different queue from lower priority traffic. This is what provides Quality of Service (QoS) to your network.

Traffic prioritization in your Switch may be applied dependent upon following factor:

- **The level of service requested by an end-station** — the transmitting end-station sets the priority of each stream of traffic. Received traffic at the Switch is forwarded through the appropriate queue depending on its priority level for onward transmission across the network.

A QoS network can differentiate between time critical data, business critical data and opportunistic data (such as email, File Transfer Protocol (FTP) and Web traffic). A QoS network also has the ability to stop unauthorized usage of the network, such as online gaming.

To achieve quality of service the Switch will use four processes:

- **Traffic Classification** — a QoS network examines the traffic to identify which application or device generated the traffic.
- **Traffic Marking** — after traffic is identified, it is Marked so that other network devices can identify the data and give it the correct level of service.
- **Traffic Re-marking** — if a traffic packet enters the Switch with a priority marking requesting an unacceptable level of service, the Switch can Re-mark it with a different priority value to downgrade its level of service.
- **Traffic Prioritization** — once the network can differentiate types of traffic, for example, a telephone conversation from Web surfing, prioritization can ensure that a large download from the Internet does not disrupt the telephone conversation.



The Switch is configured to handle priority tagged packets and NBX phone traffic.

Traffic Classification

To determine the service level to be applied to each incoming traffic type, each packet or frame must first be classified. Traffic classification is the means of identifying which application, device or user generated the traffic.

The Switch employs several methods of classifying (identifying) traffic. These can be based on any combination of fields in the first 64 bytes of the packet, and at different levels of the 7 layer OSI model as shown in [Table 5](#).

Table 5 Attributes on which incoming traffic can be classified (identified)

OSI Layer and Protocols	Summary of Protocols
Layer 2 <ul style="list-style-type: none"> ■ IEEE 802.1D priority ■ EtherType 	Chatty protocols such as AppleTalk and IPX, used by a small number of older devices, can cause traffic delays. Identifying and prioritizing data based on these protocols can reduce delays. AppleTalk can be identified by its EtherType of 0x809B, and IPX can be identified by EtherType 0x8137.
Layer 3 <ul style="list-style-type: none"> ■ Destination IP address ■ Source IP address ■ IP protocols: (ICMP, IGMP, RSVP, UDP, TCP, etc) ■ DiffServ code point (DSCP) 	Many applications are identified by their Source IP address, or IP protocol. Because servers are sometimes dedicated to single applications, such as email, the Source IP address or protocol in a packet can identify which application generated the packet. As well as being a traffic marking mechanism, the DSCP field in the IP header can also be used to classify traffic.
Layer 4 <ul style="list-style-type: none"> ■ UDP / TCP Source and Destination ports for IP applications 	Many applications use certain TCP or UDP sockets to communicate. By examining the socket number in the IP packet, the intelligent network can determine what type of application generated the packet. This is also known as Layer 4 switching.

Traffic Marking

After traffic has been identified through classification, it must be Marked to ensure that other devices such as Layer 2 switches or routers on the network know how to prioritize the application, device or user that generated it. The Switch uses two of the industry-standard methods of marking network traffic:

- **IEEE 802.1D** — a layer 2 marking scheme.
- **Differentiated Services (DiffServ)** — a layer 3 marking scheme.

IEEE 802.1D Traffic Marking

The IEEE Std 802.1D, 1998 Edition marking scheme is an enhancement to the IEEE Std 802.1D to enable Quality of Service in the LAN. Traffic service levels are defined in the IEEE 802.1Q 4-byte tag, which is used to carry VLAN identification as well as IEEE 802.1p priority information. The 4 byte tag immediately follows the destination MAC address and Source MAC address.

The IEEE Std 802.1D, 1998 Edition priority marking scheme assigns each frame with an IEEE 802.1p priority level between 0 and 7, which

determines the level of service that type of traffic should receive. Refer to [Table 6](#) for an example of how different traffic types can be mapped to the eight IEEE 802.1p priority levels.

Table 6 IEEE recommendation for mapping 802.1p priority levels to 802.1D traffic types

IEEE 802.1p Priority Level	IEEE 802.1D Traffic Type
0	Best Effort (default)
1	Background
2	Standard (spare)
3	Excellent Effort (business critical)
4	Controlled Load (streaming multimedia)
5	Video (interactive media), less than 100 milliseconds latency and jitter
6	Voice (interactive voice), less than 10 milliseconds latency and jitter
7	Network Control Reserved traffic



The traffic marking and prioritization supported by the Switch using layer 2 information is compatible with the relevant sections of the IEEE Std 802.1D, 1998 Edition (incorporating IEEE 802.1p).

The IEEE 802.1D standard is the most widely used prioritization scheme in the LAN environment, but it does however have some restrictions:

- It requires an additional 4-byte tag in the frame, which is normally optional in Ethernet networks. Without this tag, the scheme cannot work.
- The tag is part of the IEEE 802.1Q header, so to implement QoS at layer 2, the entire network has to implement IEEE 802.1Q VLAN tagging.
- It is only supported on a LAN and not across routed WAN links, because the IEEE 802.1Q tags are removed when the packets pass through a router.

Differentiated Services (DiffServ) Traffic Marking

DiffServ is a Layer 3 marking scheme that uses the DiffServ Code Point (DSCP) field in the IP header to store the packet priority information. DSCP is an advanced intelligent method of traffic marking because you can choose how your network prioritizes different types of traffic. DSCP

uses 64 values that map to user-defined service levels, allowing you to establish more control over network traffic.

Advantages of DiffServ over IEEE 802.1D are:

- No extra tags are required in the packet.
- DSCP uses the IP header of a packet and therefore priority is preserved across the Internet.
- DSCP is backward compatible with IPv4 TOS, which allows operation with any existing devices with layer 3 TOS enabled prioritization scheme in use.

Traffic Re-Marking

Traffic entering the Switch may get downgraded depending on the network policies. If for example a traffic packet enters the Switch with a priority marking higher than the Quality of Service (QoS) configuration policies, the Switch will Re-Mark the packet with a different 802.1D priority or new DSCP value.

Traffic Prioritization

Your Switch supports Basic and Advanced Quality of Service (QoS) traffic prioritization. Basic traffic prioritization classifies traffic based on layer 2 of the OSI 7 layer model, and the Switch will prioritize the received traffic according to the priority information defined in the received packet. Advanced traffic prioritization can classify traffic at layers 2, 3 and 4 of the OSI 7 layer model.

Basic Traffic Prioritization

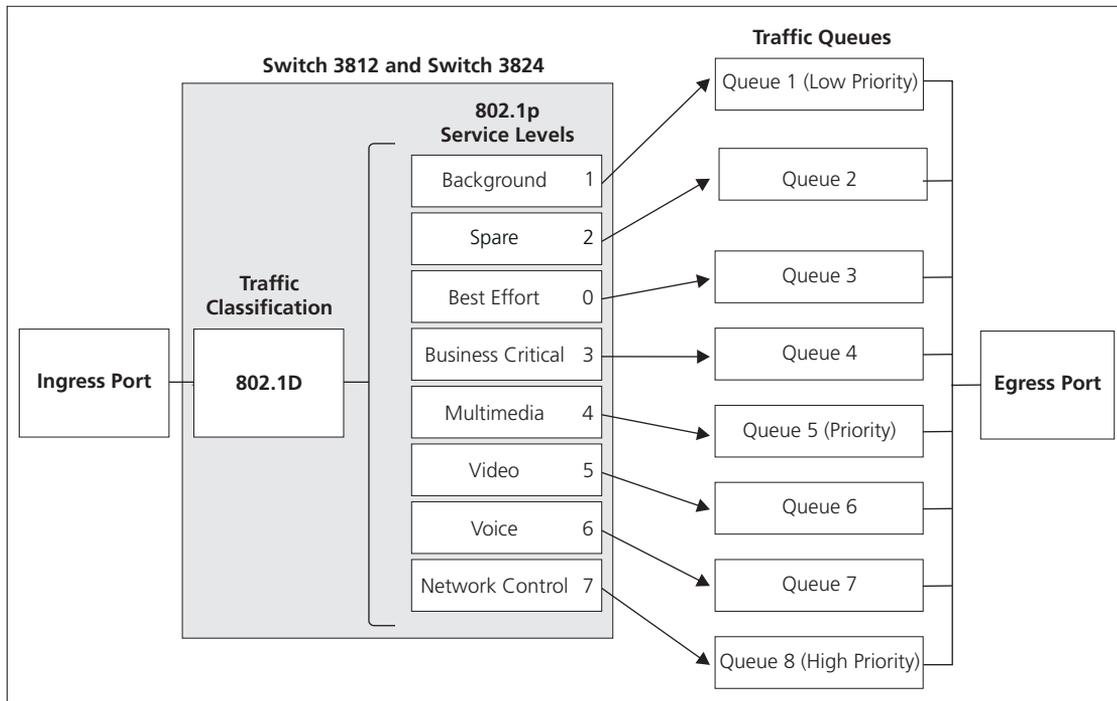
Incoming traffic is classified based upon the IEEE 802.1D frame and is assigned to the appropriate priority queue based upon the IEEE 802.1p service level value defined in that packet. Service level markings (values) are defined in the IEEE 802.1Q 4-byte tag, and therefore traffic will only contain 802.1p priority markings if the network is configured with VLANs and VLAN tagging. The Switch 3812 and Switch 3824 both support basic traffic prioritization. The traffic flow through the Switch is as follows:

- 1 A packet received by the Switch may or may not have an 802.1p tag associated with it. If it does not, then it is given a default 802.1p tag (which is 0). The packet may be remarked with a new 802.1p value, which will result in all knowledge of the old 802.1p tag being lost.
- 2 Because the 802.1p priority levels are fixed to the traffic queues (as shown in [Figure 14](#) on [page 53](#)), the packet will be placed in the appropriate priority queue, ready for transmission through the

appropriate egress port(s). When the packet reaches the head of its queue and is about to be transmitted the device determines whether or not the egress port is tagged for that VLAN. If it is, then the new 802.1p tag is used in the extended 802.1D header.

The IEEE 802.1D standard specifies eight distinct levels of priority (0 to 7), each of which relates to a particular type of traffic. The priority levels and their traffic types are shown in [Figure 14](#) in order of increasing priority. The mapping from 802.1p level to traffic queue in the Switch is proprietary and is slightly different to the recommended IEEE mapping.

Figure 14 IEEE 802.1p priority levels and recommended IEEE 802.1D traffic types



The number of queues and their mappings to the 8 levels is proprietary and can even vary between Switches from the same vendor.



You cannot alter the mapping between the IEEE 802.1p priorities and the traffic queues. These are calculated to be the most efficient, and are fixed as illustrated in [Figure 14](#).

[Figure 14](#) shows how traffic prioritization works at layer 2. The Switch will check a packet received at the ingress port for IEEE 802.1D traffic classification, and then prioritize it based upon the IEEE 802.1p value (service levels) in that tag. It is this 802.1p value that determines which traffic queue the packet is mapped to.

Advanced Traffic Prioritization

Incoming traffic can be classified based on packet attributes at different layers of the OSI 7 layer model. The Switch can look in the packet for layer 2, 3 and 4 attributes to identify incoming traffic.

Most of the current applications, for example Microsoft Word, Lotus Notes and NetMeeting, are not QoS-aware and do not apply a service level to the traffic that they send. Being an intelligent Switch, your Switch can use its own rules to classify and mark the traffic. If the incoming traffic has pre-defined service level markings, the advanced traffic prioritization of your Switch will assign the appropriate DSCP and 802.1D service level markings to that incoming traffic.

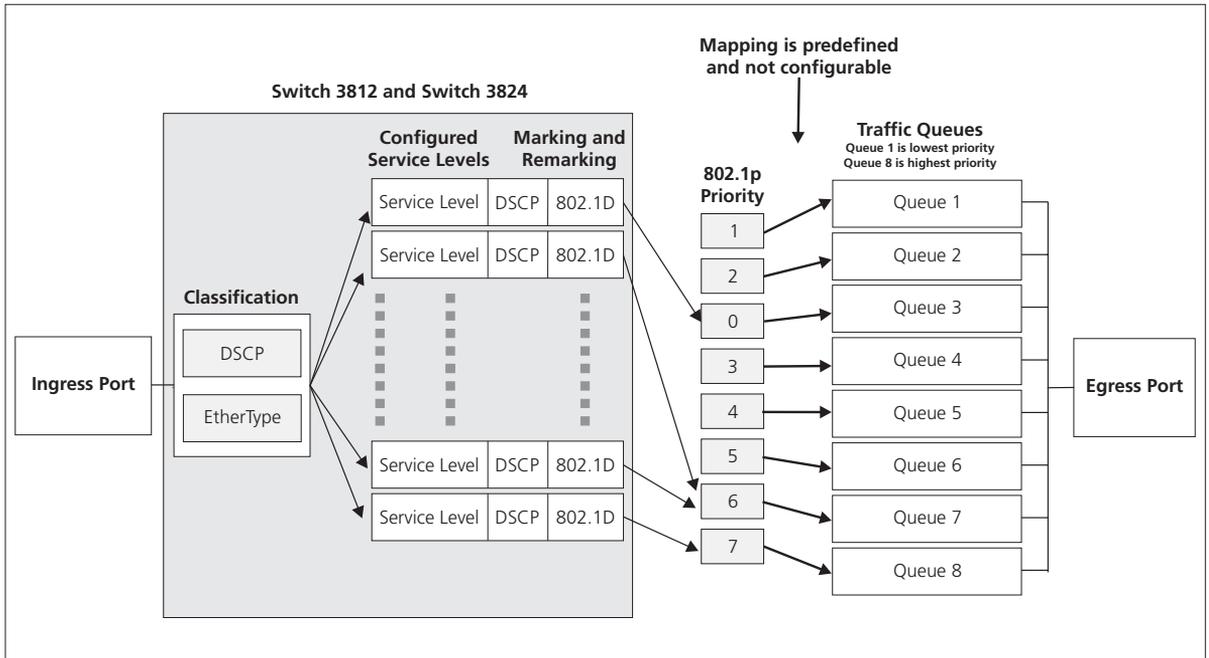
The advanced traffic prioritization in the Switch will:

- Classify traffic based on different packet attributes. The four common methods of classification are DSCP, TCP/UDP ports, IP Address and Protocol.
- Mark traffic as it enters the Switch with the appropriate DSCP and 802.1D markings.

[Figure 15](#) shows how traffic is treated using the advanced traffic prioritization in the Switch.



The DSCP field in the IP packet header can be used to classify (identify) traffic as well as carrying the priority markings, as shown in Figure 15.

Figure 15 Advanced traffic prioritization and marking

- 1 The packet received at the ingress port is checked for any of the supported traffic classification methods (DSCP, EtherType, Protocol) to identify the traffic.
- 2 The classification in an incoming packet will be compared with the predefined classifications in the Switch, and if there is a match, the configured service level associated with the classified traffic will be applied.
- 3 The service level associated with the classifier may cause the 802.1p tag to be remarked, if the packet already has an 802.1p tag, and the DSCP value in IP packets to be remarked.
- 4 The Switch will remark the 802.1p tag and DSCP field.
- 5 It is the priority associated with the packet that is used to direct it to the appropriate queue. This is determined as follows:
 - If the packet matches a classifier with a configured service level specifying that the DSCP or 802.1p tag should be re-marked, then the packet is re-marked with the configured DSCP value and or the 802.1p priority.

- Otherwise, if there are no other classifiers except the 802.1p tag, then the packet will pass through the Switch with the original 802.1p priority tag.
- Otherwise, if the received packet does not have an 802.1p tag, then a default 802.1p tag (which is usually 0) is assigned to it.

Traffic Queues

It is the multiple traffic queues within the Switch hardware that allow packet prioritization to occur. Higher priority traffic can pass through the Switch without being delayed by lower priority traffic. As each packet arrives in the Switch, it passes through any ingress processing (which includes classification or marking/remarking), and is then sorted into the appropriate queue. The Switch then forwards packets from each queue. It is worth noting that each egress port has its own set of queues, so that if one port is congested it does not interfere with the queue operation of other ports.

The Switch uses the Weighted Round Robin (WRR) queuing mechanism. This method services all the traffic queues, giving priority to the higher priority queues. Under most circumstances, this method gives high priority precedence over low-priority, but in the event that high-priority traffic exceeds the link capacity, lower priority traffic is not blocked.



Traffic queues cannot be enabled on a per-port basis on the Switch 3812 and Switch 3824.

Important QoS Considerations

Before implementing QoS on your network you need to consider the following points:

- Your Switch:
 - Has a predefined Classifier for NBX traffic, which is DSCP service level 46. If the profile assigned to the port on which the NBX traffic is received has an NBX classifier in it such as the default profile does, the Switch will automatically detect NBX telephone voice traffic and prioritize accordingly. The Switch also has an NBX classifier for Ethernet Type 0x8868, which is the layer 2 NBX traffic identifier. NBX is layer 2 out of the box and has to be configured by the user to be layer 3, so the DSCP 46 classifier may in fact be used in fewer NBX installations than the Ethernet Type 0x8868.
 - Can map between IEEE 802.1D and DSCP to support legacy devices in the network that only support IEEE 802.1D.

- Has eight traffic queues, but it is important to note that not all Switches have the same number of priority queues.
- QoS is about providing a consistent, predictable data delivery service. It should not be used as an alternative to deploying sufficient bandwidth. The recommended configuration for most networks is 10/100 Mbps switching to the desktop, Gigabit connections for servers, and non-blocking Gigabit backbones.
- QoS requires the support of every network device from end-to-end. All devices in the network should support QoS. If there is just one section in the data path that does not support QoS, it can produce bottlenecks and slowdowns, although a performance improvement will be noticed over the parts of the network that do support QoS.
- Ensure that all QoS devices are configured the same way. Mismatches will cause the same traffic to be prioritized in one section and not in another.
- Only use Switches or hardware-based routers in the LAN. Hubs cannot prioritize traffic, and software-based routers can cause bottlenecks.
- Use Switches and hardware-based routers that understand both the IEEE 802.1D (incorporating IEEE 802.1p) and DSCP marking schemes.
- Classify traffic as soon as it enters the network. If traffic is not classified until it gets to the WAN router or firewall, end-to-end prioritization cannot be guaranteed. The ideal place for traffic classification is within the Switch.
- Traffic Marking is performed as a result of classification, and so you should aim to perform the marking only once to reduce the additional requirements that QoS places upon the capabilities of your network infrastructure.
- As DSCP uses a field in the IP header, it is only possible to use the DSCP in IP packets. It does not apply, for example, to AppleTalk, IPX or NetBEUI.
- Because DSCP is a redefinition of the use of the TOS byte in the IP header, there are some issues with interaction with IP TOS based networks.

Default QoS Configurations

The Switch is pre-configured with the following settings:

Table 7 Default traffic classifiers configured in your Switch

Classifier Name	Classifier Type	Protocol Identifier	IEEE 802.ID Priority	DSCP Marking
3Com NBX Voice-LAN	EtherType	0x8868	6	46
3Com NBX Voice-IP	DSCP	46	6	46
Internet Network Control	DSCP	48	7	48
Network Control	DSCP	56	7	56

7

STATUS MONITORING AND STATISTICS

This chapter contains details of the Remote Monitoring ([RMON](#)) feature that assists you with status monitoring and statistics.



For detailed descriptions of the Web interface operations and the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch.

RMON

Using the RMON capabilities of a Switch allows you to improve your network efficiency and reduce the load on your network.

This section explains more about RMON. It covers the following topics:

- [What is RMON?](#)
- [Benefits of RMON](#)
- [RMON and the Switch](#)

What is RMON?

RMON is a system defined by the IETF (Internet Engineering Task Force) that allows you to monitor the traffic of LANs or VLANs.

RMON is an integrated part of the Switch software agent and continually collects statistics about a LAN segment or VLAN, and transfers the information to a management workstation on request or when a pre-defined threshold is crossed. The workstation does not have to be on the same network as the Switch and can manage the Switch by in-band or out-of-band connections.

The RMON Groups

The IETF define groups of Ethernet RMON statistics. This section describes the four groups supported by the Switch, and details how you can use them.

Statistics

The Statistics group provides traffic and error statistics showing packets, bytes, broadcasts, multicasts and errors on a LAN segment or VLAN.

Information from the Statistics group is used to detect changes in traffic and error patterns in critical areas of your network.

History

The History group provides historical views of network performance by taking periodic samples of the counters supplied by the Statistics group.

The group is useful for analyzing the traffic patterns and trends on a LAN segment or VLAN, and for establishing the normal operating parameters of your network.

Alarms

The Alarms group provides a mechanism for setting thresholds and sampling intervals to generate events on any RMON variable.

Alarms are used to inform you of network performance problems and they can trigger automated responses through the Events group.

Events

The Events group provides you with the ability to create entries in an event log and send SNMP traps to the management workstation. Events are the action that can result from an RMON alarm. In addition to the standard five traps required by SNMP (link up, link down, warm start, cold start, and authentication failure), RMON adds two more: rising threshold and falling threshold.

Effective use of the Events group saves you time; rather than having to watch real-time graphs for important occurrences, you can depend on the Event group for notification. Through the SNMP traps, events can trigger other actions, therefore providing a way to automatically respond to certain occurrences.

Benefits of RMON

Using the RMON features of your Switch has three main advantages:

- **It improves your efficiency**

Using RMON allows you to remain at one workstation and collect information from widely dispersed LAN segments or VLANs. This means that the time taken to reach a problem site, set up equipment, and begin collecting information is largely eliminated.

- **It allows you to manage your network in a more proactive manner**

If configured correctly, RMON can deliver information before problems occur. This means that you can take action before they affect users. In addition, probes record the behavior of your network, so that you can analyze the causes of problems.

- **It reduces the load on the network and the management workstation**

Traditional network management involves a management workstation polling network devices at regular intervals to gather statistics and identify problems or trends. As network sizes and traffic levels grow, this approach places a strain on the management workstation and also generates large amounts of traffic.

RMON, however, autonomously looks at the network on behalf of the management workstation without affecting the characteristics and performance of the network. RMON reports by exception, which means that it only informs the management workstation when the network has entered an abnormal state.

RMON and the Switch

The RMON support provided by your Switch 3812 and Switch 3824 is detailed in [Table 8](#).

Table 8 RMON support supplied by the Switch

RMON group	Support supplied by the Switch
Statistics	A new or initialized Switch has one Statistics session per port.
History	A new or initialized Switch has two History sessions per port. These sessions provide the data for the Web interface history displays: <ul style="list-style-type: none"> ■ 10min intervals, 6 historical samples stored ■ 1 hour intervals, 6 historical samples stored

Table 8 RMON support supplied by the Switch

RMON group	Support supplied by the Switch
Alarms	A new or initialized Switch has the following alarm(s) defined for each port: For more information about the alarms setup on the Switch, see “Alarm Events” on page 62 .
Events	A new or initialized Switch has Events defined for use with the default alarm system.

When using the RMON features of the Switch, note the following:

- After the default sessions are created, they have no special status. You can delete or change them as required.
- The greater the number of RMON sessions, the greater the burden on the management resources of the Switch. If you have many RMON sessions, the forwarding performance of the Switch is not affected but you may experience slow response times from the Web interface.

Alarm Events

You can define alarms for the Switch. The events that you can define for each alarm and their resulting actions are listed in [Table 9](#).

Table 9 Alarm Events

Event	Action
No action	
Notify only	Send Trap.
Notify and filter port	Send Trap. Block broadcast and multicast traffic on the port. Recovers with the <i>unfilter port</i> event.
Notify and disable port	Send Trap. Turn port off.
Notify and enable port	Send Trap. Turn port on.
Disable port	Turn port off.
Enable port	Turn port on.
Notify and switch resilient port	Send Trap. If port is the main port of a resilient link pair then move to standby.
Notify and unfilter port	Send Trap. Stop blocking broadcast and multicast traffic on the port.
System started	

8

SETTING UP VIRTUAL LANs

Setting up Virtual LANs (VLANs) on your Switch increases the efficiency of your network by dividing the LAN into logical, rather than physical, segments which are easier to manage.

This chapter explains more about the concept of VLANs and explains how they can be implemented on your Switch. It covers the following topics:

- [What are VLANs?](#)
- [Benefits of VLANs](#)
- [VLANs and Your Switch](#)
- [VLAN Configuration Examples](#)

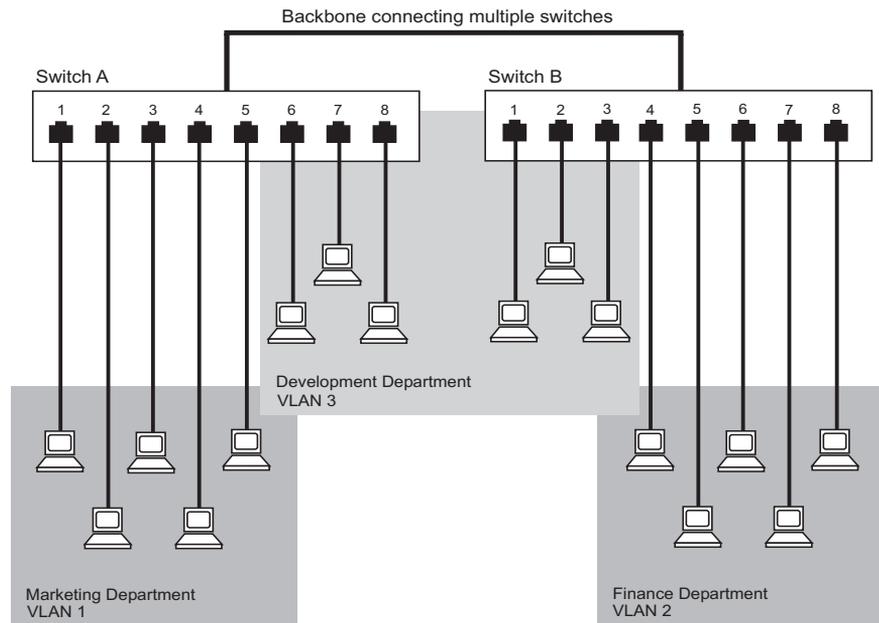


For detailed descriptions of the Web interface operations and the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the Management Interface Reference Guide supplied in HTML format on the CD-ROM that accompanies your Switch.

What are VLANs?

A VLAN is a flexible group of devices that can be located anywhere in a network, but which communicate as if they are on the same physical segment. With VLANs, you can segment your network without being restricted by physical connections — a limitation of traditional network design. As an example, with VLANs you can segment your network according to:

- **Departmental groups** — For example, you can have one VLAN for the Marketing department, another for the Finance department, and another for the Development department.
- **Hierarchical groups** — For example, you can have one VLAN for directors, another for managers, and another for general staff.
- **Usage groups** — For example, you can have one VLAN for users of e-mail, and another for users of multimedia.

Figure 16 A network setup showing three VLANs

Benefits of VLANs

The main benefit of VLANs is that they provide a network segmentation system that is far more flexible than any traditional network. Using VLANs also provides you with three other benefits:

- ### ■ VLANs ease the movement of devices on networks

With traditional networks, network administrators spend much of their time dealing with moves and changes. If users move to a different subnetwork, the addresses of each endstation must be updated manually.

With a VLAN setup, if an endstation in VLAN *Marketing* for example is moved to a port in another part of the network, and retains its original subnet membership, you only need to specify that the new port is in VLAN *Marketing*. You do not need to carry out any re-cabling.

- ### ■ VLANs provide extra security

Devices within each VLAN can only communicate with other devices in the same VLAN. If a device in VLAN *Marketing* needs to communicate with devices in VLAN *Finance*, the traffic must pass through a routing device or Layer 3 Switch.

- **VLANs help to control traffic**

With traditional networks, congestion can be caused by broadcast traffic that is directed to all network devices whether they require it or not. VLANs increase the efficiency of your network because each VLAN can be set up to contain only those devices that need to communicate with each other.

VLANs and Your Switch

Your Switch provides support for VLANs using the IEEE Std 802.1Q-1998. This standard allows traffic from multiple VLANs to be carried across one physical link.

The IEEE Std 802.1Q-1998 allows each port on your Switch to be placed in:

- Any one VLAN defined on the Switch.
- Several VLANs at the same time using 802.1Q tagging.

The standard requires that you define the following information about each VLAN on your Switch before the Switch can use it to forward traffic:

- *VLAN Name* — This is a descriptive name for the VLAN (for example, Marketing or Management).
- *802.1Q VLAN ID* — This is used to identify the VLAN if you use 802.1Q tagging across your network.

The Default VLAN

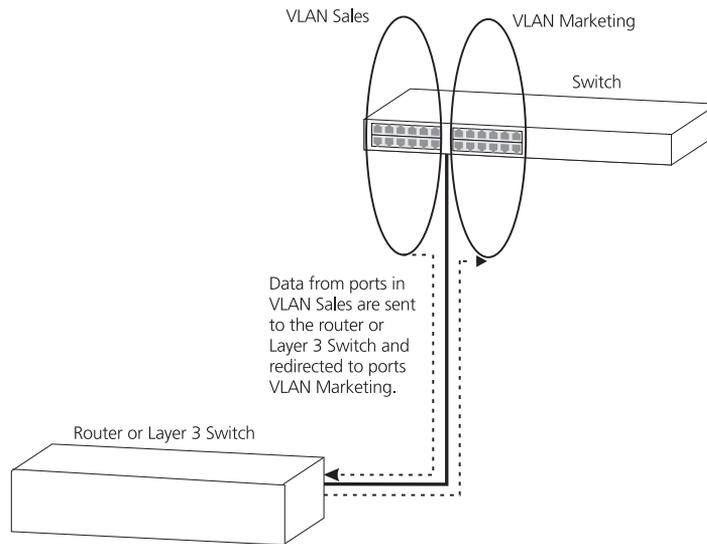
A new or initialized Switch contains a single VLAN, the Default VLAN. This VLAN has the following definition:

- *VLAN Name* — Default VLAN
- *802.1Q VLAN ID* — 1 (if tagging required)

All the ports are initially placed in this VLAN, and it is the only VLAN that allows you to access the management software of the Switch over the network.

Communication Between VLANs

If the devices placed in a VLAN need to communicate to devices in a different VLAN, a router or Layer 3 switching device with connections to both VLANs needs to be installed. Communication between VLANs can only take place if they are all connected to a routing or Layer 3 switching device.

Figure 17 Two VLANs connected via a router**Creating New VLANs**

If you want to move a port from the Default VLAN to another VLAN, you must first define information about the new VLAN on your Switch.

VLANs: Tagged and Untagged Membership

Your Switch supports 802.1Q VLAN tagging, a system that allows traffic for multiple VLANs to be carried on a single physical (backbone) link.

When setting up VLANs you need to understand when to use untagged and tagged membership of VLANs. Quite simply, if a port is in a single VLAN it can be an untagged member, but if the port needs to be a member of multiple VLANs, tagged membership must be defined. Typically endstations (for example, clients) will be untagged members of one VLAN, while inter-Switch connections will be tagged members of all VLANs.

The IEEE Std 802.1Q-1998 defines how VLANs operate within an open packet-switched network. An 802.1Q compliant packet carries additional information that allows a Switch to determine to which VLAN the port belongs. If a frame is carrying the additional information, it is known as *tagged*.

To carry multiple VLANs across a single physical (backbone) link, each packet must be tagged with a VLAN identifier so that the Switches can

identify which packets belong in which VLANs. To communicate between VLANs a router must be used.

VLAN Configuration Examples

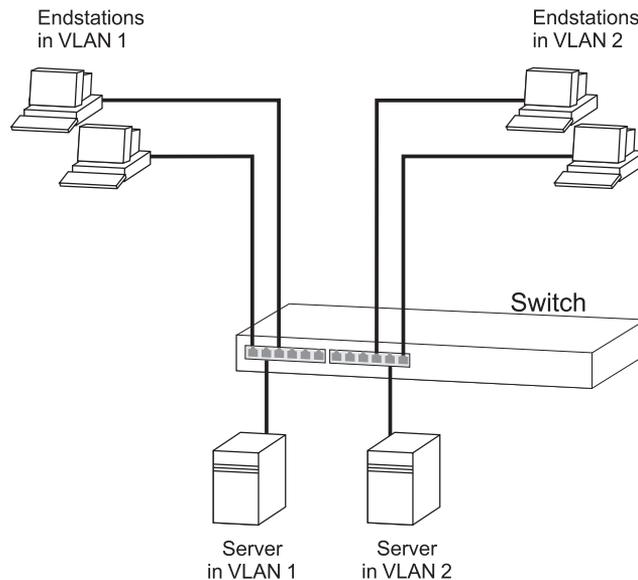
This section contains examples of VLAN configurations. It describes how to set up your Switch to support simple untagged and tagged connections.

Using Untagged Connections

The simplest VLAN operates in a small network using a single switch. In this network there is no requirement to pass traffic for multiple VLANs across a link. All traffic is handled by the single Switch and therefore untagged connections can be used.

The example shown in [Figure 18](#) illustrates a single Switch connected to endstations and servers using untagged connections. Ports 1, 2 and 3 of the Switch belong to VLAN 1, ports 10, 11 and 12 belong to VLAN 2. VLANs 1 and 2 are completely separate and cannot communicate with each other. This provides additional security for your network.

Figure 18 VLAN configuration example: Using untagged connections



To set up the configuration shown in [Figure 18](#):

1 Configure the VLANs

Define VLAN 2 on the Switch. VLAN 1 is the default VLAN and already exists.

2 Add ports to the VLANs

Add ports 10, 11 and 12 of the Switch as untagged members to VLAN 2.



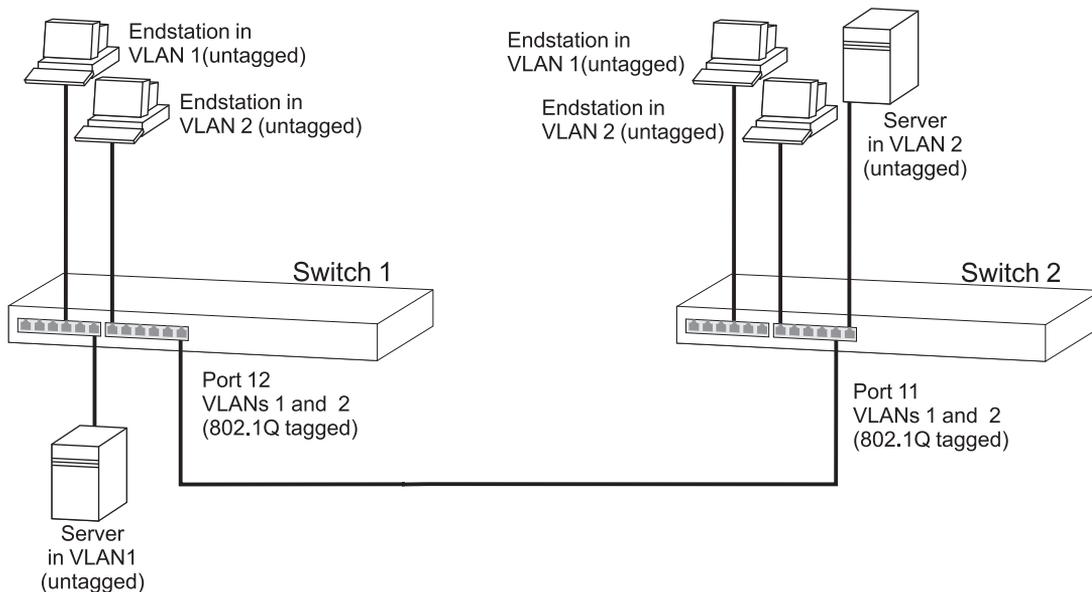
You can use the Switch Web Interface to change VLAN configuration. VLAN configuration can be found at Bridge > VLAN.

Using 802.1Q Tagged Connections

In a network where the VLANs are distributed amongst more than one Switch, you must use 802.1Q tagged connections so that all VLAN traffic can be passed along the links between the Switches. 802.1Q tagging can only be used if the devices at both ends of a link support IEEE 802.1Q.

The example shown in [Figure 19](#) illustrates two Switch units. Each Switch has endstations and a server in VLAN 1 and VLAN 2. All endstations in VLAN 1 need to be able to connect to the server in VLAN1 which is attached to Switch 1 and all endstations in VLAN 2 need to connect to the server in VLAN2 which is attached to Switch 2.

Figure 19 VLAN configuration example: 802.1Q tagged connections



To set up the configuration shown in [Figure 19](#):

1 Configure the VLANs on Switch 1

Define VLAN 2. VLAN 1 is the default VLAN and already exists.

2 Add endstation ports on Switch 1 to the VLANs

Place the endstation ports in the appropriate VLANs as untagged members.

3 Add port 12 on Switch 1 to the VLANs

Add port 12 on Switch 1 as a tagged member of both VLANs 1 and 2 so that all VLAN traffic is passed over the link to Switch 2.

4 Configure the VLANs on Switch 2

Define VLAN 2. VLAN 1 is the default VLAN and already exists.

5 Add endstation ports on Switch 2 to the VLANs

Place the endstation ports in the appropriate VLANs as untagged members.

6 Add port 11 on Switch 2 to the VLANs

Add port 11 on Switch 2 as a tagged member of both VLANs 1 and 2 so that all VLAN traffic is passed over the link to Switch 1.

7 Check the VLAN membership for both Switches

The relevant ports should be listed in the VLAN members summary.

8 Connect the Switches

Connect port 12 on Switch 1 to port 11 on Switch 2.

The VLANs are now configured and operational and the endstations in both VLANs can communicate with their relevant servers.

9

USING AUTOMATIC IP CONFIGURATION

This chapter explains more about IP addresses and how the automatic configuration option works. It covers the following topics:

- [How Your Switch Obtains IP Information](#)
- [How Automatic IP Configuration Works](#)
- [Important Considerations](#)



For detailed information on setting up your Switch for management, see the [Getting Started Guide](#) that accompanies your Switch.



For detailed descriptions of the Web interface operations and the Command Line Interface (CLI) commands that you require to manage the Switch please refer to the [Management Interface Reference Guide](#) supplied in HTML format on the CD-ROM that accompanies your Switch.



For background information on IP addressing, see [Appendix C "IP Addressing"](#).

How Your Switch Obtains IP Information

Your Switch has two ways to obtain its IP address information:

- **Automatic IP Configuration** (default) — the Switch attempts to configure itself by communicating with a DHCP server on the network.
- **Manual IP Configuration** — you can manually input the IP information (IP address, subnet mask, and default gateway).



If you select an option for no IP configuration the Switch will not be accessible from a remote management workstation on the LAN. In addition, the Switch will not be able to respond to SNMP requests.

How Automatic IP Configuration Works

When your Switch is powered up for the first time the IP configuration setting is set to `auto` — this is the default setting.

If your Switch has been powered up before, whichever of the three options for IP configuration (`manual`, `auto`, `none`) was last configured is activated when the Switch powers up again.



You can switch to manual IP configuration at any time using a serial port connection to set up the IP information. For more information see the Getting Started Guide that accompanies your Switch.

Automatic Process

To detect its IP information using the automatic configuration process, the Switch continually attempt to contact a DHCP server on the network requesting IP information from the server.

If a DHCP server is on the network and working correctly it responds to the clients request with an IP address (allocated from a pool of available addresses) and other parameters such as a subnet mask, default gateway, lease time, and any other options configured in the DHCP server.



The way a DHCP server responds is dependant on the DHCP server settings. Therefore the way your DHCP server responds may be different to the process outlined.

Important Considerations

This section contains some important points to note when using the automatic IP configuration feature.



The dynamic nature of automatically configured IP information means that a Switch may change its IP address whilst in use.

Server Support

Your Switch has been tested to interoperate with DHCP servers that use the following operating systems:

- Microsoft Windows 2000 Server
- Microsoft Windows NT4 Server
- Sun Solaris v2.5.1

If you want DHCP to be the method for automatic configuration, make sure that your DHCP servers are operating normally before you power on your Switch.

**Event Log Entries
and Traps**

An event log will be generated and an SNMP trap will be sent if the IP address configuration is changed manually.

A

CONFIGURATION RULES

Configuration Rules for Gigabit Ethernet

Gigabit Ethernet is designed to run over several media:

- Single-mode fiber optic cable, with connections up to 5 km (3.1 miles). Support for distances over 5 km is supported depending on the module specification.
- Multimode fiber optic cable, with connections up to 550 m (1804 ft).
- Category 5 cabling, with connections up to 100 m (328 ft).

The different types of Gigabit Ethernet media and their specifications are detailed in [Table 10](#).

Table 10 Gigabit Ethernet cabling

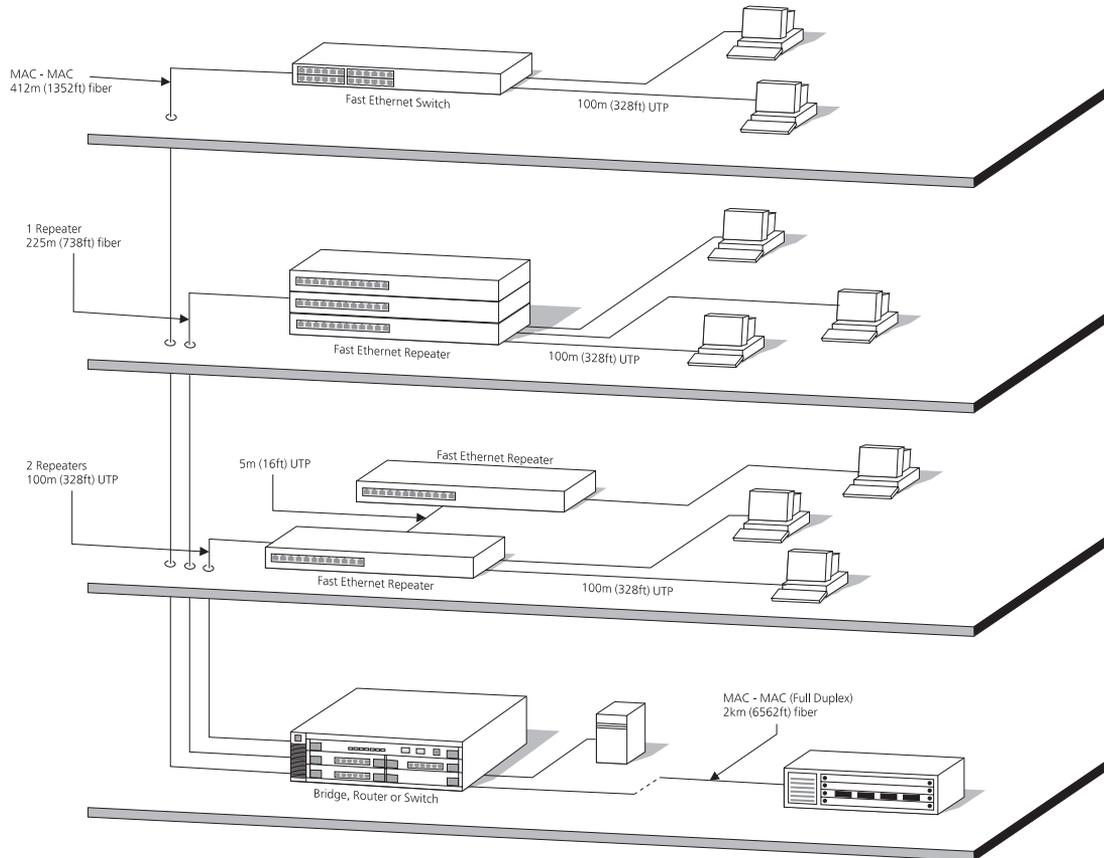
Gigabit Ethernet Transceivers	Fiber Type	Modal Bandwidth (MHz/km)	Lengths Supported Specified by IEEE (meters)
1000BASE-LX	62.5 μ m MM	500	2–550
	50 μ m MM	400	2–550
	50 μ m MM	500	2–550
	10 μ m SM	N/A	2–5000
1000BASE-SX	62.5 μ m MM	160	2–220
	62.5 μ m MM	120	2–275
	50 μ m MM	400	2–500
	50 μ m MM	500	2–550
1000BASE-T	N/A	N/A	100

MM = Multimode SM = Single-mode

Configuration Rules for Fast Ethernet

The topology rules for 100 Mbps Fast Ethernet are slightly different to those for 10 Mbps Ethernet. [Figure 20](#) illustrates the key topology rules and provides examples of how they allow for large-scale Fast Ethernet networks.

Figure 20 Fast Ethernet configuration rules



The key topology rules are:

- Maximum UTP cable length is 100 m (328 ft) over Category 5 cable.
- A 412 m (1352 ft) fiber link is allowed for connecting switch-to-switch, or endstation-to-switch, using half-duplex 100BASE-FX.

- A total network span of 325 m (1066 ft) is allowed in single-repeater topologies (one hub stack per wiring closet with a fiber link to the collapsed backbone). For example, a 225 m (738 ft) fiber link from a repeater to a router or switch, plus a 100 m (328 ft) UTP link from a repeater out to the endstations.

Configuration Rules with Full Duplex

The Switch provides full duplex support for all its ports, including Expansion Module ports. Full duplex allows packets to be transmitted and received simultaneously and, in effect, doubles the potential throughput of a link.

With full duplex, the Ethernet topology rules are the same, but the Fast Ethernet rules are:

- Maximum UTP cable length is 100 m (328 ft) over Category 5 cable.
- A 2 km (6562 ft) fiber link is allowed for connecting switch-to-switch, or endstation-to-switch.

B

NETWORK CONFIGURATION EXAMPLES

This chapter contains the following sections:

- [Simple Network Configuration Examples](#)
 - [Desktop Switch Example](#)
- [Advanced Network Configuration Examples](#)
 - [Improving the Resilience of Your Network](#)
 - [Enhancing the Performance of Your Network](#)

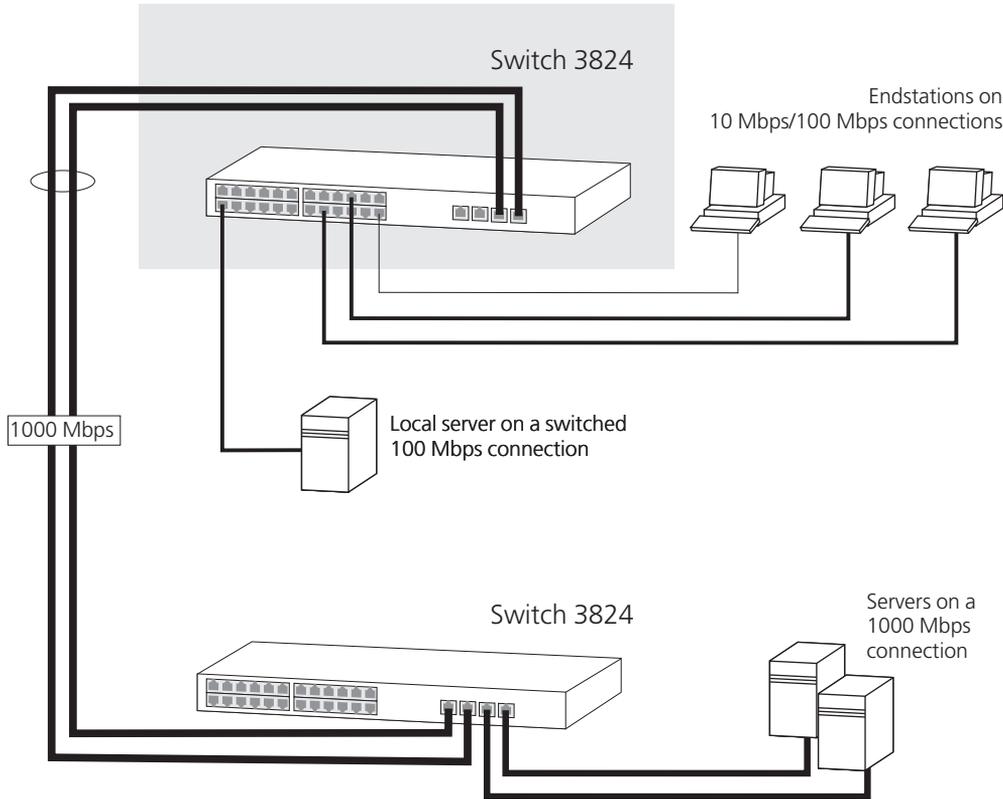
Simple Network Configuration Examples

The following illustrations show some simple examples of how the Switch 3812 and Switch 3824 can be used in your network.

Desktop Switch Example

The example in [Figure 21](#) shows how a Switch 3812 and Switch 3824 can be used for a group of users that require dedicated 10 Mbps 100 Mbps or 1000 Mbps connections to the desktop.

Figure 21 Using the Switch 3812 and Switch 3824 in a desktop environment



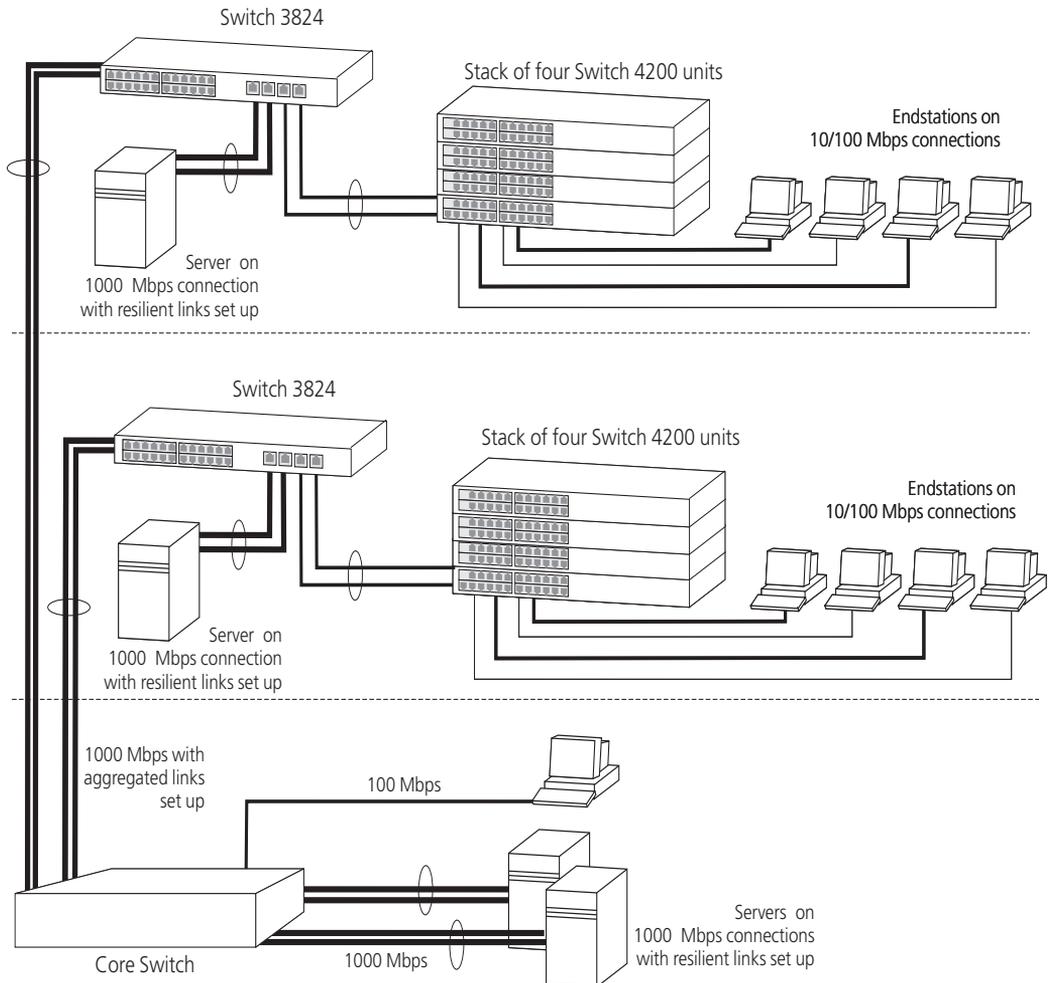
Advanced Network Configuration Examples

This section shows some network examples that illustrate how you can set up your network for optimum performance using some of the features supported by your Switch.

Improving the Resilience of Your Network

[Figure 22](#) shows how you can set up your network to improve its resilience using Spanning Tree Protocol (STP) and aggregated links also Aggregated links increase bandwidth available and also provide extra resilience.

Figure 22 Network set up to provide resilience

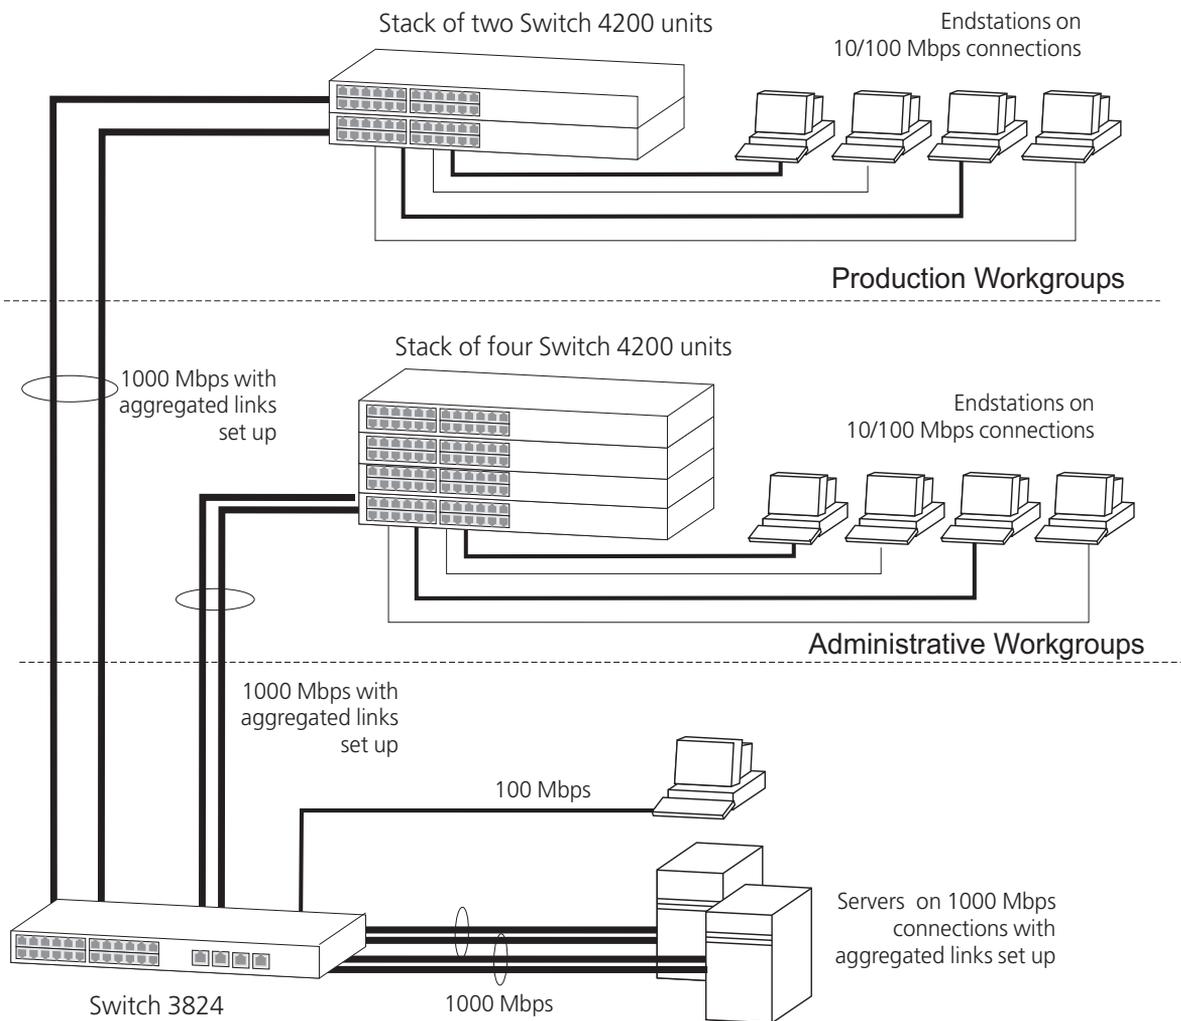


Enhancing the Performance of Your Network

Figure 23 shows how you can set your network up to enhance its performance.

All ports are auto-negotiating and will therefore pass data across the network at the optimum available speed and duplex mode. Flow control will help avoid packet loss during periods of network congestion. A Gigabit Ethernet backbone is set up between the Switch 3824 and each Switch in the workgroups to increase the bandwidth, and therefore the overall network performance.

Figure 23 Network set up to enhance performance



C

IP ADDRESSING

This chapter provides some background detail on the IP information that needs to be assigned to your Switch to enable you to manage it across a network. The topics covered are:

- [IP Addresses](#)
- [Subnets and Subnet Masks](#)
- [Default Gateways](#)



IP addressing is a vast topic and there are white papers on the World Wide Web and publications available if you wish to learn more about IP addressing.

IP Addresses

This IP address section is divided into two parts:

- [Simple Overview](#) — Gives a brief overview of what an IP address is.
- [Advanced Overview](#) — Gives a more in depth explanation of IP addresses and the way they are structured.

Simple Overview

To operate correctly, each device on your network must have a unique IP address. IP addresses have the format $n.n.n.n$ where n is a decimal number between 0 and 255. An example IP address is '192.168.100.8'.

The IP address can be split into two parts:

- The first part, called the network part, ('192.168' in the example) identifies the network on which the device resides.
- The second part, called the host part, ('100.8' in the example) identifies the device within the network.

If your network is internal to your organization only, you may use any arbitrary IP address. 3Com suggests you use addresses in the series

192.168.100.X (where X is a number between 1 and 254) with a subnet mask 255.255.255.0.



These suggested IP addresses are part of a group of IP addresses that have been set aside specially for use “in house” only.



CAUTION: *If your network has a connection to the external IP network, you must apply for a registered IP address. This registration system ensures that every IP address used is unique; if you do not have a registered IP address, you may be using an identical address to someone else and your network will not operate correctly.*

Obtaining a Registered IP Address

InterNIC Registration Services is the organization responsible for supplying registered IP addresses. The following contact information is correct at time of publication:

World Wide Web site: <http://www.internic.net>

Advanced Overview

IP addresses are 32-bit addresses that consist of a *network part* (the address of the network where the host is located) and a *host part* (the address of the host on that network).

Figure 24 IP Address: Network Part and Host Part



The boundary between network and host parts depends on the class of IP network.

IP addresses differ from Ethernet MAC addresses, which are unique hardware-configured 48-bit addresses. A central agency, such as the InterNIC Registration Services mentioned above, assigns the network part of the IP address, and you assign the host part. All devices that are connected to the same network share the same network part (also called the *prefix*).

Dotted Decimal Notation

The actual IP address is a 32-bit number that is stored in binary format. These 32 bits are segmented into 4 groups of 8 bits — each group is

referred to as a *field* or an *octet*. Decimal notation converts the value of each field into a decimal number, and the fields are separated by dots.

Figure 25 Dotted Decimal Notation for IP Addresses

10011110.01100101.00001010.00100000 = Binary notation
158.101.10.32 = Decimal notation



The decimal value of an octet whose bits are all 1s is 255.

Network Portion

The location of the boundary between the network part and the host part depends on the class that the central agency assigns to your network. The three primary classes of IP addresses are as follows:

- **Class A address** — Uses 8 bits for the network part and 24 bits for the host part. Although only a few Class A networks can be created, each can contain a very large number of hosts.
- **Class B address** — Uses 16 bits for the network part and 16 bits for the host part.
- **Class C address** — Uses 24 bits for the network part and 8 bits for the host part. Each Class C network can contain only 254 hosts, but many such networks can be created.

The high-order bits of the network part of the address designate the IP network class. See [Table 11](#).

Table 11 How Address Class Corresponds to the Address Number

Address Class	High-order Bits	Address Number (Decimal)
A	0nnnnnnn	0-127
B	10nnnnnn	128-191
C	11nnnnnn	192-254

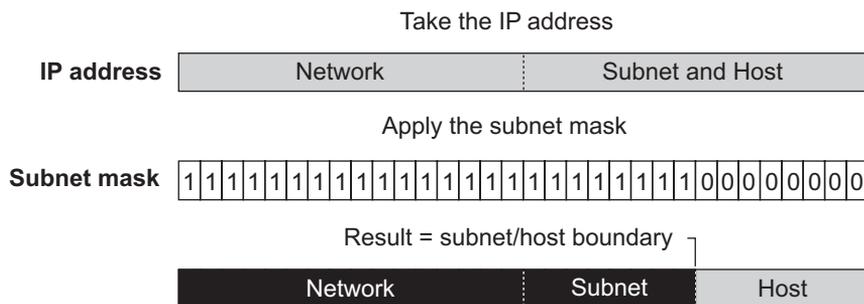
Subnets and Subnet Masks

You can divide your IP network into sub-networks also known as subnets. Support for subnets is important because the number of bits assigned to the device part of an IP address limits the number of devices that may be addressed on any given network. For example, a Class C address is restricted to 254 devices.

The IP address can also contain a *subnetwork part* at the beginning of the host part of the IP address. Thus, you can divide a single Class A, B, or C network internally, allowing the network to appear as a single network to other external networks. The subnetwork part of the IP address is visible only to hosts and gateways on the subnetwork.

When an IP address contains a subnetwork part, a *subnet mask* identifies the bits that constitute the subnetwork address and the bits that constitute the host address. A subnet mask is a 32-bit number in the IP address format. The 1 bits in the subnet mask indicate the network and subnetwork part of the address. The 0 bits in the subnet mask indicate the host part of the IP address, as shown in [Figure 26](#).

Figure 26 Subnet Masking



[Figure 27](#) shows an example of an IP address that includes network, subnetwork, and host parts. Suppose the IP address is *158.101.230.52* with a subnet mask of *255.255.255.0*. Since this is a Class B address, this address is divided as follows:

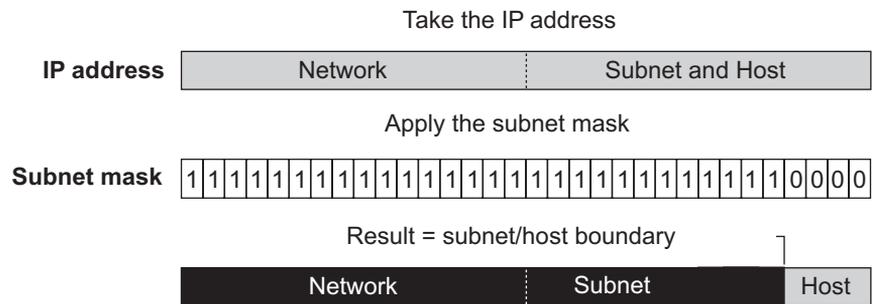
- *158.101* is the network part
- *230* is the subnetwork part
- *52* is the host part



As shown in this example, the 32 bits of an IP address and subnet mask are usually written using an integer shorthand. This notation translates four consecutive 8-bit groups (octets) into four integers that range from 0 through 255. The subnet mask in the example is written as 255.255.255.0.

Traditionally, subnet masks were applied to octets in their entirety. However, one octet in the subnet mask can be further subdivided so that part of the octet indicates an *extension* of the network number, and the rest of the same octet indicates the host number, as shown in [Figure 27](#).

Figure 27 Extending the Network Prefix



Using the Class B IP address from [Figure 26](#) (158.101.230.52), the subnet mask is 255.255.255.240.

The number that includes both the Class B natural network mask (255.255) and the subnet mask (255.240) is sometimes called the *extended network prefix*.

Continuing with the previous example, the subnetwork part of the mask uses 12 bits, and the host part uses the remaining 4 bits. Because the octets are actually binary numbers, the number of subnetworks that are possible with this mask is 4,096 (2^{12}), and the number of hosts that are possible in each subnetwork is 16 (2^4).

Subnet Mask Numbering

An alternate method to represent the subnet mask numbers is based on the number of bits that signify the network portion of the mask. Many Internet Service Providers (ISPs) now use this notation to denote the subnet mask. See [Table 12](#).

Table 12 Subnet Mask Notation

Standard Mask Notation	Network Prefix Notation
100.100.100.100 (255.0.0.0)	100.100.100.100/8
100.100.100.100 (255.255.0.0)	100.100.100.100/16
100.100.100.100 (255.255.255.0)	100.100.100.100/24



The subnet mask 255.255.255.255 is reserved as the default broadcast address.

Default Gateways

A gateway is a device on your network which is used to forward IP packets to a remote destination. An alternative name for a gateway is a Router. “Remote” refers to a destination device that is not directly attached to the same network segment as the source device.

The source device cannot send IP packets directly to the destination device because it is in a different network segment. Instead you configure it to send the packets to a gateway which is attached to multiple segments.

When it receives the IP packets, the gateway determines the next network hop on the path to the remote destination, and sends the packets to that hop. This could either be the remote destination or another gateway closer towards the destination.

This hop-by-hop process continues until the IP packets reach the remote destination.

If manually configuring IP information for the Switch, enter the IP address of the default gateway on the local subnet in which the Switch is located. If no default gateway exists on your network, enter the IP address `0.0.0.0` or leave the field blank.

GLOSSARY

3Com Network Supervisor	The 3Com network management application used to manage 3Com's networking solutions.
10BASE-T	The IEEE specification for 10 Mbps Ethernet over Category 3, 4 or 5 twisted pair cable.
100BASE-FX	The IEEE specification for 100 Mbps Fast Ethernet over fiber-optic cable.
100BASE-TX	The IEEE specification for 100 Mbps Fast Ethernet over Category 5 twisted-pair cable.
1000BASE-T	The IEEE specification for 1000 Mbps Gigabit Ethernet over four-pair Category 5 twisted-pair cable.
1000BASE-SX	The IEEE specification for 1000 Mbps Gigabit Ethernet over fiber-optic cable.
aging	The automatic removal of dynamic entries from the Switch Database which have timed-out and are no longer valid.
Aggregated Links	Aggregated links allow a user to increase the bandwidth and resilience between switches by using a group of ports to carry traffic between the switches.
auto-negotiation	A feature on twisted pair ports that allows them to advertise their capabilities for speed, duplex and flow control. When connected to a port that also supports auto-negotiation, the link can automatically configure itself to the optimum setup.
backbone	The part of a network used as a primary path for transporting traffic between network segments.
bandwidth	The information capacity, measured in bits per second, that a channel can transmit. The bandwidth of Ethernet is 10 Mbps, the bandwidth of

Fast Ethernet is 100 Mbps, and the bandwidth of Gigabit Ethernet is 1000 Mbps.

baud The signalling rate of a line, that is, the number of transitions (voltage or frequency changes) made per second. Also known as *line speed*.

bridge A device that interconnects two LANs of a different type to form a single logical network that comprises of two network segments. Bridges learn which endstations are on which network segment by examining the source addresses of packets. They then use this information to forward packets based on their destination address. This process is known as filtering.

broadcast A packet sent to all devices on a network.

broadcast storm Multiple simultaneous broadcasts that typically absorb all the available network bandwidth and can cause a network to fail. Broadcast storms can be due to faulty network devices.

cache Stores copies of frequently accessed objects locally to users and serves them to users when requested.

collision A term used to describe two colliding packets in an Ethernet network. Collisions are a part of normal Ethernet operation, but a sudden prolonged increase in the number of collisions can indicate a problem with a device, particularly if it is not accompanied by a general increase in traffic.

CSMA/CD Carrier-sense Multiple Access with Collision Detection. The protocol defined in Ethernet and IEEE 802.3 standards in which devices transmit only after finding a data channel clear for a period of time. When two devices transmit simultaneously, a collision occurs and the colliding devices delay their retransmissions for a random length of time.

DHCP Dynamic Host Control Protocol. A protocol that lets you centrally manage and automate the assignment of Internet Protocol (IP) addresses in an organization's network.

DNS Domain Name System. This system maps a numerical Internet Protocol (IP) address to a more meaningful and easy-to-remember name. When you need to access another device on your network, you enter the name of the device, instead of its IP address.

endstation	A computer, printer or server that is connected to a network.
Ethernet	A LAN specification developed jointly by Xerox, Intel and Digital Equipment Corporation. Ethernet networks use CSMA/CD to transmit packets at a rate of 10 Mbps over a variety of cables.
Ethernet address	See <i>MAC address</i> .
Fast Ethernet	An Ethernet system that is designed to operate at 100Mbps.
forwarding	The process of sending a packet toward its destination using a networking device.
Forwarding Database	See <i>Switch Database</i> .
filtering	The process of screening a packet for certain characteristics, such as source address, destination address, or protocol. Filtering is used to determine whether traffic is to be forwarded, and can also prevent unauthorized access to a network or network devices.
flow control	A mechanism that prevents packet loss during periods of congestion on the network. Packet loss is caused when devices send traffic to an already overloaded port on a Switch. Flow control prevents packet loss by inhibiting devices from generating more traffic until the period of congestion ends.
FTP	File Transfer Protocol. A protocol based on TCP/IP for reliable file transfer.
full duplex	A system that allows packets to be transmitted and received at the same time and, in effect, doubles the potential throughput of a link.
gateway	See router .
Gigabit Ethernet	IEEE standard 802.3z for 1000 Mbps Ethernet; it is compatible with existing 10/100 Mbps Ethernet standards.
half duplex	A system that allows packets to be transmitted and received, but not at the same time. Contrast with <i>full duplex</i> .
hub	A device that regenerates LAN traffic so that the transmission distance of that signal can be extended. Hubs are similar to repeaters, in that they connect LANs of the same type; however they connect more LANs than a repeater and are generally more sophisticated.

- HTTP** Hypertext Transfer Protocol. This is a set of rules for exchanging files (text, graphic images, sound, video, and other multimedia files) on the World Wide Web.
- IEEE** Institute of Electrical and Electronics Engineers. This American organization was founded in 1963 and sets standards for computers and communications.
- IEEE Std 802.1D, 1998 Edition** A standard that defines the behavior of bridges in an Ethernet network.
- IEEE Std 802.1p** A standard that defines traffic prioritization. 802.1p is now incorporated into the relevant sections of the IEEE Std 802.1D, 1998 Edition.
- IEEE Std 802.1Q-1998** A standard that defines VLAN tagging.
- IEEE Std 802.3ad** A standard that defines link aggregation. 802.3ad is now incorporated into the relevant sections of the IEEE Std 802.3-2002.
- IEEE Std 802.3x** A standard that defines a system of flow control for ports that operate in full duplex. 802.3x is now incorporated into the relevant sections of the IEEE Std 802.3-2002.
- IEEE Std 802.1w-2001** A standard that defines Rapid Spanning Tree Protocol (RSTP) behavior.
- IEEE Std 802.1X-2001** A standard that defines port-based network access control behavior.
- IETF** Internet Engineering Task Force. An organization responsible for providing engineering solutions for TCP/IP networks. In the network management area, this group is responsible for the development of the SNMP protocol.
- IGMP snooping** A mechanism performed by an intermediate device, such as a Layer 2 Switch, that optimizes the flow of multicast traffic. The device listens for IGMP messages and build mapping tables and associated forwarding filters, in addition to reducing the IGMP protocol traffic.
- Internet Group Management Protocol** Internet Group Management Protocol (IGMP) is a protocol that runs between hosts and their immediate neighboring multicast routers. The protocol allows a host to inform its local router that it wishes to receive transmissions addressed to a specific multicast group. Based on group membership information learned from the IGMP, a router is able to

determine which if any multicast traffic needs to be forwarded to each of its subnetworks.

- Intranet** An Intranet is an organization wide network using Internet protocols such as web services, TCP/IP, HTTP and HTML. An Intranet is normally used for internal communication and information, and is not accessible to computers on the wider Internet.
- IP** Internet Protocol. IP is a layer 3 network protocol that is the standard for sending data through a network. IP is part of the TCP/IP set of protocols that describe the routing of packets to addressed devices.
- IPX** Internetwork Packet Exchange. IPX is a layer 3 and 4 network protocol designed for networks that use Novell® Netware®.
- IP address** Internet Protocol address. A unique identifier for a device attached to a network using TCP/IP. The address is written as four octets separated with periods (full-stops), and is made up of a network section, an optional subnet section and a host section.
- Jitter** An expression often used to describe the end-to-end delay variations during the course of a transmission. See also *latency*.
- LAN** Local Area Network. A network of endstations (such as PCs, printers, servers) and network devices (hubs and switches) that cover a relatively small geographic area (usually not larger than a floor or building). LANs are characterized by high transmission speeds over short distances (up to 1000 m).
- LLC** Logical Link Control. A sublayer of the IEEE data link layer that is located above the MAC sublayer. The LLC sublayer is responsible for MAC sublayer addressing, flow control, error control, and framing.
- latency** The delay between the time a device receives a packet and the time the packet is forwarded out of the destination port.
- line speed** See *baud*.
- loop** An event that occurs when two network devices are connected by more than one path, thereby causing packets to repeatedly cycle around the network and not reach their destination.
- MAC** Media Access Control. A protocol specified by the IEEE for determining which devices have access to a network at any one time.

- MAC address** Media Access Control address; also called hardware or physical address. A layer 2 address associated with a particular network device. Most devices that connect to a LAN have a MAC address assigned to them as they are used to identify other devices in a network. MAC addresses are 6 bytes long.
- main port** The port in a resilient link that carries data traffic in normal operating conditions.
- MDI** Medium Dependent Interface. An Ethernet port connection where the transmitter of one device is connected to the receiver of another device.
- MDI-X** Medium Dependent Interface Cross-over. An Ethernet port connection where the internal transmit and receive lines are crossed.
- MIB** Management Information Base. A collection of information about the management characteristics and parameters of a networking device. MIBs are used by the Simple Network Management Protocol (SNMP) to gather information about the devices on a network. The Switch contains its own internal MIB.
- multicast** A packet sent to a specific group of endstations on a network.
- multicast filtering** A system that allows a network device to only forward multicast traffic to an endstation if it has registered that it would like to receive that traffic.
- NIC** Network Interface Card. A circuit board installed in an endstation that allows it to be connected to a network.
- POST** Power On Self Test. An internal test that a Switch carries out when it is powered-up.
- protocol** A set of rules for communication between devices on a network. The rules dictate format, timing, sequencing and error control.
- RADIUS** Remote Authentication Dial-In User Service. An industry standard protocol for carrying authentication, authorization and configuration information between a network device and a shared authentication server.

Rapid Spanning Tree Protocol	An enhanced version of the Spanning Tree Protocol that allows faster determination of Spanning Tree topology throughout the bridged network.
repeater	A simple device that regenerates LAN traffic so that the transmission distance of that signal can be extended. Repeaters are used to connect two LANs of the same network type.
resilient link	A pair of ports that can be configured so that one takes over data transmission should the other fail. See also <i>main port</i> and <i>standby port</i> .
RMON	IETF Remote Monitoring MIB. A MIB that allows you to remotely monitor LANs by addressing up to nine different groups of information.
router	A router is a device on your network which is used to forward IP packets to a remote destination. An alternative name for a router is a gateway.
RSTP	See <i>Rapid Spanning Tree Protocol</i> .
SAP	Service Access Point. A well-defined location that identifies the user of services of a protocol entity.
segment	A section of a LAN that is connected to the rest of the network using a switch or bridge.
server	A computer in a network that is shared by multiple endstations. Servers provide endstations with access to shared network services such as computer files and printer queues.
SMTP	Simple Mail Transfer Protocol. An IETF standard protocol used for transferring mail across a network reliably and efficiently (as defined in RFC 821).
SNMP	Simple Network Management Protocol. The current IETF standard protocol for managing devices on an TCP/IP network.
Spanning Tree Protocol (STP)	A bridge-based system for providing fault tolerance on networks. STP works by allowing you to implement parallel paths for network traffic, and ensure that redundant paths are disabled when the main paths are operational and enabled if the main paths fail.
standby port	The port in a resilient link that takes over data transmission if the main port in the link fails.

- STP** See *Spanning Tree Protocol (STP)*.
- subnet mask** A subnet mask is used to divide the device part of the IP address into two further parts. The first part identifies the subnet number. The second part identifies the device on that subnet.
- switch** A device that interconnects several LANs to form a single logical LAN that comprises of several LAN segments. Switches are similar to bridges, in that they connect LANs of a different type; however they connect more LANs than a bridge and are generally more sophisticated.
- Switch Database** A database that is stored by a switch to determine if a packet should be forwarded, and which port should forward the packet if it is to be forwarded. Also known as Forwarding Database.
- TCP/IP** Transmission Control Protocol/Internet Protocol. This is the name for two of the most well-known protocols developed for the interconnection of networks. Originally a UNIX standard, TCP/IP is now supported on almost all platforms, and is the protocol of the Internet. TCP relates to the content of the data travelling through a network — ensuring that the information sent arrives in one piece when it reaches its destination. IP relates to the address of the endstation to which data is being sent, as well as the address of the destination network.
- Telnet** A TCP/IP application protocol that provides a virtual terminal service, letting a user log into another computer system and access a device as if the user were connected directly to the device.
- TFTP** Trivial File Transfer Protocol. Allows you to transfer files (such as software upgrades) from a remote device using the local management capabilities of the Switch.
- traffic prioritization** A system which allows data that has been assigned a high priority to be forwarded through a switch without being obstructed by other data.
- unicast** A packet sent to a single endstation on a network.
- VLAN** Virtual LAN. A group of location- and topology-independent devices that communicate as if they are on the same physical LAN.
- VLAN tagging** A system that allows traffic for multiple VLANs to be carried on a single link.

WAN Wide Area Network. A communications network that covers a wide area. A WAN can cover a large geographic area, and may contain several LANs within it.

INDEX

A

addresses
 classes 85
 IP 83
 aggregated links 12, 19
 aging time, definition 46
 alarm events 62
 Alarms (RMON group) 60, 61
 automatic IP configuration 72
 auto-negotiation 12, 18

B

Backup 15
 bandwidth 17
 BPDUs. *See* Bridge Protocol Data Units
 Bridge Identifier 37
 Bridge Protocol Data Units 37
 Broadcast Storm Control 15

C

cable
 maximum length 76, 77
 Capture (RMON group) 61
 Configuration
 Restore 15
 Save 15
 Contents 3
 conventions
 notice icons, About This Guide 8
 text, About This Guide 8

D

default gateway 88
 Default VLAN 65
 Designated Bridge 38
 Designated Bridge Port 38

E

event notification 14
 Events (RMON group) 60, 61
 extended network prefix 87

F

Fast Ethernet configuration rules 76
 Filter (RMON group) 60, 61
 flow control 18
 full duplex configuration rules 77

G

Gigabit Ethernet configuration rules 75
 glossary 89

H

Hello BPDUs 39
 History (RMON group) 60, 61
 Hosts (RMON group) 61
 Hosts Top N (RMON group) 61

I

IEEE Std 802.1Q-1998 65
 IEEE Std 802.3-2002 flow control 13, 18
 IGMP
 default setting 29
 query mode 29
 snooping mode 29
 IGMP multicast filtering 30
 Internet
 addresses 83
 InterNIC 84
 IP (Internet Protocol)
 addresses 84
 IP address 72, 83
 classes of 85
 defined 84
 derivation 84
 division of network and host 84
 example 86
 obtaining 84
 subnet mask 86
 subnetwork portion 86
 IP multicast
 addressing 27
 IP routing
 address classes 85

L

learned SDB entries 46

M

MAC (Media Access Control)

addresses

IP address 84

manual configuration 72

masks

subnet 86

Matrix (RMON group) 61

Max Age 39

multicast filtering 27

IGMP 30

multicasts, description 27

N

network

addresses 83

network configuration examples 80, 81

non-aging learned SDB entries 46

O

obtaining

registered IP address 84

P

path costs. *See* port costs

permanent SDB entries 46

port costs, default 38

port security 12

port trunks

example 24

priority in STP 37

Q

QoS (see Quality of Service) 14, 47

Quality of Service 14, 47

R

Rapid Spanning Tree Protocol (RSTP) 13, 35

registered IP address, obtaining 84

Remote Monitoring. *See* RMON

Restore 15

RMON 14

alarm events 62

benefits 61

groups 59

Root Bridge 37

Root Path Cost 38

Root Port 38

S

Save 15

SDB. *See* Switch Database

segment, maximum length 76

Spanning Tree Protocol (STP) 13

Spanning Tree Protocol, *see* STP 34

Statistics (RMON group) 60, 61

STP 34

avoiding the subdivision of VLANs 43

Bridge Identifier 37

Bridge Protocol Data Units 37

default port costs 38

default priority 37

Designated Bridge 38

Designated Bridge Port 38

example 39

Hello BPDUs 39

Max Age 39

priority 37

Root Bridge 37

Root Path Cost 38

Root Port 38

using on a network with multiple VLANs 43

subnet mask 86

defined 86

example 86

numbering 87

subnets 86

subnetworking

defined 86

subnet mask 86

sub-networks. *See* subnets

Switch Database 45

T

- topology rules for Fast Ethernet 76
- topology rules with full duplex 77
- traffic prioritization 14, 47, 48
 - advanced 54
 - basic 52
 - classification 49
 - default configurations 58
 - differentiated services 50, 51
 - DiffServ Code Point (DSCP) 51
 - IEEE Std 802.1D, 1998 Edition 50
 - marking 50
 - queues 56
 - re-marking 52
 - rules, application-based 47
 - rules, device-based 47
 - traffic queues 48

V

- VLANs 63
 - benefits 64
 - Default 65
 - defining the information for 66
 - IEEE Std 802.1Q-1998 65

