User Guide for the Avaya P580 and P882 Multiservice Switches

Software Version 6.1

Doc. No. 10-300077
Issue 2
May 2004

User Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1

© Copyright Avaya Inc., 2004 ALL RIGHTS RESERVED

Produced in USA, May 2004

The products, specifications, and other technical information regarding the products contained in this document are subject to change without notice. All information in this document is believed to be accurate and reliable, but is presented without warranty of any kind, express or implied, and users must take full responsibility for their application of any products specified in this document. Avaya disclaims responsibility for errors which may appear in this document, and it reserves the right, in its sole discretion and without notice, to make substitutions and modifications in the products and practices described in this document.

P550R is a registered trademark of Avaya Inc.

Microsoft, Windows, Windows NT, Windows 95, Windows 98, and Internet Explorer are trademarks or registered trademarks of Microsoft Corporation in the U.S. and/or other countries.

Netscape and Netscape Navigator are registered trademarks of Netscape Communications Corporation in the United States and other countries.

Sybase is a registered trademark of Sybase, Inc.

Novell, NDS, Netware, and Novell Directory Services are registered trademarks of Novell, Inc.

Solaris is a trademark of Sun Microsystems, Inc.

Intel and Pentium are registered trademarks of Intel Corporation.

ALL OTHER TRADEMARKS MENTIONED IN THIS DOCUMENT ARE PROPERTY OF THEIR RESPECTIVE OWNERS.

Table of Contents

Preface . . . . . ***xvii***

Preface Contents xvii
Overview of User Guide Contents. xvii
Technical Support xx
Avaya Products and Services. xx
P580 and P882 Documentation xx
Available Documentation xx
Installing the Online Help Files xxi
Setting the Location of the Online Help Files xxi
Document Conventions xxii
Notes, Cautions, and Warnings xxiii
Terminology xxiv
Documentation Feedback xxiv

Chapter 1 — Introduction ***1-1***

Overview 1-1
Switch Descriptions 1-1
Hardware Components 1-2
Overview 1-2
Avaya P580 Multiservice Switch Chassis 1-3
Avaya P882 Multiservice Switch Chassis 1-5
The Power System 1-6
The M8000R-SUP Supervisor Module 1-8
Media Modules 1-10
Software Features 1-15
Security 1-16
VLANs 1-16
Hunt Groups 1-17
Spanning Tree 1-17
Extensive Fault Tolerance 1-18
Routing Protocols 1-19
Access Lists 1-19
IP Multicast 1-19
50-Series Buffers and Queues 1-20
80-Series QoS 1-20
System Management 1-21
Routing Overview 1-21
Compatibility with Layer 2 Modules 1-21
Routing with Layer 2 and Layer 3 Modules 1-22

Chapter 2 — Setting Up the Switch 2-1

Overview 2-1
Setting Up Your Laptop or PC 2-2
Performing Initial Setup of the Switch 2-2
Using Telnet to Manage the Switch 2-5
Setting up a Management Interface 2-6
Establishing a Telnet Session 2-6
Terminating a Telnet session 2-7
Viewing Active Telnet Sessions 2-7
Telnet Session Expiration Timer 2-7
Using the Web Agent to Configure the Switch 2-8
Opening the Web Agent 2-9
Logging Out of the Web Agent 2-11
Configuring Custom Access Types 2-12
Overview 2-12
Creating a Custom Access Type 2-13
Modifying a Custom Access Type 2-15
Deleting a Custom Access Type 2-16
Configuring User Accounts 2-17
Configuring User Account Security 2-17
Creating a User Account 2-19
Modifying a User Account 2-22
Changing Your Password 2-24
Disabling a User Account 2-25
Deleting a User Account 2-25
Changing the Console Serial Port Settings 2-26
Configuring the Serial Console Port as a TTY Console 2-27
Configuring the Serial Console Port as a PPP Console 2-29
Regaining Configuration Access to the PPP Serial Port Console 2-32
Configuring Dial-Up Networking 2-32
Using Dial-Up Networking with a PPP Serial Port Console 2-34
Changing the TCP Ports for HTTP and Telnet 2-35
Overview 2-35
Changing the TCP Port Number for HTTP and Telnet Requests 2-35
Starting a Telnet Session 2-37
Opening the Web Agent 2-37
Managing Configuration Files 2-38
Viewing the Running Configuration 2-39
Viewing the Startup Configuration 2-39
Viewing the Script Execution Log File 2-39
Copying Configuration Files 2-40
Copying Files 2-42
Viewing the Status of a TFTP Transfer 2-45

Chapter 3 — Configuring System Information 3-1

Overview 3-1
Entering General System Information 3-2

Enabling the Simple Network Time Protocol 3-3
Setting Summer Time Hours 3-5
Setting Recurring Summer Time Hours 3-5
Setting One-Time Summer Time Hours 3-10
Setting the System Clock 3-11
Setting the Temperature System 3-13
Displaying the Power System Statistics 3-15
Displaying Cooling System Statistics 3-16
Performing a Reset 3-17

Chapter 4 — Security 4-1

Overview 4-1
Secure Mode 4-1
Overview 4-1
Enabling Secure Mode 4-3
Disabling Secure Mode 4-3
Viewing the Secure Mode Setting 4-3
SSHv2 4-3
Overview 4-3
Enabling SSH on a TCP Port 4-4
Disabling SSH 4-5
Generating an SSH Server Key 4-6
Displaying the Public SSH Key 4-7
Configuring SSH Server 4-7
Displaying SSH Connections 4-9
Using the SSH Client 4-9
HTTPS Using SSLv3 or TLSv1 4-10
Overview 4-10
Supported Cipher Suites 4-11
Viewing Cipher Suites 4-12
Viewing the SSL Configuration 4-12
Creating a Self-Signed SSL Server Certificate 4-13
Viewing the Server Certificate 4-15
Enabling SSL/HTTPS 4-17
Disabling SSL/HTTPS 4-17
Restarting SSL/HTTPS 4-18
Reverting to a Backup Certificate 4-18
RADIUS Client Support 4-19
Overview 4-19
Realms and Groups 4-20
Login Order of Operations 4-21
Avaya VSAs 4-22
RADIUS Server Files 4-24
Configuring a RADIUS Client 4-26

Chapter 5 — Configuring SNMP 5-1

Overview 5-1
Authentication and Encryption 5-3

Localized Keys 5-3
Engine ID 5-3
Timeliness Checks 5-4
Views 5-4
Overview 5-4
Predefined Views 5-5
User Groups 5-5
Overview 5-5
Predefined Groups 5-6
Migration of Existing Community Strings 5-7
ATM Uplink Module 5-8
Initial SNMPv3 User and SNMPv3 Administrator 5-9
Changing the Engine ID of the Switch 5-10
Viewing the Engine ID 5-11
Creating an SNMPv3 Administrator 5-11
Configuring Views 5-12
Creating or Modifying a View 5-12
Viewing Configured Views 5-13
Deleting a View 5-13
Configuring Groups 5-13
Creating or Modifying a Group 5-14
Viewing Configured Groups 5-14
Deleting a Group 5-15
Configuring an SNMPv3 User 5-15
Creating or Modifying a User 5-15
Adding a User to a Group 5-16
Removing a User from a Group 5-17
Viewing Configured Users 5-17
Deleting a User 5-17
Changing a User Password 5-18
Configuring SNMPv1 or v2 Community Strings 5-18
Creating or Modifying a Community String for the Switch 5-18
Creating or Modifying a Community String for the ATM-Uplink Module 5-19
Viewing Configured Community Strings 5-20
Setting the Trap Receiver 5-20
Deleting a Community String 5-20
Setting the Administrative Contact 5-21
Setting the Physical Location of the Switch 5-22
Disabling or Reenabling SNMP 5-23
Viewing the SNMP Status 5-23

Chapter 6 — Using VLANs, Hunt Groups, and VTP Snooping 6-1

Overview 6-1
VLAN Introduction 6-1
What is a VLAN? 6-2
Ingress Rules 6-4
Forwarding Rules 6-5

Egress Rules 6-6
Creating and Implementing VLANs 6-7
VLAN Considerations 6-7
Creating a VLAN 6-7
Configuring VLAN Parameters 6-10
Assigning Ports to VLANs 6-13
Using Hunt Groups to Aggregate Bandwidth 6-13
Overview 6-13
Hunt Group Features 6-14
Hunt Group Terminology 6-15
How Hunt Groups Load Share 6-15
Hunt Group Example 6-17
Optimizing Hunt Group Throughput 6-19
Hunt Group Configuration Considerations 6-20
Hunt Groups and Multicast Traffic 6-21
Creating Hunt Groups 6-21
Adding Ports to the Hunt Group 6-23
Viewing Hunt Group Members 6-27
Removing Ports from Hunt Groups 6-28
Configuring VTP Snooping 6-29

Chapter 7 — Configuring Rapid Spanning Tree 7-1

Overview 7-1
How RSTP Achieves Rapid Recovery 7-3
Spanning Tree Configurations 7-4
Configuring Global Spanning Tree Options 7-7
Viewing Global Spanning Tree Information 7-10
Configuring Spanning Tree Bridges 7-11
Viewing Bridge Port Information 7-14
Configuring a Bridge Port 7-17
Configuring Spanning Tree Settings for Switch Ports 7-21

Chapter 8 — Configuring Ports 8-1

Overview 8-1
Two Categories of Port Settings 8-1
Physical Port Features 8-2
Switch Port Features 8-7
Relationship Between Different Switch Port Parameters 8-7
Assigning VLANs to a Port and Associated Issues 8-9
Configuring Physical Port Settings 8-11
Configuring Physical Ports on 10-Gigabit Ports 8-11
Configuring Physical Ports on Gigabit Ports 8-16
Configuring Physical Ports on Fast Ethernet Ports 8-21
Using the All Ports Configuration Web Page 8-28
Configuring Switch Ports Settings 8-28
Configuring Switch Ports 8-29
Automatic VLAN Creation 8-35
Configuring MAC Address Lock and Intrusion Detection 8-37

Using the All Module Ports Configuration Web Page 8-39
Viewing Switch Port Settings 8-41
Configuring Port Redundancy 8-42
Overview 8-42
Creating a Port Redundancy Pair 8-43
Deleting a Port Redundancy Pair 8-44
Disabling or Enabling Port Redundancy 8-45
GBIC Identification 8-46
Network Error Detection and Recovery 8-47
Overview 8-47
Enabling NEDR 8-48
Internal Error Detection and Recovery 8-51
Overview 8-51
Enabling IEDR 8-52
Enabling IEDR for Hunt Groups 8-52
Setting the IEDR Threshold for Internal Errors 8-53
Viewing IEDR Settings 8-53
Viewing IEDR Settings for Hunt Groups 8-53
Error Messages 8-54
System Messages 8-54

Chapter 9 — Managing the Address Forwarding Table 9-1

Overview 9-1
Relationship between VLANs, AFT and Hash Table Sizes 9-2
Hash Table guidelines for creating VLANs 9-3
Examples of Configuring VLANs and Hash Table Size 9-6
AFT Default Settings 9-7
Address Forwarding Table, Auto-Sizing, Auto Increment and Threshold 9-7
Total Entries, Address Memory, Age and Super-Age Timers 9-9
Overview 9-9
Configuring the AGE and Super Age Timers 9-9
Instance Table Information 9-10
Searching the Switch AFT 9-13
Adding Static Entries to the AFT 9-19
Option 82 for DHCP 9-21

Chapter 10 — Configuring Redundancy Options 10-1

Overview 10-1
Configuring Redundant Hardware 10-3
Redundant Switch Controllers and Elements 10-3
Installing and Enabling Redundant Hardware 10-5
Installing Redundant Hardware 10-5
Enabling Redundant Hardware 10-6
Replacing the Primary Controller 10-7
Replacing an Element 10-9
Configuring Supervisor Module Redundant Ethernet Console IP Addresses 10-10
Synchronizing the Active and Standby CPUs 10-12
Overview 10-12

Supervisor Module Redundancy Statistics 10-13
Active and Standby Roles. 10-13
Synchronization Procedure 10-15

Chapter 11 — Configuring DNS Client 11-1

Overview 11-1
DNS Client Support 11-1
Overview. 11-1
Sample Recursive DNS Query 11-3
DNS Client on the Avaya Multiservice Switch 11-4
Order of Operations for DNS on the Switch 11-5
Configuring the DNS Client. 11-6

Chapter 12 — Configuring IP Routing 12-1

Overview 12-1
Routing Function 12-2
Requirements for IP Routing 12-3
Hardware Requirements 12-3
Software Requirements. 12-3
Minimum Configuration Requirements 12-3
Routing Configuration Quickstart 12-4
Displaying Existing IP Interfaces. 12-4
Creating and Assigning IP Interfaces to the VLAN. 12-6
Enabling IP Routing Global Parameters. 12-10
IP Multinetting 12-14
Creating a Multinet Interface 12-15
Configuring Short-Lived IP Protocol Filters 12-15
Overview. 12-15
Enabling a Short-Lived IP Protocol Filter 12-16
Disabling a Short-Lived IP Protocol Filter. 12-17
Displaying Short-Lived IP Protocol Filters 12-18
Creating IP Static Routes 12-19
Creating a Static Route to a Null Interface. 12-20
Overview. 12-20
Procedure 12-22
Creating IP Static ARP Entries. 12-22
Creating a BOOTP/DHCP Server Entry 12-24
Configuring Route Redistribution 12-25
Overview. 12-25
Creating an IP Redistribute List Entry 12-26
Viewing IP Redistribute List Entries 12-28
Modifying an IP Redistribute List Entry 12-28
Deleting an IP Redistribute Entry. 12-29
IP Multicast 12-30
Configuring IGMP 12-31
Globally Enabling IGMP 12-31
Modifying IGMP Interfaces 12-32
Configuring and Modifying DVMRP. 12-35

Configuring the DVMRP Global Configuration 12-35
Modifying a DVMRP Interface 12-38
Monitoring Switch Performance Using IP Statistics 12-40
Displaying Global IP Routing Statistics 12-40
Searching the IP Routing Table 12-46
Displaying the IP Routing Table Statistics 12-48
Searching the IP ARP Cache 12-49
IP Multicast Statistics 12-51
Displaying IGMP Global Statistics 12-51
Displaying IGMP Interface Statistics 12-52
Displaying the IGMP Group Membership Table 12-55
Displaying the IGMP Local Multicast Forwarding Cache 12-56
Displaying DVMRP Global Statistics 12-57
Displaying DVMRP Interface Statistics 12-59
Displaying DVMRP Neighbor Routers 12-61
Displaying DVMRP Routing Table Statistics 12-62
Displaying the DVMRP Route Table 12-63
Displaying the DVMRP Upstream Routers 12-65
Displaying the DVMRP Designated Forwarder(s) Table 12-66
Displaying the DVMRP Downstream Dependent Routers 12-67
Displaying the DVMRP Multicast Forwarding Cache 12-67
Configuring VRRP 12-72
Globally Enabling VRRP 12-72
Enabling VRRP on an Interface 12-73
Creating a VRRP Virtual Router 12-74
Displaying VRRP Statistics 12-78
VRRP Configuration Considerations 12-80
Configuring IRDP 12-83
Enabling IRDP on an Interface 12-83
Configuring LDAP 12-86
Configuring LDAP Settings 12-86
Viewing LDAP Statistics 12-89
Configuring a Static Route for the PPP Console 12-90
Configuring the IP Interface for the PPP Console 12-91

Chapter 13 — Configuring Access Lists 13-1

Overview 13-1
Creating Access Rules 13-6
Creating Standard Access Rules 13-6
Creating Extended Access Rules 13-8
Enabling an Access List 13-12
Example: Using an Access Control Rule to Filter Web Traffic 13-13
Logging ACL Activity 13-15
Overview 13-15
Enabling ACL Logging 13-16
Setting a Logging Interval 13-17
Optimizing Switch Performance 13-18
Overview 13-18

Recognizing Performance Issues 13-20
Evaluating System Performance 13-21
Enabling Routing at the Module 13-23
Designing Safe, Efficient ACLs 13-24
Identifying the Ports 13-25
Configuring Hash Mode 13-26
Using Protocol or Port IDs in Access Rules 13-27
Managing F-chip Memory 13-28

Chapter 14 — Configuring RIP Routing 14-1

Overview 14-1
Configuring RIP on the Switch 14-1
Modifying RIP Interfaces 14-3
Creating Trusted RIP Neighbors 14-6
Key Chains 14-7
Overview 14-7
Key Chains on the Avaya Multiservice Switch 14-8
Configuring Key Chains 14-8
Viewing RIP Statistics 14-11
NBMA IP Interfaces 14-12

Chapter 15 — Configuring the OSPF Routing Protocol 15-1

Overview 15-1
What is OSPF? 15-1
Configuring OSPF 15-2
Configuring the OSPF Global Configuration 15-3
Creating OSPF Areas 15-5
Deleting OSPF Areas 15-7
Modifying OSPF Areas 15-8
Enabling OSPF on an IP Interface 15-8
Modifying OSPF Interfaces 15-9
Configuring an OSPF Passive-Interface 15-12
Creating OSPF Virtual Links 15-12
Deleting OSPF Virtual Links 15-15
Modifying OSPF Virtual Links 15-15
Creating OSPF Summaries 15-16
Deleting OSPF Summaries 15-17
Modifying OSPF Summaries 15-18
Monitoring Switch Performance Using OSPF Statistics 15-18
Displaying OSPF Statistics 15-19
Displaying OSPF Links 15-21
Displaying OSPF Neighbors 15-22
Searching the OSPF Link State Database 15-24
NBMA IP Interfaces 15-28
Overview 15-28
Setting Up an NBMA IP Interface 15-29
Setting Up NBMA Neighbors 15-30
Removing NBMA Neighbors 15-31

Chapter 16 — Configuring IPX Routing 16-1

Overview 16-1
IPX Overview 16-1
IPX Datagram Structure 16-2
Configuring the Avaya Switch as an IPX Router 16-3
Hardware Requirements 16-4
Configuring IPX Routing Globally 16-4
Configuring IPX Interfaces 16-5
Creating IPX Static Routes 16-9
Deleting IPX Static Routes 16-11
Modifying IPX Static Routes 16-12
Creating IPX Static Services 16-12
Deleting IPX Static Service 16-15
Modifying IPX Static Service 16-15
Monitoring Your Switch Using IPX 16-16
Displaying IPX Global Statistics 16-16
Searching the IPX Route Table 16-18
Displaying the IPX Route Table 16-20
Displaying IPX Route Table Statistics 16-20
Searching the IPX Service Table 16-21
Displaying the IPX Service Table 16-23
Displaying IPX Service Table Statistics 16-24

Chapter 17 — Configuring the IPX RIP Protocol 17-1

Overview 17-1
Configuring IPX RIP Interfaces 17-1
Creating and Modifying IPX RIP Filters 17-3
Viewing RIP Interface Statistics 17-8

Chapter 18 — Configuring the IPX SAP Protocol 18-1

Overview 18-1
Configuring IPX SAP Interfaces 18-1
Creating IPX SAP Filters 18-3
Creating IPX SAP Name Filters 18-4
Creating IPX SAP Network Filters 18-7
Interpreting IPX SAP Interface Statistics 18-12

Chapter 19 — Configuring AppleTalk Routing 19-1

Overview 19-1
Implementation 19-1
What is AppleTalk Routing? 19-1
Need for AppleTalk Routing 19-3
Configuring AppleTalk Routing 19-3
Enabling AppleTalk Global Routing 19-4
Creating an AppleTalk Routing Interface 19-4
Editing AppleTalk Interfaces 19-7

Deleting an AppleTalk Interface 19-9
Creating an AppleTalk Static Route 19-10
Editing AppleTalk Static Routes 19-12
Deleting an AppleTalk Static Route 19-13
Creating an AppleTalk Name-Binding Protocol (NBP) Filter 19-14
Editing an AppleTalk NBP Filter 19-16
Adding or Deleting Interfaces to an NBP Filter 19-17
Creating an AppleTalk Zone Filter 19-19
Editing an AppleTalk Zone Filter 19-20
Adding or Deleting Interfaces to a Zone Filter 19-21
Viewing AppleTalk Statistics 19-22
Viewing AppleTalk Global Statistics 19-23
Viewing the AppleTalk Interface Statistics Table 19-25
Viewing the AppleTalk Route Table 19-26
Viewing AppleTalk Route Table Statistics 19-28
Viewing the AppleTalk ARP Cache Table 19-29
Viewing the AppleTalk Zone Table 19-31
Viewing AppleTalk Zone Table Statistics 19-32
Viewing the AppleTalk NBP Table 19-33

Chapter 20 — Managing Intelligent Multicasting 20-1

Overview 20-1
Introduction 20-1
Manually Configured Intelligent Multicasting 20-2
Dynamic Intelligent Multicasting 20-2
Configuring Intelligent Multicasting 20-5
Configuring Global Intelligent Multicasting 20-5
Displaying Router Ports 20-8
Configuring Static Router Ports 20-9
Searching for Intelligent Multicast Sessions 20-10
Deleting an Intelligent Multicast Session 20-13
Deleting a Multicast Session Client Port 20-14
Creating a Static Multicast Session 20-15
Deleting Static Multicast Sessions 20-17
Creating Static Client Ports 20-18
Deleting Static Client Ports 20-19
Managing IGMP Snooping 20-20
Enabling IGMP Snooping 20-20
Viewing IGMP Snooping Statistics 20-21
Managing the LGMP Server 20-22
Configuring the LGMP Server 20-22
Viewing the LGMP Server Statistics 20-25
Viewing the LGMP Server Statistics per VLAN 20-26
Managing the LGMP Client 20-28
Enabling the LGMP Client 20-30
Viewing LGMP Clients Statistics 20-30
Viewing LGMP Clients per VLAN statistics 20-31
Managing CGMP Snooping 20-32

Enabling CGMP Snooping 20-33
Viewing CGMP Snooping 20-35

Chapter 21 — Monitoring the Avaya Multiservice Switch 21-1

Overview 21-1
Interpreting Front Panel LED Displays 21-1
Checking Active Alarms 21-2
Using the Event Subsystem 21-3
Overview 21-3
Configuring Event Notification 21-4
Configuring Notification of Protocol Events 21-8
Setting Log Size 21-10
Viewing the Event and Shutdown Logs 21-11
Clearing the Event Log 21-13
Viewing Event Statistics 21-13
Configuring Syslog Event Reporting 21-14
Overview 21-14
Configuring Syslog 21-15
Configuring Utilization Monitoring 21-17
Overview 21-17
Enabling Utilization Monitoring 21-17
Enabling Event Logging of Utilization 21-18
Setting a Utilization Threshold 21-19
Viewing Utilization Settings 21-20
Viewing CPU Utilization Statistics 21-20
Viewing Statistics for Forwarding Engine Utilization 21-21

Chapter 22 — Monitoring and Configuring the Forwarding Cache 22-1

Overview 22-1
Configuring the Forwarding Cache 22-2
Monitoring the Forwarding Cache Statistics 22-4
Displaying Frame Forwarding Statistics 22-4
Displaying and Searching the L3 Forwarding Cache for an Entry 22-6
Displaying the Forwarding Cache 22-9

Chapter 23 — Using RMON and Ethernet Statistics to Analyze Network Performance 23-1

Overview 23-1
Viewing Network Statistics 23-1
Setting Up Port Mirroring 23-7
Setting Up a Port Mirror on a Switch in Fabric Mode 1 23-8
Setting Up Port Mirroring on a Switch in Fabric Mode 2 23-12
Removing a Fabric Mode 2 Port Mirror 23-18
Viewing Information about a Fabric Mode 2 Port Mirror 23-18

Chapter 24 — Managing Buffers and Queues on 50-Series Modules 24-1

Overview 24-1
How Queues Work 24-1
Managing Buffers and Queues 24-2

Chapter 25 — 80-Series QoS 25-1

Overview 25-1
Why implement QoS? 25-2
How Does QoS Work? 25-3
QoS Process for Ingress Traffic 25-3
QoS Process for Egress Traffic 25-4
Diagram of QoS Process 25-6
Classification of Traffic 25-7
Default Priority 25-8
Classifying Traffic by Layer 2 Characteristics 25-9
Classifying Traffic by Layer 3 or Layer 4 Characteristics 25-9
Diffserv 25-11
Precedence of Priorities 25-11
Supported Number of Queues 25-12
Setting the Priority of a Physical Port 25-13
Setting a Physical Port to Ignore Tag Priority 25-14
Setting the Priority of a MAC Address 25-16
Displaying the Priority of a MAC Address 25-18
Setting a Physical Port to Use DiffServ 25-19
Setting a Physical Port to Mask DiffServ Bits 25-20
Assigning a Priority to a DSCP 25-21
Displaying the DiffServ Table 25-21
Displaying the QoS Settings for a Physical Port 25-22
Setting Up an ACL Rule 25-22
Setting Up a Default ACL Rule 25-28
Displaying ACL Rules 25-29
Ingress Policing 25-29
Setting Up Policing 25-30
Displaying the Policing Settings 25-32
Queue-Servicing Algorithms 25-33
WFQ 25-33
Strict Priority Queueing 25-35
CBQ 25-35
CBWFQ 25-35
Management Traffic 25-36
Setting Up WFQ 25-37
Setting Up Strict Priority Queueing 25-38
Setting Up CBQ 25-39
Setting Up CBWFQ 25-40
Displaying the Queue-Service Settings 25-43
QoS Statistics 25-44
Displaying QoS Statistics 25-50

Resetting the QoS Statistics 25-51
Displaying the Buffer Settings for Egress Queues 25-52

Appendix A — Upgrading the Application Software A-1

Overview A-1
Backing Up the Current Software A-2
Backing Up the Previous Configuration A-3
Downloading Application Software A-5
Setting the Startup Image A-7
Synchronizing the Active and Standby Supervisor Modules A-7
Resetting the Active Supervisor A-8
Resetting the Standby Supervisor A-9
Verifying the Upgrade A-10
Upgrading a Switch with Multiple VLANs
Statically Bound to Hunt Group Ports A-10

Appendix B — Boot Mode B-1

Accessing BOOT Mode B-1
Accessing BOOT Mode During Power Up B-1
Accessing BOOT Mode with Corrupted Operational Images B-2
Password Recovery B-2
Downloading New Operational Code B-3

Appendix C — Supported MIB Groups C-1

AppleTalk C-1
Application Software Management C-1
ATM Uplink C-2
Bridging C-2
DVMRP C-2
General Private MIBs C-2
IGMP C-2
IP C-3
IPX C-3
Monitoring C-4
SNMP C-4

Appendix D — FCC Notice D-1

Index IN-1

Preface

Preface Contents

This Preface contains the following information:

- Overview of User Guide Contents
- Technical Support
- Avaya Products and Services
- P580 and P882 Documentation
- Document Conventions
- Terminology
- Documentation Feedback

Overview of User Guide Contents

This guide provides information and procedures for configuring the Avaya P580 and P882 Multiservice switches. You can use the Web Agent, command line interface (CLI), or SNMP to configure the switches. This guide provides information for the Web Agent and CLI. The Avaya switches support both layer 2 and layer 3 functionality.

For the most current issue of this document and software downloads, see http://support.avaya.com.

*** Note:** For detailed information about the CLI, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.

This guide is divided in to the following chapters:

- Chapter 1, Introduction — Provides an overview of hardware and software used on the P580 and P882.
- Chapter 2, Setting Up the Switch — Explains how to initially configure the P580 and P882, create user accounts, configure protocols, manage configuration files, and configure the switches for dial-up networking.
- Chapter 3, Configuring System Information — Explains how to view and configure switch information such as system clock settings, and temperature, power, and cooling statistics.
- Chapter 4, Security — Provides information and procedures for configuring secure mode, SSHv2, and HTTPS Using SSLv3 or TLSv1.
- Chapter 5, Configuring SNMP —Provides information and procedures for configuring SNMPv1, v2, and v3.
- Chapter 6, Using VLANs, Hunt Groups, and VTP Snooping — Provides detailed information on how to create and configure VLANs, hunt groups, and VTP Snooping on your switch.
- Chapter 7, Configuring Rapid Spanning Tree — Provides information and procedures for configuring Rapid Spanning Tree Protocol (RSTP).
- Chapter 8, Configuring Ports — Explains how to configure physical port and switch port parameters on your switch.
- Chapter 9, Managing the Address Forwarding Table — Provides information on how to configure and manage the address forwarding tables in your switch.
- Chapter 10, Configuring Redundancy Options — Explains how to install, enable and configure redundancy hardware options on your switch.
- Chapter 11, Configuring DNS Client — Explains how to configure DNS and RADIUS clients on your switch.
- Chapter 12, Configuring IP Routing — Explains how to configure your switch for IP routing and interpret IP routing statistics.
- Chapter 13, Configuring Access Lists — Provides information and procedures for creating access rules, enabling an access list, logging ACL activity, and optimizing switch performance when an access list is enabled.
- Chapter 14, Configuring RIP Routing — Explains how to configure IP RIP on your switch.

- Chapter 15, Configuring the OSPF Routing Protocol — Explains how to configure OSPF routing protocol on your switch. Also provided is information on OSPF statistical displays.
- Chapter 16, Configuring IPX Routing — Explains how to configure IPX on your switch.
- Chapter 17, Configuring the IPX RIP Protocol — Explains how to configure IPX RIP protocol on your switch.
- Chapter 18, Configuring the IPX SAP Protocol — Explains how to configure IPX SAP on your switch.
- Chapter 19, Configuring AppleTalk Routing — Explains how to configure AppleTalk parameters and view AppleTalk statistics.
- Chapter 20, Managing Intelligent Multicasting — Explains how to configure and manage intelligent multicast sessions on your switch.
- Chapter 21, Monitoring the Avaya Multiservice Switch — Explains how to use the Web Agent to assess your switch's current operational status.
- Chapter 22, Monitoring and Configuring the Forwarding Cache — Explains how to interpret and monitor forwarding operations that occur in the address cache of the multilayer media modules.
- Chapter 23, Using RMON and Ethernet Statistics to Analyze Network Performance — Provides information on how to interpret the statistics counter values displayed in your switch.
- Chapter 24, Managing Buffers and Queues on 50-Series Modules — Explains how to use your switch's buffer management features to optimize traffic throughput through the switch fabric.
- Chapter 25, 80-Series QoS — Provides information about Quality of Service (QoS) and explains how to configure QoS on the switch.
- Appendix A, Upgrading the Application Software — Explains how to update the operational code on your switch.
- Appendix B, Boot Mode — Explains how to use boot mode on the Avaya Multiservice switches.
- Appendix C, Supported MIB Groups — Provides information on the supported MIB groups.
- Appendix D, FCC Notice — Provides the FCC notice statement.
- Index

Technical Support

To contact Avaya's technical support:

- **From the United States:**

  1-800-237-0016

- **From North America:**

  1-800-242-2121

- **Outside North America:**

  Contact your distributor

Avaya Products and Services

See the Avaya World Wide Web site at http://www.avaya.com for information about Avaya products and services.

P580 and P882 Documentation

This section contains the following information:

- Available Documentation
- Installing the Online Help Files
- Setting the Location of the Online Help Files

Available Documentation

Documents

The following documents provide additional information on the Avaya P580 and P882 Multiservice switches:

- *Installation Guide for the Avaya P550R, P580, P880, and P882 Multiservice Switches* describes how to install and set up the family of Avaya Multiservice switches.
- *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1* provides details and command syntax for the Command Line Interface (CLI).

- *Release Notes for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1* lists new software features, resolved issues, and known issues in v6.1 application software.
- *Functional Restrictions for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1* lists functional restrictions of switches that are running v6.1 application software. Functional restrictions are issues that restrict the functionality of a particular feature. For most of these issues a workaround exists.

For the most current issues of these documents and this User Guide, see http://support.avaya.com.

Online Help

You must have a browser installed on your system to view the online help.

To access the Web Agent online help, you must install the help files on an existing Web server in your network. Once the help files are installed, you must point the P580 or P882 to the location of the help files. The following two sections provide procedures for these tasks:

- Installing the Online Help Files
- Setting the Location of the Online Help Files

Installing the Online Help Files

To install the online help files on a Web server:

1. Download the online help from http://support.avaya.com to the directory where you want to store the help files.

   By default, the Web Agent is set to retrieve online help files from the **help** subdirectory of the Web server's root directory. If you install the online help files in a different directory, you must change the Web Agent's setting for the help file location. For information on changing this setting, see the following section, "Setting the Location of the Online Help Files."

2. Unzip the zip file.

Setting the Location of the Online Help Files

To set the location of the online help files:

1. Launch a Web browser and connect to the switch.
2. Enter your user name and password, and click **OK**.

3. In the navigation pane, expand the **System > Configuration** folders, and then click **Online Help**.

   The Online Help Configuration Web page is displayed in the content pane.

4. In the **HTTP Server Location** field, enter the URL of the Web server followed by the TCP port that the Web server uses. You can enter either the host name or IP address for the URL. For example, **http:// 10.10.10.10:80.**

   The default setting for the TCP port is 2010.

5. In the **HELP Directory Location** field, enter the directory where the files are stored. The default setting is **help**.

6. Click **APPLY**.

Document Conventions

The following conventions are used throughout this guide:

Convention	Represents	Examples
User Input	User entered text.	To create a new password, type **store password root**.
Boldface Text	A menu command, a keyword to be acted upon, or a button name.	In the **Interface** field, select **Default**. Click **Cancel** to stop the installation.
System Output	Text displayed by the system.	If you attempt to find the physical location of port 30, the system displays `Unit 2 Port 2`
"Save the running configuration to the startup configuration"	Always save the running configuration to startup configuration before you turn off the switch **Note:** Never save the startup configuration to the running configuration.	"Save the running configuration to the startup configuration and turn off the switch".
		1 of 2

Convention	Represents	Examples
Using the Web Agent.	The Web Agent consists of folders that you must expand to perform specific operations. The folders are located in the navigation pane of the Web Agent.	In the navigation pane, expand **Modules & Ports** folder, and then click **Configuration.**
Web Agent vs. Command Line Interface (CLI) steps.	The procedures in this guide provide detailed steps for the Web Agent. A series of CLI commands are also available that accomplish the same actions. These commands are listed after the Web Agent procedure. For more information about these commands, see the *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.	In the **Ports** field, select the port number of the module. To view switch port settings, use the following CLI command: > **show port <*mod-num*>**.
Knowledge level	All procedures in this guide are written with the assumption that the reader knows how to navigate a windows-type environment (use scroll bars, open pull-down menus, etc.).	Scroll to the “Configuring IP” section. Open the **In-band IPC** pull-down menu.
		2 of 2

Notes, Cautions, and Warnings

The following text and icon conventions are used in this guide for Notes, Cautions, and Warnings.

*** Note:** Provides additional information about a procedure or topic.

WARNING: Indicates a condition that may cause bodily injury or death.

CAUTION: Indicates a condition that *may damage* hardware or software.

Terminology

Throughout this guide, the term *layer 2* refers to switching capability. For example, *layer 2 module* is a module that provides switching capability.

The term *layer 3* refers to the combined ability to switch and route. For example, *layer 3 module* is a module that provides both switching and routing capabilities.

Documentation Feedback

If you have comments about the technical accuracy or general quality of this document, please send an e-mail message to techpubs@avaya.com.

Please cite the document title, part number (if applicable), and page number.

1 Introduction

Overview

The following information is provided in this chapter:

- Switch Descriptions
- Hardware Components
- Software Features
- Routing Overview

Switch Descriptions

The Avaya P580 and P882 Multiservice switches are high performance Gigabit Ethernet switches for LAN backbone applications. These switches are the centerpiece for Avaya data, voice, and multimedia LANs. The P580 switching capacity ranges from 46Gbps in Fabric mode 1 to 55Gbps in Fabric mode 2. The P882 switching capacity ranges from 56Gbps, in Fabric mode 1, to 139Gbps, in Fabric mode 2.

The Avaya Multiservice switch family supports specific hardware components and various software features, including:

- Increased bandwidth
- Elimination of bottlenecks
- Better manageability
- Routing
- Dependable multimedia support
- Redundant switch-to-switch trunks (or *hunt groups*)

Hardware Components

Overview

Chassis

The Avaya switch chassis comes in two models:

- The P580 chassis.
- The P882 chassis.

The P580 and P882 chassis can run in two modes:

- Fabric Mode 1 = 55MHz or Fabric_mode 1
- Fabric Mode 2 = 66Mhz or Fabric_mode 2

Fabric Mode 1

Fabric mode 1 has a switching fabric performance of 1.76 Gigabits per second per fabric port. Fabric mode 1 modules operate at 55MHz only.

In Fabric mode 1 the P580 chassis will have a 13x13 crossbar switching fabric, while the P882 chassis will have a 16x16 cross bar switching fabric. In addition, the P882 chassis is limited to sixteen functional slots in this mode (slot 17 can not be used).

Fabric Mode 2

The 80-series (also called *series 2*) platform is an evolution of the 50-series (also called *series 1*) architecture. It offers a 20% increase in switching fabric performance or 2.11 Gigabits per second per fabric port. This allows greater 10/100/1000 port density on a per crossbar port basis.

This increase in performance is accomplished with Avaya Application-Specific Integrated Circuits (ASICs) that operate at 66MHz (as well as 55MHz capable) and consolidates and enriches the functionality of a number of ASICs used in the 50-series modules.

If any 50-series modules are installed in the chassis, they do not turn on when the switch is in Fabric mode 2.

In Fabric mode 2, the P580 chassis will have a 13x13 crossbar switching fabric, while the P882 chassis will have a 33x33 cross bar switching fabric. All 17 slots in the P882 chassis can be used in Fabric 2 mode.

Contents

This section covers the following hardware components:

- Avaya P580 Multiservice Switch Chassis
- Avaya P882 Multiservice Switch Chassis

- The Power System
- The M8000R-SUP Supervisor Module
- Media Modules

Avaya P580 Multiservice Switch Chassis

The Avaya P580 Multiservice Switch chassis consists of:

- Seven slots (1 slot for the supervisor module and six payload slots for media modules)
- Fan Tray
- Up to three power supplies
- A 13x13 crossbar switch fabric

Slots

There are seven slots in the P580 chassis. Slot 1 is dedicated for a supervisor module and the remaining six are payload slots. If a redundant supervisor is used in the chassis, it must be inserted in slot 2.

Ports

The P580 chassis configured with 50-series modules provides support for:

- Up to 120 10/100BASE-TX ports (RJ-45 connector, autosensing)
- Up to 288 10/100BASE-TX ports (Telco connector, autosensing)
- Up to 60 100BASE-FX ports
- Up to 24 gigabit-speed Ethernet ports

The P580 chassis configured with 80-series modules provides support for:

- Up to 144 10/100BASE-TX ports (RJ-45 connector, autosensing)
- Up to 288 10/100BASE-TX ports (Telco connector, autosensing)
- Up to 144 100BASE-FX ports (MT-RJ connector)
- Up to 48 gigabit-speed Ethernet ports

Fan Tray

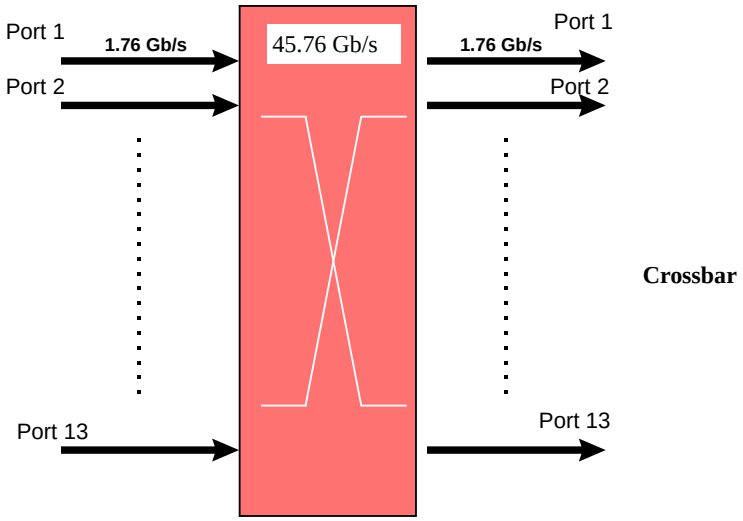
There are two fan assemblies in the chassis. One fan assembly is located on the left side of the chassis to provide air flow to the media modules. This fan assembly is hot-swappable. The other fan assembly is located in the rear of the chassis to provide air flow to the switch fabric. This fan assembly is not field replaceable.

Switch Fabric

In Fabric mode 1, the switching fabric provides 45.76 Gbps aggregate bandwidth. In Fabric mode 2 (P580 only), the switching fabric provides 54.91 Gbps aggregate bandwidth.

The crossbar switch matrix provides low latency, high throughput packet switching using a crossbar architecture (Figure 1-1).

Figure 1-1. Crossbar Architecture



The crossbar supports:

- 13 fabric ports (two per I/O module slot, one for the supervisor module).
- 1.76 Gbps (in and out) on each fabric port in 50-series mode and 2.11 Gbps (in and out) on each fabric port in Fabric mode 2.
- 45.76 Gbps backplane capacity in Fabric mode 1 and 54.91 Gbps backplane capacity in Fabric mode 2.
- Under-subscribed switching fabric in most configurations.
- Single copy replication—Input frames destined for multiple output switch ports pass through the crossbar only once and are copied by the crossbar to each destination.
- Hardware-assisted multicast pruning—The switch only forwards to appropriate destination switch ports.

Avaya P882 Multiservice Switch Chassis

The Avaya P882 Multiservice Switch Chassis consists of:

- Seventeen slots (1 slot for the supervisor module and sixteen payload slots for media modules)
- Fan Tray
- Up to three power supplies
- A 33x33 crossbar switch fabric

Slots

The P882 chassis provides seventeen slots. Slot 1 is dedicated for a supervisor module and the remaining sixteen are payload slots.

In Fabric mode 1, the P882 Chassis supports up to 15 payload slots. Slot seventeen is unsupported in Fabric mode 1. In Fabric mode 2, all sixteen payload slots are usable.

If a redundant supervisor is used in the chassis, the redundant supervisor must be installed in slot 2.

Ports

With 50-Series modules, the P882 Chassis supports:

- Up to 300 10/100BASE-TX ports (RJ-45 connector, autosensing)
- Up to 720 10/100BASE-TX ports (Telco connector, autosensing)
- Up to 150 100BASE-FX ports
- Up to 60 gigabit-speed Ethernet ports

With 80-Series modules, the P882 chassis supports:

- Up to 384 10/100BASE-TX ports (RJ-45 connector, autosensing)
- Up to 768 10/100BASE-TX ports (Telco connector, autosensing)
- Up to 384 100BASE-FX ports (MT-RJ connector)
- Up to 128 gigabit-speed Ethernet ports

Fan Tray

There are two fan assemblies in the chassis. One fan assembly is located on the left side of the chassis to provide air flow to the media modules. The other fan assembly is located in the rear of the chassis to provide air flow to the switch fabric. Both fan assemblies are hot-swappable and field replaceable.

Switch Fabric

In Fabric mode 1, the switching fabric provides a 16x16 crossbar with 56.32 Gbps aggregate bandwidth. In Fabric mode 2, the switching fabric provides a 33x33 crossbar with 139.392 Gbps aggregate bandwidth.

The crossbar switch matrix provides low latency, high throughput packet switching using a crossbar architecture.

The Avaya P882 Multiservice switch chassis crossbar supports:

- 33 fabric ports (two per I/O module slot, plus one for the supervisor module - Fabric mode 2).
- 1.76 Gb/s (in and out) on each fabric port in Fabric mode 1 and 2.11 Gb/s (in and out) on each fabric port in Fabric mode 2.
- 56.32 Gb/s backplane capacity in Fabric mode 1 and 139.392 Gb/s backplane capacity in Fabric mode 2.
- Under-subscribed switching fabric in most configurations.
- Single copy replication—Input frames destined for multiple output switch ports pass through the crossbar only once and are copied by the crossbar to each destination.
- Hardware-assisted multicast pruning—The switch only forwards to appropriate destination switch ports

The Power System

Overview

Two power supplies are required to operate a fully-loaded P580 or P882 chassis. You can install a third power supply for backup if one of the other two fails.

Features of the power system include:

- Load sharing by the installed power supplies.
- Power supplies can be warm swapped. (Before removing a power supply, you must turn it off and unplug it; however, the other power supplies can remain turned on).
- Both AC and DC power supplies are supported.
- Static power allocation or dynamic power allocation if insufficient power is available for all modules in the chassis.

Static Power Allocation

The switch uses *static power allocation* at startup when the switch has insufficient power to turn on all modules in the chassis.

When using static power allocation, the switch turns on the module in slot 2 (slot 3 if slot 2 contains a redundant supervisor) and then turns on each consecutive module for which sufficient power is available. When the switch reaches a module for which insufficient power is available, the switch stops attempting to turn on modules. All modules below the first module for which insufficient power is available remain turned off.

For example:

1. You turn on a P882 that has only 700 watts available.

- Modules 1 through 10 require 660 watts.
- Module 11 requires 50 watts.

2. Modules 1 through 10 are turned on and operate normally.

3. Modules 11 through 17 remain turned off.

The remaining 40 watts of available power is unused.

Dynamic Power Allocation

The switch uses *dynamic power allocation* when the amount of available power changes. For example, one power supply is turned on and you add another power supply or two power supplies are turned on and you turn one off. Unlike static power allocation, dynamic power allocation maximizes the number of modules that are turned on.

When using dynamic power allocation, the switch first turns off all modules in the chassis and then turns on the modules for which sufficient power is available. The module in slot 2 (slot 3 if slot 2 contains a redundant supervisor) is turned on first and then each consecutive module for which sufficient power is available. When the switch reaches a module for which insufficient power is available, that module remains turned off, and the switch searches for and turns on the next module for which sufficient power is available.

For example:

1. One power supply fails and the switch has only 700 watts of power available.

- Modules 1 through 10 consume 660 watts.
- Module 11 and 12 each require 50 watts. But module 13 requires only 40 watts.

2. Modules 1 through 10 are turned on and operate normally.

3. Modules 11 and 12 remain turned off.
4. Module 13 is turned on.
5. Modules 14 through 17 remain turned off because all available power is now consumed.

The M8000R-SUP Supervisor Module

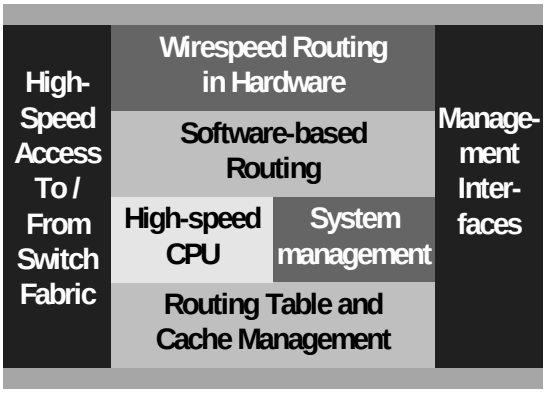
Functions

The M8000R-SUP supervisor module is responsible for learning addresses and managing the address cache, as well as protocol management and switch management.

The supervisor module is part of the path that some packets take through the system.

Figure 1-2 illustrates a conceptual diagram of the supervisor module's functions.

Figure 1-2. Layer 3 Supervisor Conceptual Operation



Multilayer Supervisor Conceptual Diagram

Hardware Features

Hardware features of the M8000R-SUP supervisor module are:

- PowerPC 750 (RISC) processor
- Memory: 4 MB Flash, 256 MB DRAM, 128 KB NVRAM, 512 KB cache
- Multiple memory configurations (see Table 1-1)

Table 1-1. Supported Memory, Boot Code, and Software Configurations

Memory (in MB)	Boot Code	Software Version	Switch Boots?	Support of Configuration
64	v5.00.01	Earlier than 6.0	Boots	Supported
64	v5.00.01	6.0 or later	Does not boot	Unsupported
64	v6.00.00	Earlier than 6.0	Boots	Unsupported
64	v6.00.00	6.0 or later	Boots, but displays the following error message: WARNING! The current version of application code (v6.01.00) requires a memory upgrade to maintain proper operation of the switch. Failure to upgrade will result in switch instability. Please refer to the user documentation for detailed instructions.	Unsupported
256	v5.00.01	Earlier than 6.0	Boots	Supported *However, the switch will use only 64 MB of memory and display only 64 MB for total memory.
256	v5.00.01	6.0 or later	Does not boot	Unsupported
256	v6.00.00	Earlier than 6.0	Boots	Unsupported
256	v6.00.00	6.0 or later	Boots	Supported

- Real-time clock
- Out-of-band console: 10/100BASE-T and RS-232
- Dot matrix display
- System management functions and interfaces
- 1.5 million packets per second of hardware-based routing for packets that arrive on layer 2-only media modules

Redundant Supervisor

The redundant supervisor module is an auxiliary (standby) supervisor module that acts as a fault-tolerant supervisor in the event that the active supervisor fails. The redundant supervisor provides a seamless solution to switch failure. Once the redundant supervisor is installed, loaded with the same software version as the active supervisor, and synchronized, it is ready to act as a backup to the active supervisor. If the active supervisor does fail, the redundant module quickly assumes control with the least possible effect on network operation. No user intervention is required for the CPU module switchover. The management view is accessible without changing IP or MAC addresses.

For more information about the redundant supervisor module, see Chapter 10, "Configuring Redundancy Options."

Media Modules

The layer 2 and layer 3 media modules implement wire speed routing and bridging in ASICs. The routing and bridging ASICs can process 1.5 million packets per second of minimum-sized Ethernet frames.

*** Note:** All layer 3 modules interoperate with layer 2 modules.

The ATM Uplink module provides LAN Emulation (LANE) connectivity over an ATM network and supports RFC 1483.

Bridging and routing are performed on the input side of each media module. Frames whose destination Media Access Control (MAC) address is on the same Virtual LAN (VLAN) as the interface on which they arrived are bridged. Frames that are going from one VLAN to another are routed. Some frames are both bridged and routed, as is the case with multicast.

The media modules' features include:

- **Onboard routing**—See "Routing Overview" later in this chapter for more information about the onboard routing capability of the media modules.
- **IEEE 802.3x full-duplex flow control—**This allows the switch ports to send a pause command before input buffers overflow. Half-duplex ports support active backpressure (jamming).
- **VLAN trunking or non-tagged access modes —**This allows the switch ports to interoperate with popular tagged trunking schemes used by other vendors.
- **Priority Queuing and Class of Service—**These features allow you to prioritize traffic between particular stations or sets of stations to support jitter-sensitive applications.

50-Series Media Modules

Table 1-2 lists the 50-series (also called *series 1*), layer 2 and layer 3, Fast Ethernet modules that the Avaya P580 and P882 Multiservice switches support.

*** Note:** HDX stands for Half Duplex, FDX stands for Full Duplex.

Table 1-2. 50-Series Fast Ethernet Modules

Model Number	Description
M5520-100TX	20-Port 10/100BASE-TX module with RJ-45 connectors (layer 2 support)– 10/100, HDX/FDX 100m
M5510-100FX	10-Port 100BASE-FX module with SC connectors (layer 2 support) – 1300 nM, HDX/FDX multimode fiber, 2Km
M5548N-100TC	48-Port, 10/100BASE-TX module with RJ-21 Telco connectors (layer 2 support) – 10/100, HDX/FDX, 100M
M5510R-100FX	10-Port 100BASE-FX module with SC connectors (layer 2 and layer 3 support),– 1300 nM, HDX/FDX, multimode fiber, 2 Km
M5512R-100TX	12-Port 10/100BASE-TX module with RJ-45 connectors (layer 2 and layer 3 support) – 10/100, HDX/FDX, 100m

Table 1-3 lists the 50-series layer 2 and layer 3 Gigabit Ethernet modules that the Avaya P580 and P882 Multiservice switches support. All of the modules have SC connectors.

Table 1-3. 50-Series Gigabit Modules

Model Number	Description
M5502-1000SX-F	2-Port, Full-Duplex1000BASE-SX module (layer 2 support), 550M, 850 nM optics, multimode fiber
M5502-1000LX-F	2-Port, Full-Duplex 1000BASE-LX module (layer 2 support) 1310 nM optics, 5Km, singlemode fiber
M5502-1000SLX-F	2-Port, Full-Duplex 1000BASE-SLX module (layer 2 support) 10 Km, 1310nM optics, singlemode fiber
M5504-1000SX-F	4-Port, Full-Duplex 1000BASE-SX module (layer 2 support) 550m, 850 nM optics, multimode fiber
M5504-1000LX-F	4-port 1000BASE-LX module (layer 2 support), 5Km, 1310nM optics, singlemode fiber
M5504-1000SLX-F	4-port 1000BASE-SLX module (layer 2 support), 10Km, 1310nM optics, singlemode fiber
	1 of 2

Table 1-3. 50-Series Gigabit Modules

Model Number	Description
M5502R-1000SX-F	2-Port, Full-Duplex 1000BASE-SX module (layer 2 and layer 3 support) 550m, 850 nM optics, multimode fiber
M5502R-1000LX-F	2-Port, Full-Duplex 1000BASE-LX module (layer 2 and layer 3 support), 1310 nM optics, 5Km, singlemode fiber
M5502R - 1000SLX-F	2-port 1000BASE-SLX module (layer 2 and layer 3 support), 1310nM optics, 10Km, singlemode fiber
	2 of 2

80-Series Media Modules

Table 1-4 lists the 80-series (also called *series 2*), 10-Gigabit modules that the Avaya P580 and P882 Multiservice switches support.

Table 1-4. 80-Series 10-Gigabit Modules

Model Number	Description
M8001R-10000LR	1-Port L3 10-Gigabit LR module, single mode fiber, 1310 nM, distances up to 10 km.
M8001R-10000ER	1-Port L3 10-Gigabit ER module, single mode fiber, 1550 nM, distances up to 40 km.

Table 1-5 lists the 80-series Gigabit modules that the Avaya P580 and P882 Multiservice switches support.

Table 1-5. 80-Series Gigabit Modules

Model Number	Description
M8004R-1000GB	4-port, Gigabit Interface Card (GBIC) based module. This module can be configured with up to 4 GBIC optical modules. The 4-port GBIC module supports the following GBIC-type connectors: • 1000BASE-SX (850 nM) — 50-micron multimode fiber (MMF) for distances up to 550m — 62.5-micron MMF for distances up to 260m • 1000BASE-LX (1300 nM) — 50-micron MMF for distances up to 550m — 62.5-micron MMF for distances up to 440m — Singlemode fiber for distances up to 10km • 1000BASE-ELX (1550 nM) for singlemode fiber for distances up to 80km **Note:** If you use the 1000BASE-SX GBIC with single mode fiber, you must use mode conditioning patch cords with a maximum distance of 500m. •
M8008R-1000GB	8-port, GBIC based module. This module can be configured with up to 8 GBIC optical modules. The 8-port GBIC module supports the following GBIC-type connectors: • 1000BASE-SX (850 nM) — 50-micron multimode fiber (MMF) for distances up to 550m — 62.5-micron MMF for distances up to 260m • 1000BASE-LX (1300 nM) — 50-micron MMF for distances up to 550m — 62.5-micron MMF for distances up to 440m — Singlemode fiber for distances up to 10km • 1000BASE-ELX (1550 nM) for singlemode fiber for distances up to 80km **Note:** If you use the 1000BASE-SX GBIC with single mode fiber, you must use mode conditioning patch cords with a maximum distance of 500m.
	1 of 2

Table 1-5. 80-Series Gigabit Modules

Model Number	Description
M8004R-1000T	4-Port 1000BASE-T module with RJ-45 connectors (layer 2 and layer 3 support), 100M
M8008R-1000T	8-Port 1000BASE-T module with RJ-45 connectors (layer 2 and layer 3 support), 100M
	2 of 2

Table 1-6 lists the 80-series Fast Ethernet modules that the Avaya P580 and P882 Multiservice switches support.

Table 1-6. 80-Series Fast Ethernet Modules

Model Number	Description
M8024R-100TX	24-Port 10/100BASE-TX module with RJ-45 connectors (layer 2 and layer 3 support) – 10/100, Half Duplex (HDX)/ Full Duplex (FDX)
M8024R-100FX	24-Port 100BASE-FX module with MT-RJ connectors (layer 2 and layer 3 support), multimode fiber, 2Km
M8048R-100TC	48-Port 10/100BASE-TX module with RJ-21 Telco connectors (layer 2 and layer 3 support) - 10/100, HDX/ FDX, 100M
M8048R-100TX	48-Port 10/100 module with RJ-45 connectors (layer 2 and layer 3 support)

ATM Uplink Module

There are four variations of the ATM Uplink module:

- 2-port Synchronous Optical NETwork (SONET)/ Synchronous Digital Hierarchy (SDH) OC-3c/STM-1 Single Mode Fiber (SMF)
- 2-port SONET/SDH OC-3c/STM-1 MMF (MultiMode Fiber)
- 2-port SONET/SDH OC-12c/STM-4c SMF
- 2-port SONET/SDH OC-12c/STM-4c MMF

Each module has two physical interfaces that are capable of being active simultaneously.

The major features of the ATM Uplink module are:

- LANE V2 Client - with LANE 1 compatibility, supports 128 interfaces
- QoS (Quality of Service) support for UBR, nrt-VBR, rt-VBR, and CBR
- ATM UNI (User-to-Network Interface) V 3.0, 3.1 and 4.0 signaling
- RFC 1483 support
- ILMI (Integrated Local Management Interface) Version 4.0
- PNNI (Private Network - to - Network Interface) non-transit
- Link Failover
- Spanning Tree
- Load sharing

*** Note:** All ATM protocols are ATM Forum standard.

See *Avaya P550R/P580/P880/P882 Multiservice Switch ATM Uplink Module User Guide* for more detailed information.

Software Features

The Avaya P580 and P882 Multiservice switches support the following software features:

- Security
- VLANs
- Hunt Groups
- Spanning Tree
- Extensive Fault Tolerance
- Routing Protocols
- Access Lists
- IP Multicast
- 50-Series Buffers and Queues
- 80-Series QoS
- System Management

Security

The P580 and P882 support the following security features:

- Secure Mode—Secure mode restricts management of the switch to the following secure protocols: HTTPS, SSH, and SNMPv3. For information on secure mode, see Chapter 4, "Security."
- SSHv2—SSH is a protocol for secure remote login and other secure services. SSH provides a secure service that is similar to Telnet. For information on SSHv2, see Chapter 4, "Security."
- HTTPS Using SSLv3 or TLSv1—SSL and TLS are protocols that provide data security between application protocols (such as HTTP, Telnet, NNTP, FTP) and TCP/IP. SSL and TLS are used to transmit encrypted data over TCP/IP networks, and authenticate sites and clients with certificates. For information on HTTPS, see Chapter 4, "Security."
- RADIUS Client Support—RADIUS is a service that authenticates users when they attempt to log in to a Network Access Device (NAD) such as an Avaya switch. For information on RADIUS, see Chapter 4, "Security."
- SNMPv3. SNMPv3 provides encryption and timeliness checks of PDUs and authentication of users. For information on SNMPv3, see Chapter 5, "Configuring SNMP."
- Enhanced user account security, including maximum number of login attempts, timeout limit, minimum password length, and expiration. For information on user account security, see Chapter 2, "Setting Up the Switch."

VLANs

VLANs provide network managers with two significant capabilities:

- The ability to segment traffic in a flat switched network. This helps prevent traffic from being forwarded to stations where it is not needed.
- The ability to ignore physical switch locations when creating workgroups. VLANs are logical constructions and can traverse physical switch boundaries.

The hardware on all multiservice switches support port-based VLANs with the following characteristics:

- Frames classified as Layer 1 (port-based) when they enter the switch
- Explicitly tagged VLAN packets — these are forwarded based on the information in the packet.
- Up to 1,000 VLANs — VLANs define a set of ports in a flooding domain. Packets that need to be flooded are sent only to ports participating in that VLAN.

For more information on VLANs, see Chapter 6, "Using VLANs, Hunt Groups, and VTP Snooping."

Hunt Groups

Hunt groups allow you to aggregate bandwidth from multiple ports so they act as one high-bandwidth switch port. Hunt groups create multi-gigabit pipes to transport traffic through the highest traffic areas of your network. You can create hunt groups that interoperate with other vendors' equipment (for example, Cisco's Etherchannel and Sun's Quad Adapter).

For more information on hunt groups, see Chapter 6, "Using VLANs, Hunt Groups, and VTP Snooping."

Spanning Tree

Spanning tree protocol is used to prevent loops from forming in your network. The spanning tree algorithm creates a single path through the network by ensuring that if more than one path exists between two parts of a network, only one of these paths is used, while the others are blocked.

The P580 and P882 support common Spanning Tree and Rapid Spanning Tree. Rapid Spanning Tree provides faster recovery from network failures than common Spanning Tree provides.

All multiservice switches support the following four Spanning Tree configurations:

- IEEE 802.1D Spanning Tree
- Per-VLAN Spanning Tree

- Dual-Layer Spanning Tree
- Global Disable

For more information about Spanning Tree, see Chapter 7, "Configuring Rapid Spanning Tree."

Extensive Fault Tolerance

All multiservice switches are designed to function as backbone switches. You can install each switch in your network's core without creating a single point of failure.

Extensive fault tolerance features include:

- Port redundancy. If the primary port fails, the secondary port becomes the active port. Failover time is less than 5 seconds. For information on port redundancy, see Chapter 8, "Configuring Ports."
- VRRP. Virtual Router Redundancy Protocol (VRRP) is used to provide fast-fail over for hosts if the default gateway fails. For information on VRRP, see Chapter 12, "Configuring IP Routing."
- N+1 power. Power supplies share the power supply load. If one fails, the remaining supplie(s) assume the load automatically and the switch management system warns of the failure.
- Hot-swappable fans and modules—Each fan and module can be changed from the switch front panel, without powering down the switch.
- Redundant switch links (using spanning tree and hunt groups).
- Front-loadable cables, modules, power supplies and fans.
- Redundant switch element and switch controller modules (fault-tolerant switches only).
- Redundant supervisor modules.

For more information hardware redundancy, see Chapter 10, "Configuring Redundancy Options."

Routing Protocols

The P580 and P882 support the following routing protocols:

- IRDP. For information on IRDP, see Chapter 12, "Configuring IP Routing."
- RIP. For information on RIP, see Chapter 14, "Configuring RIP Routing."
- OSPF. For information on OSPF, see Chapter 15, "Configuring the OSPF Routing Protocol."
- IPX. For information on IPX, see Chapter 16, "Configuring IPX Routing."
- IPX RIP. For information on IPX RIP, see Chapter 17, "Configuring the IPX RIP Protocol."
- IPX SAP. For information on IPX SAP, see Chapter 18, "Configuring the IPX SAP Protocol."
- AppleTalk. For information on AppleTalk, see Chapter 19, "Configuring AppleTalk Routing."

Access Lists

An access list is a group of access control rules. Access control rules describe *how* to forward (route) packets, as opposed to where to forward them. The *how* can be to forward the packet with a specific priority(0-7), forward the traffic with an unchanged priority, or filter packets (drop). For information on access lists, see Chapter 13, "Configuring Access Lists."

IP Multicast

IP multicast enables a single host to distribute information to multiple recipients. The P580 and P882 support the following IP multicast protocols:

- IGMP
- DVMRP

  For information on IGMP and DVMRP, see Chapter 12, "Configuring IP Routing."
- Intelligent Multicast—LGMP, CGMP Snooping, and IGMP Snooping. For information on intelligent multicast, see Chapter 20, "Managing Intelligent Multicasting."

50-Series Buffers and Queues

Buffer and queue management relieves congestion in a network. Adding gigabit speeds to existing networks means that there can be a huge disparity between link speeds. For example, anything more than a 1-percent load on a gigabit link could easily overwhelm a 10 Mb/s Ethernet link.

Each switch employs the following buffer and queue management techniques:

- Configurable active backpressure:
  - Half-duplex ports use active backpressure to jam input ports when their frame buffers are full.
  - Full-duplex links use IEEE 802.3z pause control frames to pause traffic when buffers are full.
- Packed frame buffers for optimal memory utilization. The memory management allows virtually 100% utilization of buffer memory.
- Two CoS priority queues for 50-Series modules and eight CoS priority queues for 80-Series modules provide flexible queue management algorithms to meet application requirements.
- Configurable queue depth for each prioritized packet queues.
- Configurable priority threshold.
- Configurable service ratio tunes queue priority.

For more information on 50-series buffers and queues, see Chapter 24, “Managing Buffers and Queues on 50-Series Modules.”

80-Series QoS

Quality of Service (QoS) is a set of tools that make it possible for you to manage traffic across a switch or a network. These tools protect specific traffic from the effects of network congestion. You can enable these features for a module or individual ports on a module to control the flow of traffic across your network. This control makes it possible for you to guarantee that delay-sensitive traffic such as voice over IP (VoIP) receives the priority it requires, while also ensuring that the switch services other low priority data.

QoS makes it possible for you to prioritize time-sensitive traffic and assign larger amounts of bandwidth to those applications that require it.

For more information on QoS, see Chapter 25, “80-Series QoS.”

System Management

The P580 and P882 support the following system management features:

- LDAP. For information on LDAP, see Chapter 12, "Configuring IP Routing."
- RMON. For information on RMON, see Chapter 23, "Using RMON and Ethernet Statistics to Analyze Network Performance."
- SNMP. For information on SNMP, see Chapter 5, "Configuring SNMP."
- Port mirroring. For information on port mirroring, see Chapter 23, "Using RMON and Ethernet Statistics to Analyze Network Performance."
- Syslog event reporting. For information on syslog, see Chapter 21, "Monitoring the Avaya Multiservice Switch."

Routing Overview

All P580 and P882 Multiservice switches can be configured as an IP, IPX, and AppleTalk router with virtual interfaces. Virtual interfaces are mapped to physical ports or VLANs. Layer 3 IP traffic is routed between the virtual interfaces.

Ports become members of VLANs by being assigned or by rules. Multiple VLANs can share a single trunk port. In contrast, multiple physical ports can be associated with a single VLAN. In all cases, traffic that arrives and leaves the same VLAN is bridged, not routed.

This section provides the following sections:

- Compatibility with Layer 2 Modules
- Routing with Layer 2 and Layer 3 Modules

Compatibility with Layer 2 Modules

Each switch is completely backward compatible with all of the layer 2 media modules that the switch currently supports. Layer 3 traffic is routed by sending that traffic to the supervisor module. The supervisor module routes all traffic from layer 2 media modules as described in the next section, "Routing with Layer 2 and Layer 3 Modules."

*** Note:** Layer 2 traffic that does not require routing is bridged independently of the layer 3 traffic based on the MAC address or VLAN information.

Routing with Layer 2 and Layer 3 Modules

When a switch contains a combination of layer 2 and layer 3 modules, IP and IPX routing is performed by the layer 3 media modules or the supervisor module using special ASICs present on those modules. These ASICs contain an *address cache* (forwarding table) that can contain a maximum of 20,000 cache entries. The address cache entries consist of packet addressing information and next hop information that enable the switch to effectively route the packets to their destination.

The supervisor module also maintains a *master routing table* that is kept in the its memory. The master routing table can contain up to 28,000 entries. This routing table enables the supervisor module to keep track of which entries are in each address cache. As a result, each time a change occurs in the master routing table, the supervisor module updates the appropriate address caches. For example, if a unicast route is removed from the master routing table, all matching entries in address caches are also removed.

Consequently, when you connect a switch to the network, it begins to receive frames from the network and builds a master routing table (supervisor module) and forwarding tables (address caches of media modules) based on those frames.

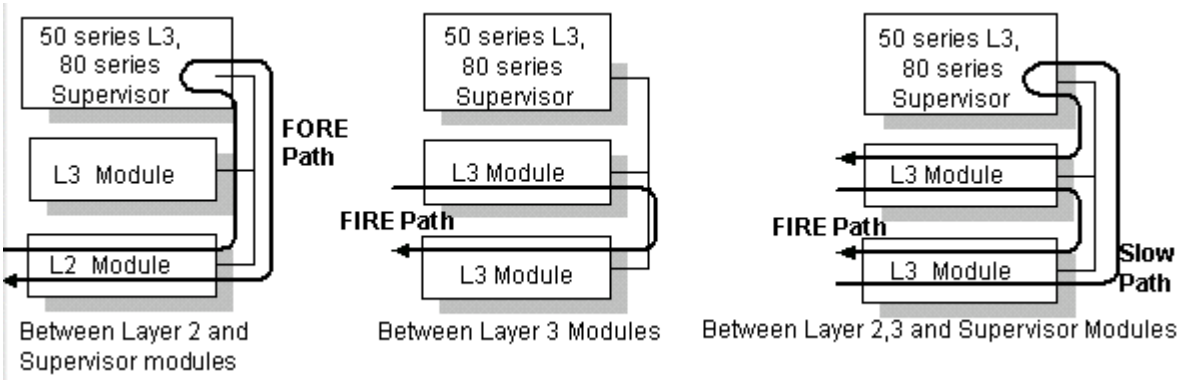
This process creates three distinct results:

- All known (learned) layer 3 traffic that requires routing, is routed directly by the 50-series layer 3 or 80-series media module without a need to traverse the switching fabric to get to the supervisor module's software routing function. This is referred to as the **FIRE** path, or Fast Inband Routing Engine, since the routing is accomplished in the hardware in the media modules.
- All unknown (not learned) layer 3 traffic must first be sent to the 50-series layer 3 or 80-series supervisor module, where information on the frame is added to the supervisor module's master routing table, then it is added to the address cache of each applicable layer 3 media module. This is referred to as the **Slow Path**, since the routing is accomplished by the CPU and software in either supervisor module instead of the hardware in the media modules.The **Slow Path** is used when the destination is unknown for IP and IPX packets and for all AppleTalk routing.

- Since layer 2 modules have no routing capability, packets that are received by a layer 2 module and require routing are forwarded by sending the packet to the 50-series layer 3 or 80-series supervisor module. The routing engine on the supervisor module then performs the routing operation for the layer 2 modules and sends the packet back through the switching fabric to the destination port.This is referred to as the **FORE** path, or Fast Out of Bands Routing Engine, since the routing is accomplished in the hardware of either supervisor module.

Figure 1-3 shows how traffic is routed in a switch.

Figure 1-3. Layer 2 and Layer 3 Routing in a Switch



2 Setting Up the Switch

Overview

The following information and procedures that are provided in this chapter are common to both layer 2 and layer 3 module:

- Setting Up Your Laptop or PC
- Performing Initial Setup of the Switch
- Using Telnet to Manage the Switch
- Using the Web Agent to Configure the Switch
- Configuring Custom Access Types
- Configuring User Accounts
- Changing the Console Serial Port Settings
- Configuring Dial-Up Networking
- Changing the TCP Ports for HTTP and Telnet
- Managing Configuration Files

*** Note:** The last step in each procedure tells you to Click **APPLY** to save the setup or changes that you made. This step saves the setup or any changes to the running configuration only. The startup configuration is not changed. Therefore, these and other changes will be lost if the switch goes down or if you turn it off.

To save any changes to the startup configuration, you must copy the running configuration to the startup configuration. For information on how to perform this task, see "Copying Configuration Files" later in this chapter.

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.

Setting Up Your Laptop or PC

You need a laptop or PC to make a serial line connection to the switch to complete initial switch setup. Table 2-1 lists the settings required for the laptop or PC to communicate with the switch.

Table 2-1. Laptop or PC Settings

Baud Rate	Stop Bits	Data Bits	Flow Control	Parity
9,600	1	8	Xon/Xoff	None

Performing Initial Setup of the Switch

Before you can use Telnet or the Web Agent to configure the switch, you must connect a laptop or PC to the console port and assign the switch an IP address and subnet mask.

Before Moving a Supervisor Module

Before moving a supervisor module (either active or standby) to another chassis, you must initialize NVRAM on the module (**nvram initialize** CLI command).

If you do not initialize NVRAM before moving a supervisor module to another chassis and assigning it a new IP address, the supervisor retains its console IP address, console MAC address, and chassis information. This problem may cause:

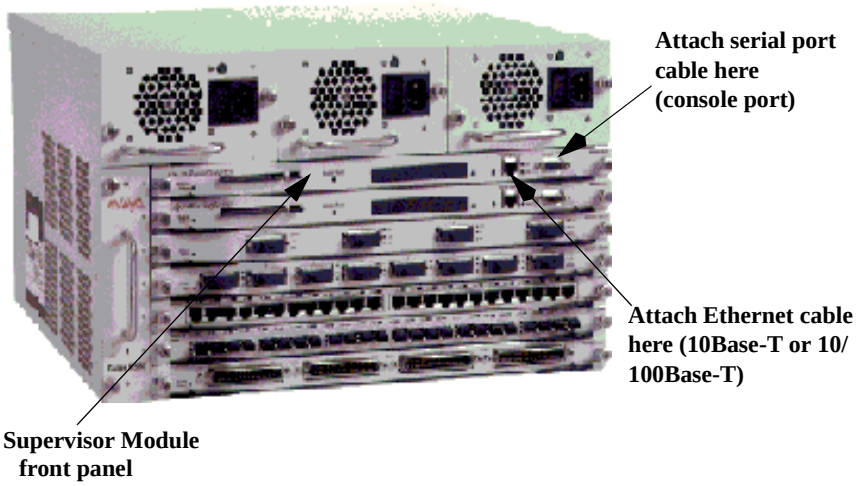
- Duplicate IP addresses
- Inaccurate chassis information that can cause the switch to not operate correctly.

Assigning an IP Address

To perform initial setup of the switch:

1. Attach a 9-pin straight-through male-to-female serial cable from the serial port on your laptop or PC to the serial port on the supervisor module's front panel (Figure 2-1).

Figure 2-1. Avaya P580 Multiservice Switch Front View



2. Run a terminal emulation program (HyperTerminal, for example) on the attached laptop or PC. Verify that the laptop or PC monitor settings match those listed in Table 2-1.

3. Turn on the switch. The switch displays a series of startup messages in the terminal emulation program:

4. At the `Login:` prompt, enter **root**. The password prompt displays.

*** Note:** Login information and password prompts are case sensitive.

5. At the `Password` prompt, enter **root**. The command line interface prompt displays. You must now change the command mode to the Global Configuration mode so that you can use the **setup** command.

6. Perform the following steps to change the command mode to the Global Configuration mode:

   a. Enter the **enable** command.

   b. Enter the **configure** command.

7. Enter the **setup** command. This command initiates a series of queries. Answer each query as follows:

   a. Press **Enter** when prompted to change the super user password. This accepts the default of **Yes**.

   b. Enter your **old password**. The system then prompts you for a new password.

   c. Enter your **new password**, then re-enter the new password to verify your choice.

*** Note:** Avaya strongly recommends that you change the default password.

d. Enter the **IP address** for the switch manager's Ethernet console.

e. Enter the **subnet mask**.

f. Enter the **default gateway** for the switch.

The following is a sample output of the **Setup** command to change or modify the Supervisor's Ethernet Console's IP address, subnet mask, default gateway or changing the super user password.

```
Welcome to Switch Setup. The brief series of
questions that follows will help you to configure
this switch. After completing this process, you
will be able to manage the switch using:

 - the switch-based HTTP server

 - the Element Management System.

Text in [] is the default answer for each
questions.

To accept the default, press ENTER.

Would you like to change the super user password
[Yes]? Y

Old Password: xxxx

New Password: xxxx

Re-type New Password: xxxx

User password changed successfully

What do you want the switch manager's console

Ethernet IP Address to be [0.0.0.0]? 10.0.0.1

What is the subnet mask for your network's

IP address [0.0.0.0]? 255.255.255.0

What is the IP address of the default gateway for
this network segment [0.0.0.0]?
```

You can now connect to the switch using the front-panel out-of-band 10Base-T connection. This allows you to log in using either the embedded Web Agent or the EMS.

See the Installation and Operation guides for instruction on establishing additional IP network connections.

Connecting to the Supervisors 10/100Base-T port

Connect one end of a crossover patch cable to the 10/100Base-T on the Supervisor module front panel (Figure 2-1). Connect the other end to an out of band device.

See Table 2-2 for an explanation of the pinouts for the 10Base-t crossover patch cable.

Table 2-2. Pinouts for 10Base-T Crossover Patch Cables

Pin #	Color	Pin #	Color
1	WO	3	WG
2	O	6	G
3	WG	1	WO
4	B	4	B
5	WB	5	WB
6	G	2	O
7	WBr	7	WBr
8	Br	8	Br

Using Telnet to Manage the Switch

You can manage the Avaya Multiservice switch several ways. In addition to managing the switch by means of the Web Agent, SNMP and serial console command line interface (CLI) commands, the switch provides Telnet server capability. Through this Telnet session, an administrator can manage the switch by means of CLI commands.

This section contains the following information:

- Setting up a Management Interface
- Establishing a Telnet Session
- Terminating a Telnet session

- Viewing Active Telnet Sessions
- Telnet Session Expiration Timer

As a more secure alternative to Telnet, you can use SSHv2. For information on using an SSH connection to manage the switch, see "SSHv2" in Chapter 4, "Security."

Setting up a Management Interface

Before you can establish a Telnet session you must configure an IP interface to allow management of the switch. Using the IP Interfaces Web page or a CLI command, create an IP interface and set the **IP Routing** option to either **Routing/Mgmt** or **Mgmt Only**. For more information on setting up an IP interface on the switch, see "Creating and Assigning IP Interfaces to the VLAN" in Chapter 12, "Configuring IP Routing."

Establishing a Telnet Session

Once you have set up an IP interface and configured it to allow management traffic, you can open a Telnet session to the switch. To open a Telnet session, you must use one of the commercially available Telnet clients. Launch the Telnet client application and open a Telnet session with the IP address of the management interface. The switch will display its login prompt.

If the TCP port for Telnet requests is set to a port number other than 23, you must specify the TCP port number in addition to the IP address or host name to start a Telnet session. For example, if you change the TCP port from 23 to 9998 on switch 192.168.0.126, enter **telnet 192.168.0.126 9998** to start a Telnet session. For information on how to change the TCP port for Telnet requests, see "Changing the TCP Ports for HTTP and Telnet" later in this chapter.

Once logged into the switch through a Telnet session, you can manage the switch using CLI commands. For more information on CLI commands, see the *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.

Supported Number of Telnet Sessions

The Avaya Multiservice Switch supports a maximum of 6 concurrent Telnet sessions. This number makes it possible for multiple network administrators to connect to and manage the switch. However, if multiple Telnet sessions are open, configuration changes from one session could be overwritten by configuration changes from another.

*** Note:** Avaya recommends that you use one Telnet session at a time to manage the switch.

Although it is possible to establish nested Telnet sessions, Avaya recommends that you establish a separate Telnet session for each switch that you want to manage. A nested Telnet session occurs when you establish a Telnet session from a client to one switch, then through that session, open another Telnet session to a second switch. Having a separate Telnet session for each switch that you want to manage helps to avoid confusion.

Terminating a Telnet session

To terminate the Telnet session gracefully, exit the management session by repeatedly typing **exit** until the switch displays the login prompt. To terminate the Telnet session from the client, enter the client-specific command to terminate the session.

Troubleshooting Tip

When you enter CLI commands in a Telnet session, memory is allocated to hold the CLI command history. Closing the Telnet session abruptly by terminating the Telnet client application does not free this memory and may cause the switch to reset. Close the Telnet session gracefully as explained in the previous paragraph.

Viewing Active Telnet Sessions

To view active Telnet sessions on the switch enter the **show sessions** command as shown in the example below. A list of active telnet sessions is displayed.

```
Welcome to the Avaya *Enhanced* CLI
Avaya> show sessions
Session ID    Line ID   Location
*0            9vty      123.123.123.123:1211
```

Telnet Session Expiration Timer

Once a telnet session is established, the session will remain open while there is activity. However, if the session is left idle, the session expires and is automatically terminated based on the setting of the Telnet inactivity timer. The default setting for this timer is 900 seconds or 15 minutes. Setting this timer to 0 deactivates the inactivity timer so the Telnet session will not expire.

To set the Telnet inactivity timer, enter Global Configuration mode and then the following CLI command:

```
(configure)# ip telnet inactivity-period <time-out in seconds>
```

Table 2-3 lists the minimum, maximum, and default values for the Telnet inactivity timer.

Table 2-3. Telnet Inactivity Timer

Minimum value:	0 seconds	Disables inactivity timer
Default value:	900 seconds	Expiration time is 15 minutes.
Maximum value:	65,536 seconds	Expiration time is approx. 18 hrs.

CAUTION: Avaya recommends that you not disable the inactivity timer. If the inactivity timer is disabled and the maximum number of Telnet sessions is open, none of them will expire. Subsequent attempts to establish a Telnet session will fail. Additionally, idle Telnet sessions may represent security risks.

Using the Web Agent to Configure the Switch

The switch provides an embedded HTTP server that allows you to set switch parameters by using the Web Agent. You can use this interface for quick and simple configuration changes.

As a more secure alternative to HTTP, you can use HTTPS. For information on using an HTTPS connection to manage the switch, see "HTTPS Using SSLv3 or TLSv1" in Chapter 4, "Security."

Although this manual provides detailed procedures for using the Web Agent to configure the switch, the CLI commands to perform the same task are listed after each procedure. See *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1* for details about these commands.

This section contains the following procedures:

- Opening the Web Agent
- Logging Out of the Web Agent

Opening the Web Agent

Although the Web Agent is supported by any frames-capable browser, the system has been specifically qualified only with the following browsers:

- Microsoft Internet Explorer 5.5 or later. To download Microsoft Internet Explorer, see http://www.microsoft.com.
- Netscape Navigator 7.0 or later. To download Netscape Navigator, see http://www.netscape.com.

To open to the Web Agent:

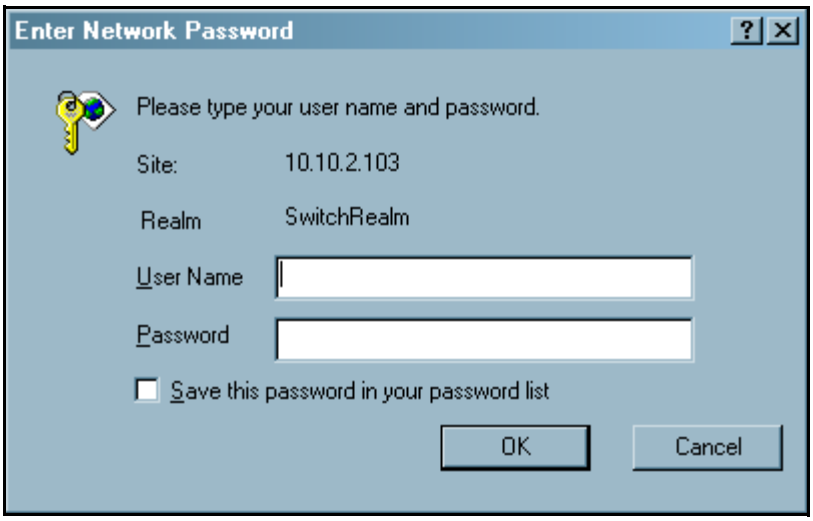
1. Start your browser.
2. Enter the IP address for the switch you want to manage in the Address bar. For example: http://10.91.7.23.

   If the TCP port for HTTP requests is set to a port number other than 80, you must specify the TCP port number in addition to the IP address to open the Web Agent. For example, if you change the TCP port from 80 to 9999 on switch 192.168.0.126, enter **http://192.168.0.126:9999** to open the Web Agent. For information about how to change the TCP port for HTTP requests, see "Changing the TCP Ports for HTTP and Telnet" later in this chapter.

   ***Note:** You must enter a separate IP address for each supervisor module interface (console or inband). For layer 3, you can use can any of the router interfaces.

3. Press **Enter**. The Welcome to the Switch Management Web page is displayed.
4. Click **Login**. The **Enter Network Password** dialog box is displayed. See Figure 2-2.

Figure 2-2. Enter Network Password Dialog Box



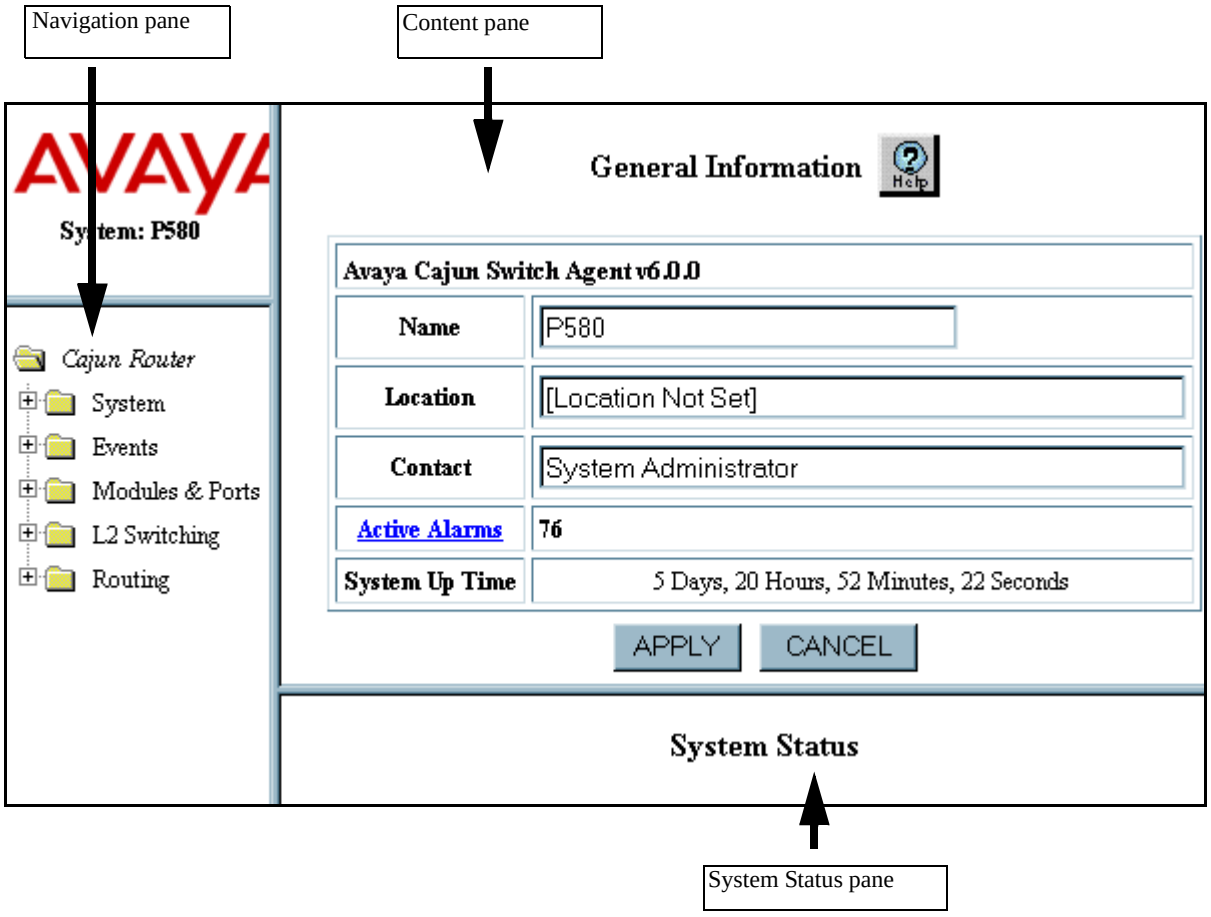
5. In the **User Name** field, enter a valid user name. The default super user name is **root**.

6. In the **Password** field, enter a valid password. The default password is **root**.

7. Click **OK**. The General Information Web page is displayed in the content pane. See Figure 2-3.

***Note:** To optimize security, change the root password for the system as soon as possible.

Figure 2-3. General Information Web Page



Logging Out of the Web Agent

To exit the Web Agent securely, you must log out of the Web Agent and close all open browser windows.

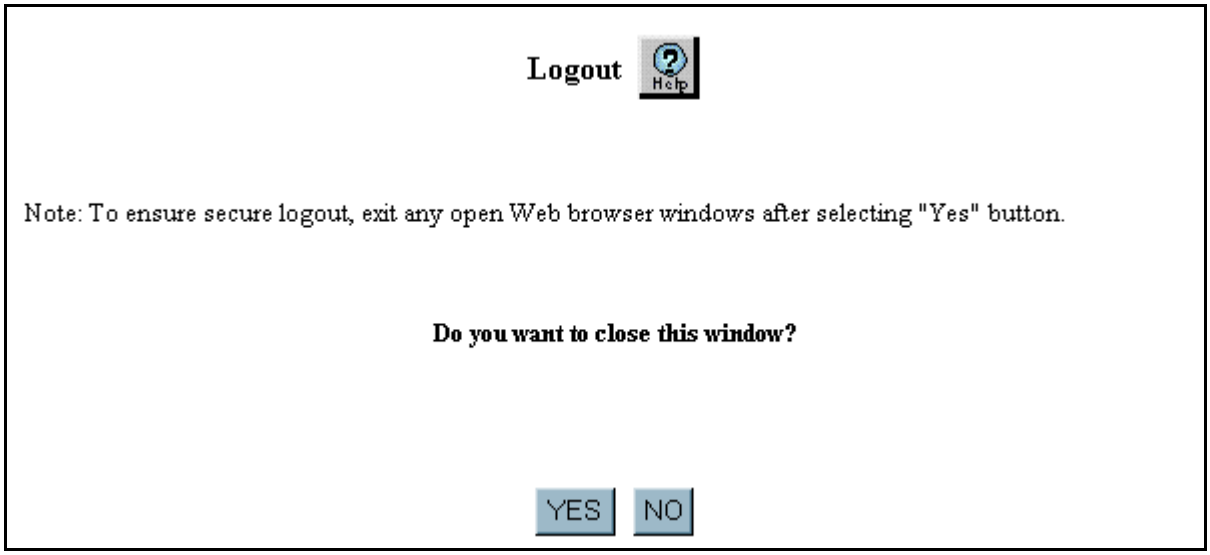
If you do not log out and close all open browser windows, it is possible that another user could access the Web Agent and change the switch configuration.

***Important:** For the logout feature to work correctly in Microsoft Internet Explorer, it must be set to run ActiveX controls.

To log out of the Web Agent:

1. In the navigation pane, expand the **System** folder.
2. Click **Logout**. The Logout Web page is displayed in the content pane. See Figure 2-4.

Figure 2-4. Logout Web Page



3. Click **YES** to log out.
4. Close all other open Web browser windows.

Configuring Custom Access Types

Overview

Custom access types make it possible for you to create a user profile that has specific functional privileges. The functional privileges that you can assign a user profile are:

- System configuration
- Module and port management
- Events management
- L2 switching
- Routing

After creating one or more custom access types (user profiles), you can create a user account and assign a custom access type to that account. The user is then able to manage the switch based on the custom access type that was assigned to his or her user account.

The switch supports a maximum of 30 custom access types.

This section contains procedures for the following tasks:

- Creating a Custom Access Type
- Modifying a Custom Access Type
- Deleting a Custom Access Type

Creating a Custom Access Type

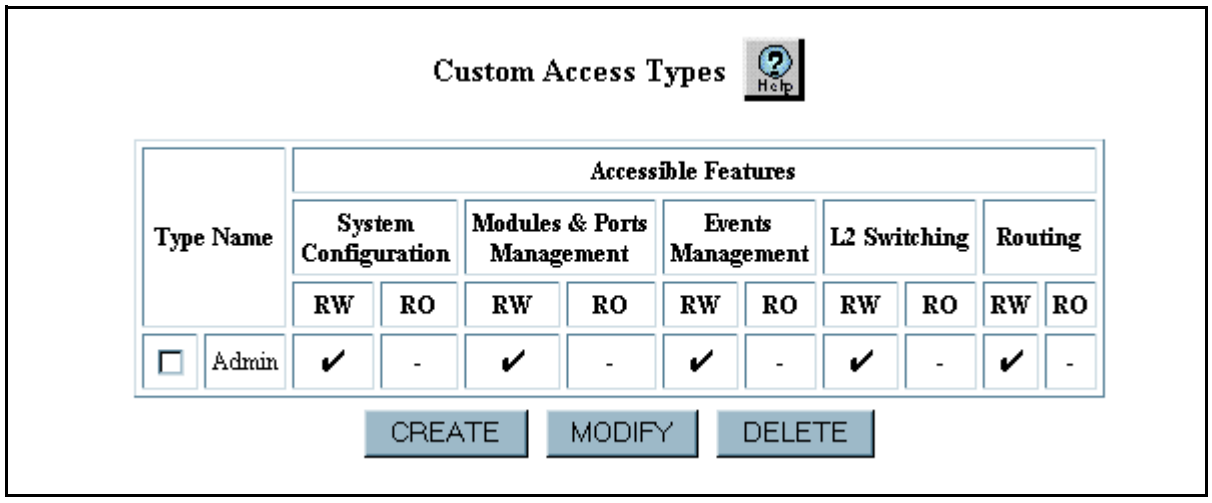
Web Agent Procedure

To create a custom access type:

1. In the navigation pane, expand the **System** > **Administration** folders, and then click **Custom Access Types**.

   The Custom Access Types Web page is displayed in the content pane. See Figure 2-5.

Figure 2-5. Custom Access Types Web Page



2. Click **Create**.

   The Create Custom Access Type Web page is displayed in the content pane. See Figure 2-6.

Figure 2-6. Create Custom Access Type Web Page



3. In the **Type Name** field, enter a name for the custom access type. You can enter up to 31 characters. Do not use spaces.

4. For each feature in the Accessible Features list, select the permission that you want to assign the custom access type:

   — **Read-Write** allows users to view and modify settings for the feature.

   — **Read-Only** allows users only to view settings for the feature.

   — **None** allows users to neither view or modify settings for the feature.

5. Click **Apply** to save the custom access type or **Cancel** to clear your entries.

CLI Command

To create a custom access type, use the following CLI command:

`(configure)#` **custom-access-type <*catName*> [sys-configuration [ro]] [module-port-mgmt [ro]] [events-mgmt [ro]] [l2-switching [ro]] [routing [ro]]**

Modifying a Custom Access Type

Web Agent Procedure

To modify a custom access type:

1. In the navigation pane, expand the **System** > **Administration** folders, and then click **Custom Access Types**.

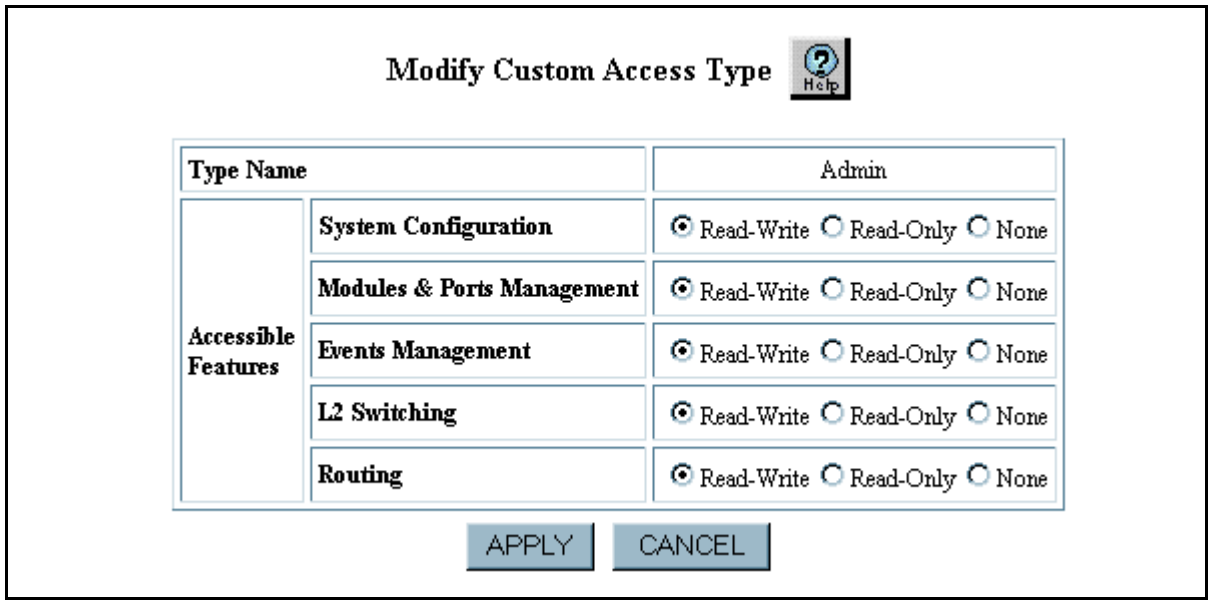
   The Custom Access Types Web page is displayed in the content pane. See Figure 2-5.

2. In the **Type Name** field, select the custom access type that you want to modify.

3. Click **Modify**.

   The Modify Custom Access Type Web page is displayed in the content pane. See Figure 2-7.

Figure 2-7. Modify Custom Access Type Web Page



4. In the **Accessible Features** list, modify the permissions as necessary:

   — **Read-Write** allows users to view and modify settings for the feature.

   — **Read-Only** allows users only to view settings for the feature.

   — **None** allows users to neither view or modify settings for the feature.

5. Click **Apply** to save the custom access type or **Cancel** to clear your entries.

CLI Command

To modify a custom access type, use the following CLI command:

`(configure)#` **set custom-access-type <*catName*> [sys-configuration {ro | rw | none}] [module-port-mgmt {ro | rw | none}] [events-mgmt {ro | rw | none}] [l2-switching {ro | rw | none}] [routing {ro | rw | none}]**

*** Note:** Unlike the **custom-access-type** command that you use to *create* a custom access type, you must specify read-only, read-write, or no permission when you use the **set custom-access-type** command to *modify* a custom access type.

Deleting a Custom Access Type

Web Agent Procedure

To delete a custom access type:

1. In the navigation pane, expand the **System** > **Administration** folders, and then click **Custom Access Types**.

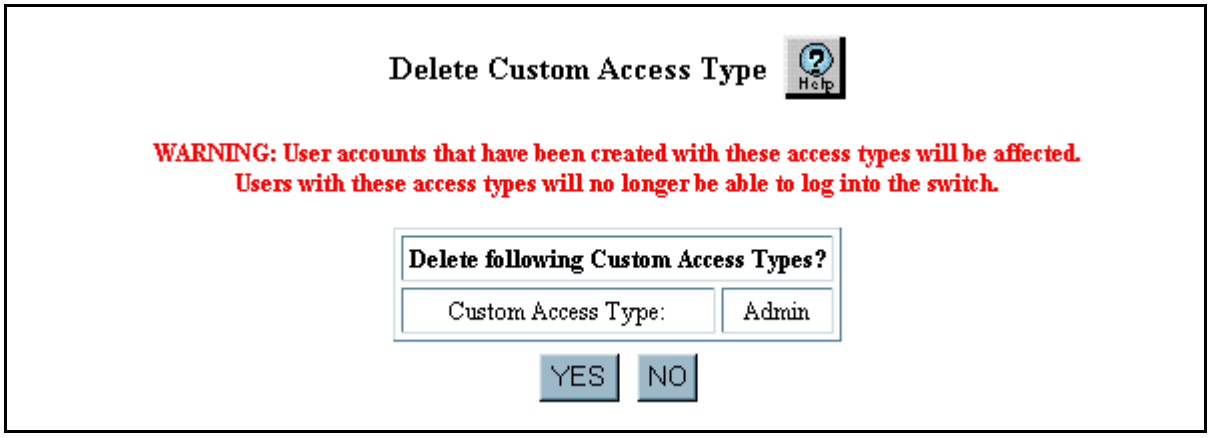
   The Custom Access Types Web page is displayed in the content pane. See Figure 2-5.

2. In the **Type Name** field, select the access type that you want to delete.

3. Click **Delete**.

   The Delete Custom Access Type Web page is displayed in the content pane. See Figure 2-8.

Figure 2-8. Delete Custom Access Type Web Page



4. Click **Yes** to delete the custom access type that is displayed, or click **No** to keep the access type.

   ***Note:** User accounts that are assigned this custom access type will no longer be able to log into the switch.

CLI Command

To delete a custom access type, use the following CLI command:

`(configure)#` **no custom-access-type** <***cat name***>

Configuring User Accounts

This section contains procedures for the following tasks:

- Configuring User Account Security
- Creating a User Account
- Modifying a User Account
- Changing Your Password
- Disabling a User Account
- Deleting a User Account

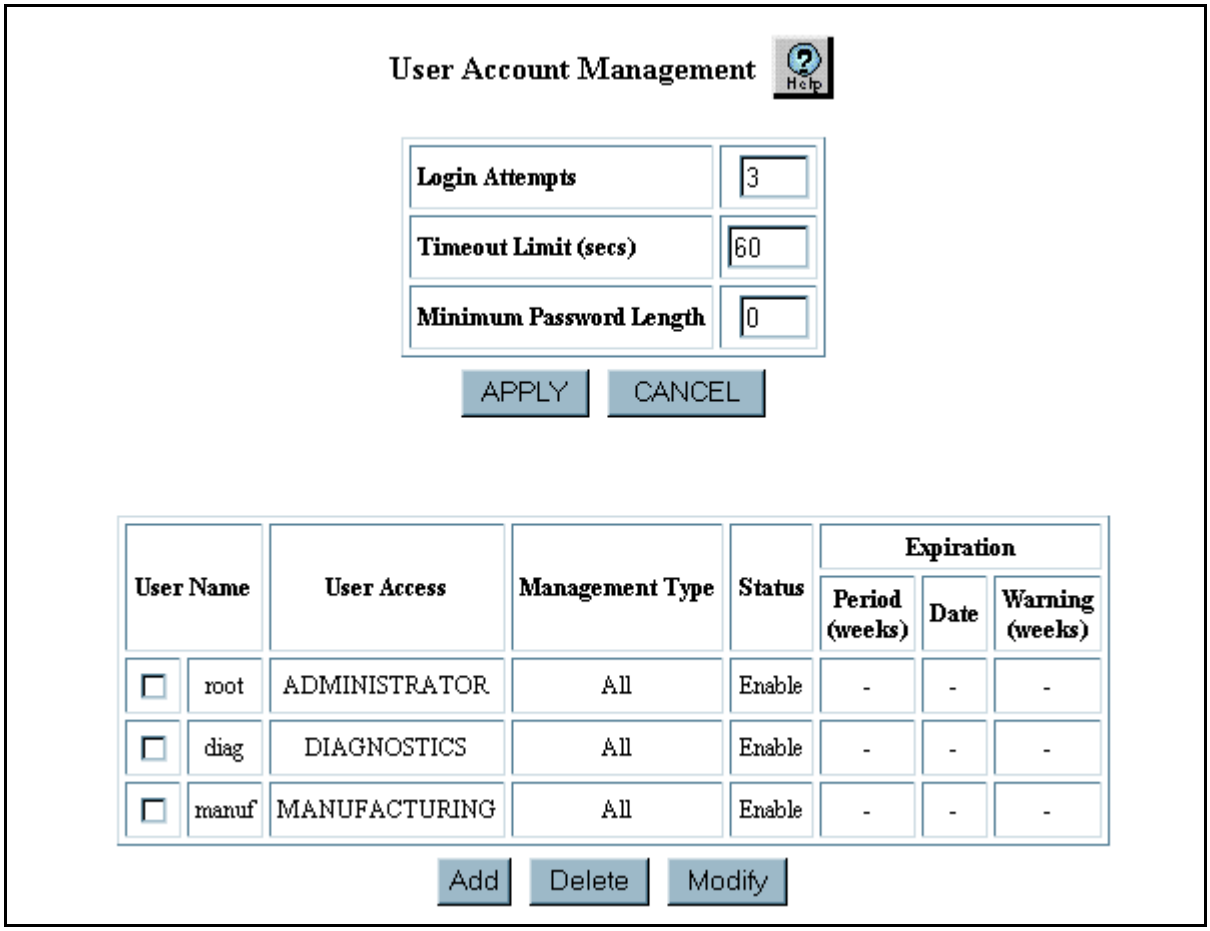
You must have administrator access to configure user accounts.

Configuring User Account Security

Web Agent Procedure

1. In the navigation pane, expand the **System** > **Administration** folders.
2. Click **User Accounts**. The User Account Management Web page is displayed in the content pane. See Figure 2-9.

Figure 2-9. User Account Management Web Page



3. In the **Login Attempts** field, enter the number of login attempts that you want to allow users. When a user exceeds the limit for login attempts, his or her user account is disabled and the switch displays an error message.

   Valid values for this field are 3 to 99 login attempts.

4. In the **Timeout Limit (secs)** field, enter the number of seconds that you want a user account disabled when the limit for login attempts is exceeded. Once the timeout limit expires, the user can attempt to login again.

5. In the **Minimum Password Length** field, enter the minimum number of characters that you want to allow for user passwords. If a user attempts to create a password with fewer characters, the switch displays the following error message:

   Password too short – must be at least <*x*> characters.

6. Click **APPLY**.

CLI Commands

To configure user account security, use the following CLI command:

`(configure)#` **set login [attempts <*num-login-attempts*>] [timeout-limit <*timeout-limit*>] [min-password-length <*min-password-length*>]**

To view the current settings for user account security, use the following CLI command:

show login

Creating a User Account

You can create up to 27 user accounts.

Only users that have administrator access can perform the following tasks:

- View the running configuration and startup configuration.
- Copy text files to or from a TFTP server.
- Download BOOT, APP1, and APP2 images to the switch.
- Update the software on a module.
- Access the Load MIB.

*** Note:** To log in to the ATM Uplink module, a user must:

- Have read-write access to the **Modules and Ports Management** features.
- Enter Privileged mode or Global Configuration mode to enter the **session <*slot number*>** command.

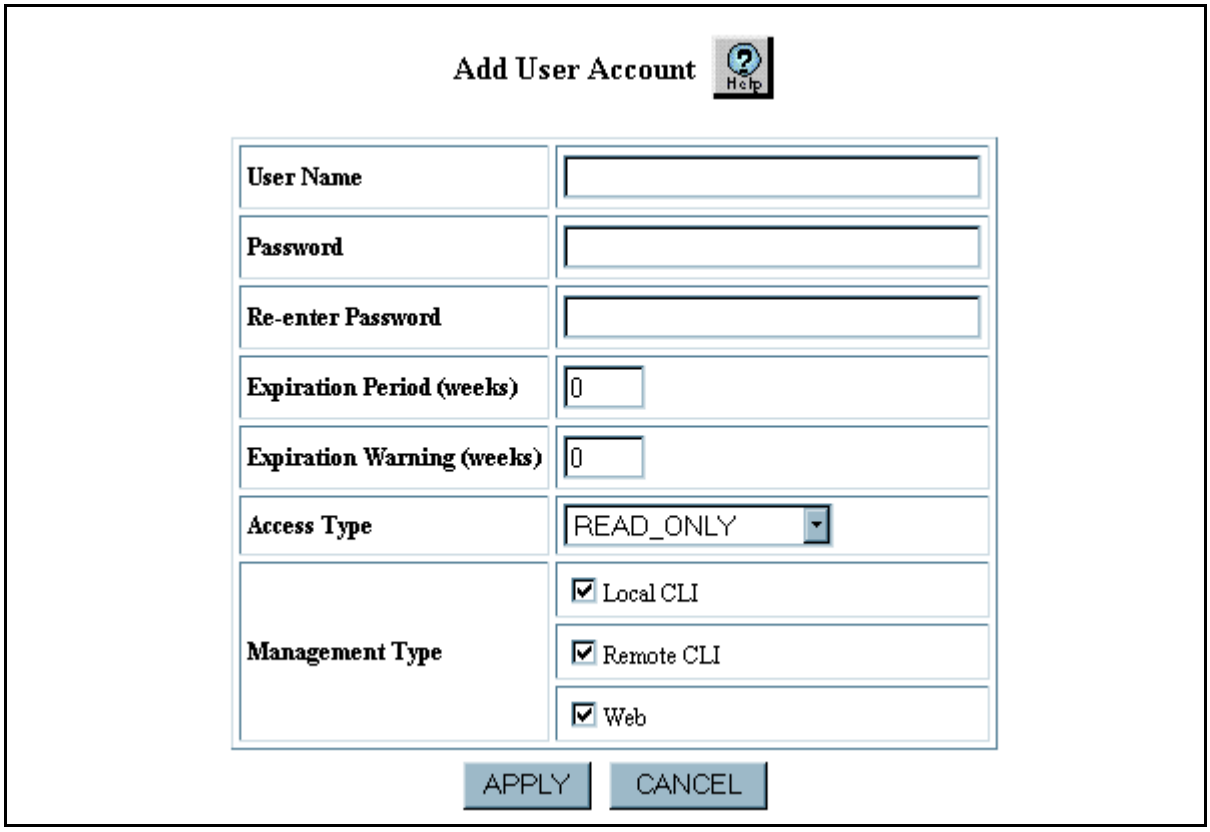
In earlier versions of application software, you could log in to the ATM Uplink module from User mode. This change in functionality prevents unauthorized users from logging in to the ATM Uplink module.

Web Agent Procedure

To create a user account:

1. In the navigation pane, expand the **System** > **Administration** folders.
2. Click **User Accounts**. The User Account Management Web page is displayed in the content pane. See Figure 2-9.
3. Click **Add**. The Add User Account Web page is displayed in the content pane. See Figure 2-10.

Figure 2-10. Add User Account Web Page



4. In the **User Name** field, enter a unique user name. User names can consist of a maximum of 31 characters.

5. In the **Password** field, enter a password for the user name. Passwords can consist of a maximum of 31 characters.

   ***Note:** Do not use a combination of the following special characters for the password ;, ?, \,(,),#, $,%, ^, &, or *.

6. In the **Re-enter Password** field, reenter the password that you entered in the **Password** field.

7. In the **Expiration Period (weeks)** field, enter the number of weeks for which the user account is valid. The expiration period can range from 3 to 999 weeks. The default setting is 0, no expiration.

   When a user account expires, you must reset the account. For information on resetting a user account, see "Modifying a User Account" later in this chapter.

8. In the **Expiration Warning (weeks)** field, enter the number of weeks before user account expiration that you want the user warned. The expiration warning can range from 0 to the expiration period. A setting of 0 indicates that no warning is generated.

9. In the **Access Type** field, select an access type for the user. The options are:

   — READ_ONLY

   — READ_WRITE

   — ADMINISTRATOR

   — Custom access types—All custom access types that are configured on the switch are displayed. For information on custom access types, see "Configuring Custom Access Types" earlier in this chapter.

10. In the **Management Type** field, select the interfaces from which the user can manage the switch. The options are:

    — Local CLI. User can manage the switch from the serial console on the supervisor module.

    — Remote CLI. User can manage the switch by means of a Telnet or SSH session.

    — Web. User can manage the switch by means of the Web Agent.

11. Click **APPLY**.

CLI Command

To create a user account, use the following CLI commands:

- To create a user account:

  `(configure)#` **username <*name*> password [encrypted-type1] <*passwd*> [access-type {read-only | read-write | admin | <*catName*>}] [mgmt-type [all] [local-cli] [remote-cli] [web]]**

- To set an expiration period and expiration warning:

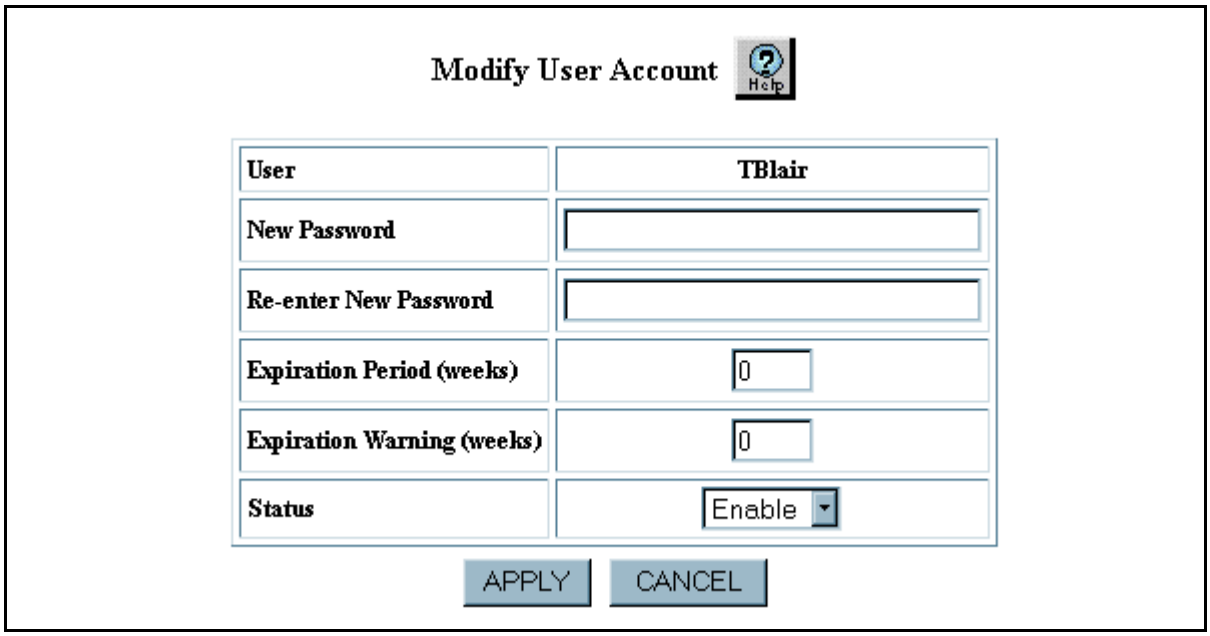
  `(configure)#` **username <*name*> [exp-period <*exp-period*>] [exp-warning <*exp-warning*>]**

Modifying a User Account

To modify a user account:

1. In the navigation pane, expand the **System** > **Administration** folders.
2. Click **User Accounts**. The User Account Management Web page is displayed in the content pane. See Figure 2-9.
3. In the **User Name** field, select the user account that you want to modify.
4. Click **Modify**. The Modify User Account Web page is displayed in the content pane. See Figure 2-11.

Figure 2-11. Modify User Account Web Page



5. To change the user password:

   a. In the **New Password** field, enter a new password for the user name. Passwords can consist of a maximum of 31 characters.

   ***Note:** Do not use a combination of the following special characters for the password ;, ?, \,(,),#, $,%, ^, &, or *.

   b. In the **Re-enter New Password** field, reenter the password that you entered in the **New Password** field.

6. To change the number of weeks for which the user account is valid, change the setting of the **Expiration Period (weeks)** field. The expiration period can range from 3 to 999 weeks. The default setting is 0, no expiration.

7. To change the number of weeks before user account expiration that the user is warned, change the setting of the **Expiration Warning (weeks)** field. The expiration warning can range from 0 to the expiration period. A setting of 0 indicates that no warning is generated.

8. In the **Status** field, select **Enable** or **Disable** to change the status of the user account.

   When a user account expires, you reset the account by changing the setting of this field from **Disable** to **Enable**.

9. Click **APPLY**.

CLI Commands

To modify a user account, use the following CLI commands:

- To modify a user account, `(configure)#` **username <*name*> password [encrypted-type1] <*passwd*> [access-type {read-only | read-write | admin | <*catName*>}] [mgmt-type [all] [local-cli] [remote-cli] [web]]**
- To modify the expiration period or expiration warning for a user account, `(configure)#` **username <*name*> [exp-period <*exp-period*>] [exp-warning <*exp-warning*>]**
- To modify the status of a user account, `(configure)#` **username <*name*> status {enable | disable}**
- To view the settings for a user account, `>` **show username [<*name*>]**

Changing Your Password

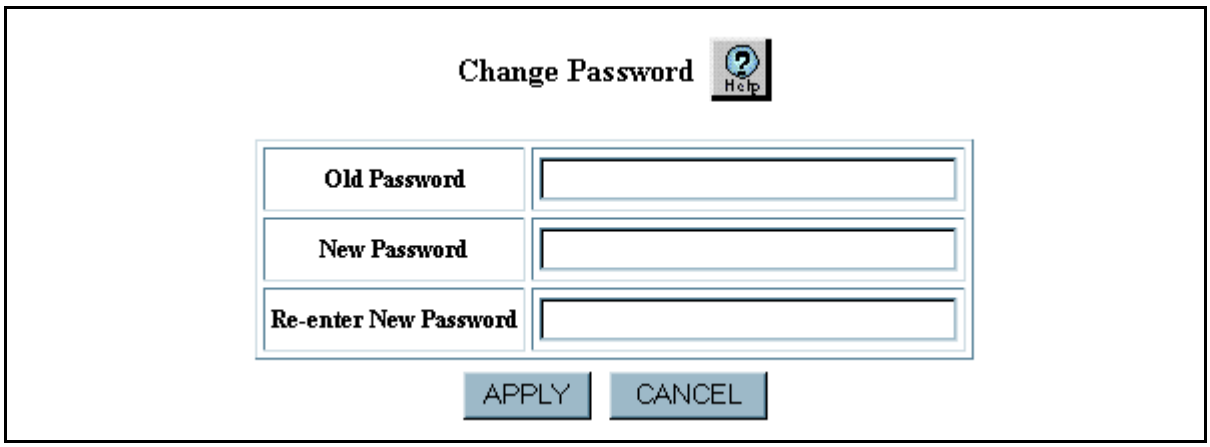
All users can change their own passwords.

Web Agent Procedure

To change your password:

1. In the navigation pane, expand the **System** folder.
2. Click **Change Password**. The Change Password Web page is displayed in the content pane. See Figure 2-12.

Figure 2-12. Change Password Web Page



3. In the **Old Password** field, enter your currently active password.
4. In the **New Password** field, enter a new password. Passwords can consist of a maximum of 31 characters.

   ***Note:** Do not use a combination of the following special characters for the password ;, ?, \,(,),#, $,%, ^, &, or *.

5. In the **Re-enter New Password** field, reenter the password that you entered in the **New Password** field.
6. Click **APPLY**.

CLI Command

To change your password, use the following CLI command:

> password <*passwd*>

Disabling a User Account

For security reasons, you may want to disable certain user accounts, such as the **manuf** and **diag** accounts.

Web Agent Procedure

To disable a user account:

1. Expand the **System** > **Administration** folders.
2. Click **User Accounts**. The User Account Management Web page is displayed in the content pane. See Figure 2-9.
3. Select the **manuf** or **diag** account that you want to disable.
4. Click **Modify**. The Modify User Account Web page is displayed in the content pane. See Figure 2-11.
5. In the **Status** field, select **Disable**.
6. Click **APPLY**.

CLI Command

To disable a user account, use the following CLI command:

`(configure)#` **username <*name*> status disable**

Deleting a User Account

Web Agent Procedure

To delete a user account:

modify a user account:

1. In the navigation pane, expand the **System** > **Administration** folders.
2. Click **User Accounts**. The User Account Management Web page is displayed in the content pane. See Figure 2-9.
3. In the **User Name** field, select the user account that you want to delete.
4. Click **Delete**. The Delete User Account Web page is displayed in the content pane. See Figure 2-13.

Figure 2-13. Delete User Account Web Page



5. Click **YES** to delete the user account.

CLI Command

To delete a user account, use the following CLI command:

`(configure)#` **no username** <***name***>

Changing the Console Serial Port Settings

You can change the communications settings for the serial port connection located on the front panel of the layer 3 supervisor module from the Web Agent. The switch's console port is initially configured as a TTY Console to support a TTY connection. The layer 2 and layer 3 supervisor modules allow you to reconfigure the console serial port as a PPP Console to support a dial-in PPP connection using a modem.

*** Note:** If you reconfigure the serial port as a PPP console, you can only change the switch's baud rate and flow control parameters. The flow control parameters are limited to None or Xon/Xoff.

This section contains procedures for the following tasks:

- Configuring the Serial Console Port as a TTY Console
- Configuring the Serial Console Port as a PPP Console
- Regaining Configuration Access to the PPP Serial Port Console

Configuring the Serial Console Port as a TTY Console

You can configure the serial port as a TTY console using either the Web Agent or the CLI.

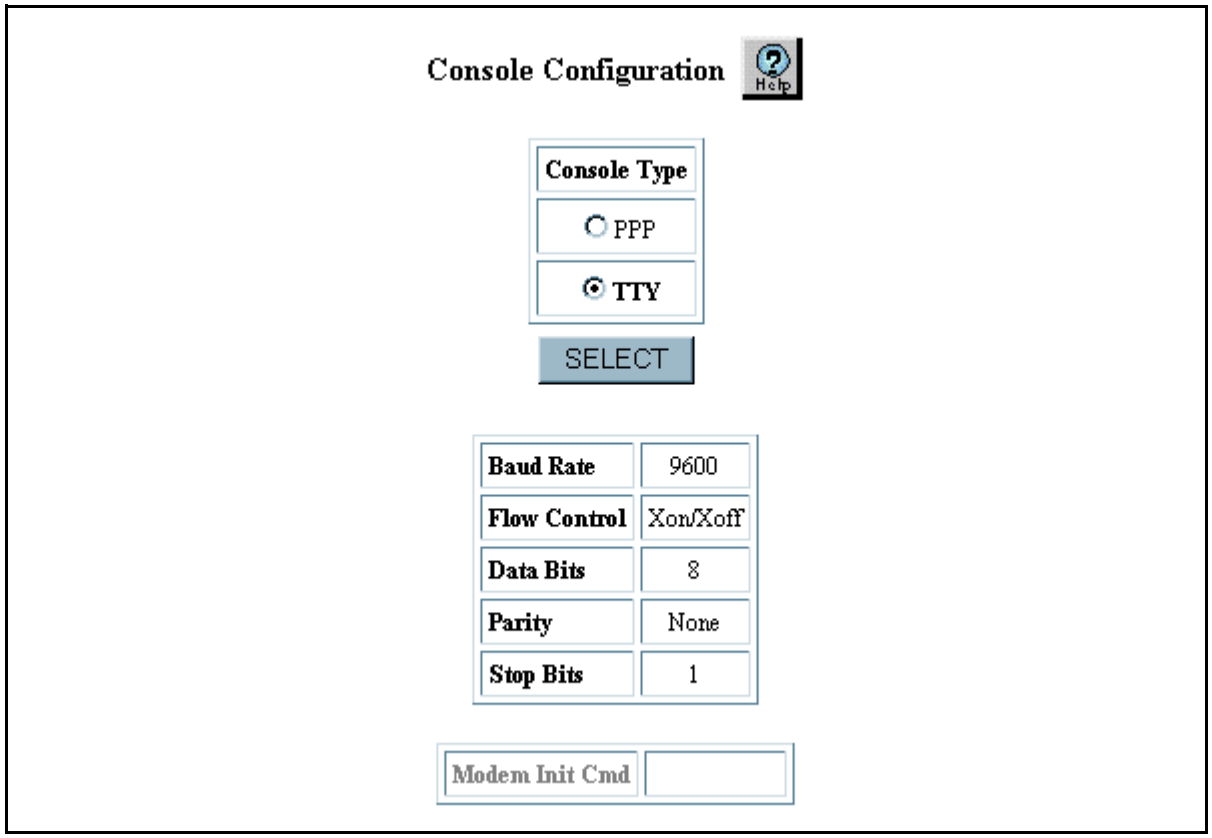
Web Agent Procedure

To configure the console serial port as a TTY Console using the Web Agent:

1. In the navigation pane, expand the **System > Configuration** folders, and then click **Console Configuration**.

   The Console Configuration Web page is displayed in the content pane. See Figure 2-14.

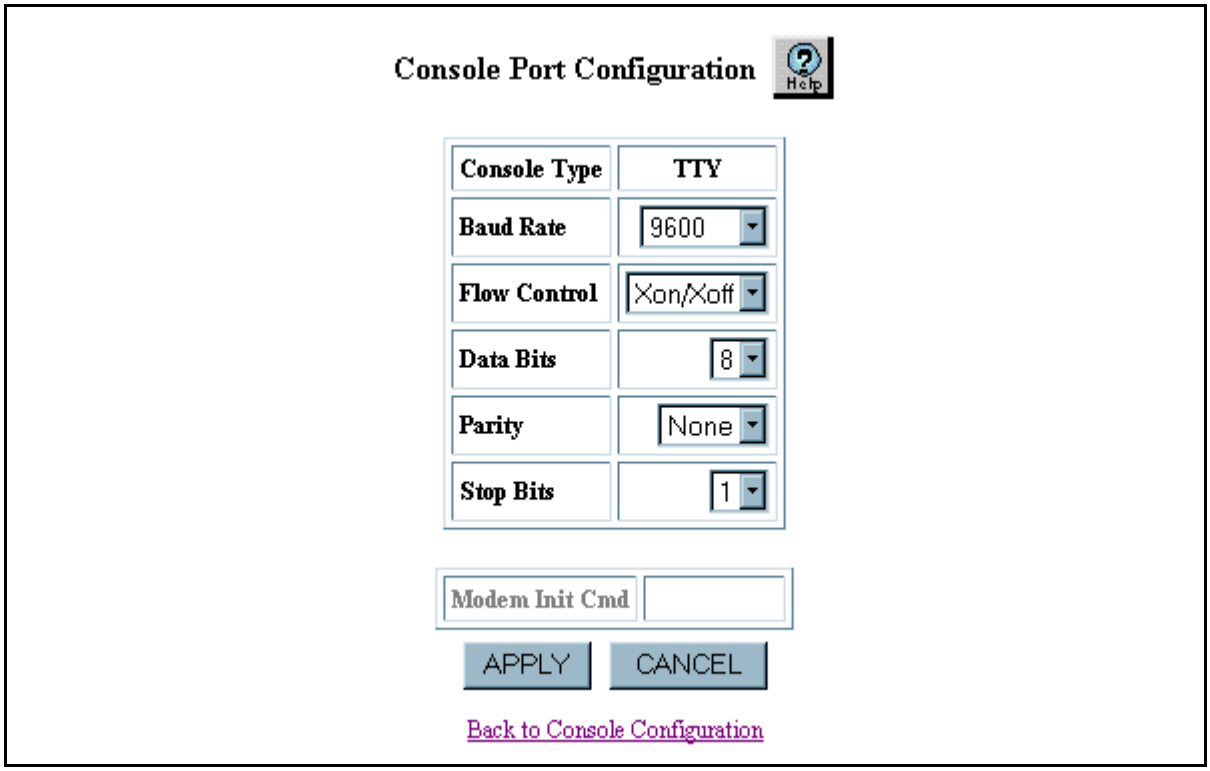
Figure 2-14. Console Configuration Web Page



2. Select **TTY** as the console type and click **SELECT**.

   The Console Port Configuration Web page is displayed in the content pane and displays **TTY** in the **Console Type** field (Figure 2-15).

Figure 2-15. Console Port Configuration Web Page



3. See Table 2-4 and configure the other Console Port Configuration Web page parameters.

4. Click **APPLY** to save your changes, or **CANCEL** to clear your selection.

Table 2-4. Console Port Configuration Web Page Parameters

Option	Default	Available Settings
Baud Rate	9600	300, 1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200
Flow Control	Xon/Xoff (TTY)	None, Xon/Xoff (TTY)
Data Bits	8	7 or 8
Parity	None	Odd, Even, or None
Stop Bits	1	1 or 2

CLI Command

To configure the console serial port as a TTY Console, use the following CLI command:

`(configure)#` **set console type {tty | ppp}**

Configuring the Serial Console Port as a PPP Console

You can reconfigure the serial port as a PPP console to support a dial-in PPP connection using either the Web Agent or the CLI. Before you can do so, however, you must connect a modem to the switch.

*** Note:** If you do not configure an IP address for the PPP serial interface, the switch immediately sends the modem configuration string. If you type TTY, you will re-access the CLI login prompt.

Connecting a Modem

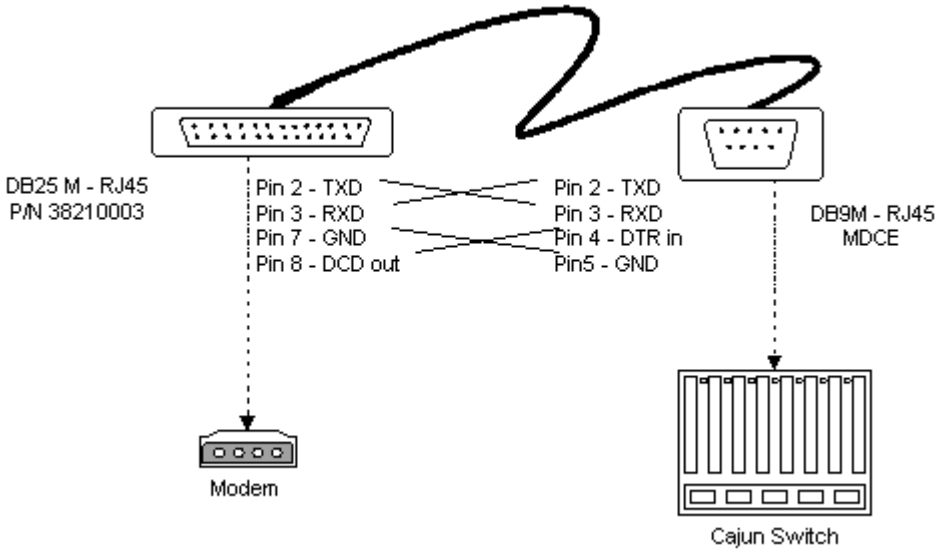
After you connect a modem to your switch with the specified serial cable and connectors, and configure the serial port in PPP mode, the switch will convert the normal Distributed Computing Environment (DCE) interface to a Data Terminal Equipment (DTE) interface that is used by modems.

It then periodically sends the modem configuration string from the serial console port. This synchronizes the baud rates between the modem and the console port and configures the modem to operate with the switch's DTE interface.

To connect a modem (Figure 2-16):

1. Attach a DB25M-RJ45 (P/N 38210003) connector to the modem.
2. Attach the DB9M-RJ45 MDCE connector to the switch's serial console port on the front panel of the switch.

Figure 2-16. Typical Modem Connection to a Switch



Web Agent Procedure

To configure the console serial port as a PPP console from the Web Agent:

1. In the navigation pane, expand the **System > Configuration** folders, and then click **Console**.

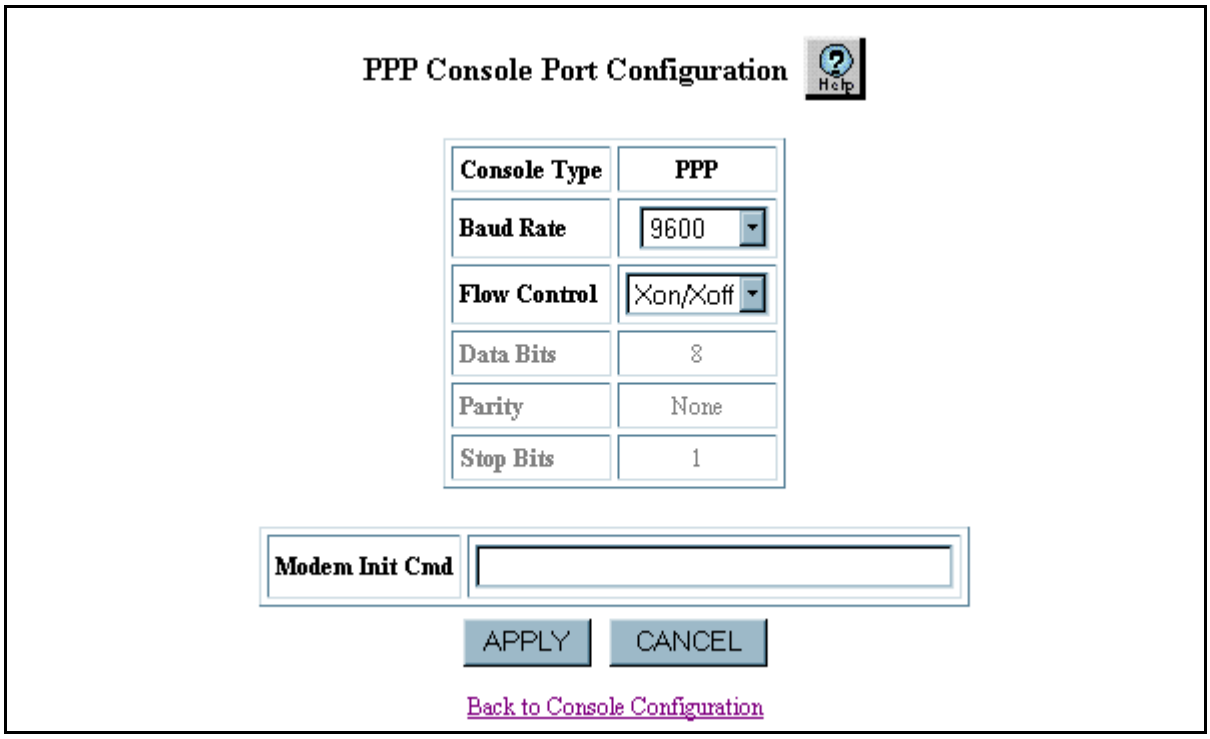
   The Console Configuration Web page is displayed in the content pane. See Figure 2-14.

2. Select **PPP** as the console type and click **SELECT**.

   The Console Port Configuration Web page is displayed in the content pane and displays **PPP** for the Console Type (Figure 2-17).

*** Note:** If you select **PPP**, the PPP Console Port Configuration Web page is displayed.

Figure 2-17. PPP Console Configuration Web Page



3. Select a baud rate from the **Baud Rate** field pull-down menu that is different from the one already selected.

4. Select an option from the **Flow Control** field pull-down menu. The options are: None and Xon/Xoff.

   ***Note:** You can only change the baud rate and flow control parameters after you configure the serial port as a PPP console.

5. Enter the modem initialization command in the **Modem Init Cmd** field. The default modem configuration init command is AT&D0SO=1.

   See Table 2-5 for a definition of the Modem Configuration Init command parameters.

Table 2-5. Modem Configuration String Parameters

Parameter	Definition
&D0	Disable DTR
S0=1	Auto-answer mode (one ring)
CD follows carrier	Depends on modem
E0	Disable local echo
Software Flow Control (Receive and Transmit)	Depends on modem

*** Note:** If you misconfigure the PPP serial console port, you can regain CLI access to correct the configuration parameters. For more information, see "Regaining Configuration Access to the PPP Serial Port Console" later in this chapter.

Safety Tip: To successfully dial-in with PPP to the switch, you must also configure an IP address and interface for the PPP Serial Interface (Serial-Console). See "Managing Configuration Files," later in this chapter for more information.

6. Click **APPLY** to save your changes, or **CANCEL** to clear your selection.

*** Note:** If you click **Apply**, the changes are saved in the Running config only. The Startup config has not changed. Therefore, these and other changes will be lost if your switch goes down or if you power it off.
To save these changes, you must copy the Running config to the Startup config. See "Copying Configuration Files" later in this section.

CLI Command

To configure the console serial port as a TTY console after you install a mode, use the following CLI command:

`(configure)#` **set console type {tty | ppp}**

Regaining Configuration Access to the PPP Serial Port Console

If you incorrectly configure the PPP serial port console, you cannot regain configuration access to the Web Agent or the CLI. This is a temporary situation and can corrected.

To regain configuration access to the Web Agent or the CLI when the console port is in PPP mode:

1. Enter `TTY` in the console window. The CLI login prompt displays. It may be necessary to press Enter several times to see the login prompt. You may see the modem init command string.
2. Enter your user name at the **Login** prompt. The password prompt displays.
3. Enter your password at the **Password** prompt. The switch CLI prompt displays.

   Enter the PPP configuration commands necessary to start PPP. See *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1* for details about these PPP commands.
4. Enter **`exit`** at the CLI prompt after you complete your configuration settings to reinvoke the modem control software and exit CLI mode.

*** Note:** You do not need to exit from CLI if the serial port console has been configured as a TTY console, or if you do not intend to re-attempt connecting using PPP.

Configuring Dial-Up Networking

To configure your PC for dial-up networking with a PPP serial port console:

1. Open **My Computer** on your PC.
2. Double-click **Dial-Up Networking**. The **Dial-Up Networking** window opens.

   ***Note:** You must have dial-up networking installed on your PC.
3. Double-click **Make New Connection** to configure your modem. The **Make New Connection** wizard opens.
4. Enter a connection name for the computer you are dialing in the Type a name for the computer you are dialing field.

5. Open the **Select a Modem** pull-down menu.

6. Select your modem and select **Configure**. The **Modem Properties** window opens.

7. Select the **Connection** tab. The **Connection** window opens.

8. Select the following from the **Connection Preferences** window:

   — Open the Data bits pull-down menu and select 8.

   — Open the Parity pull-down menu and select None.

   — Open the Stop bits pull-down menu and select 1.

9. Select **Advanced**. The **Advanced Connection Settings** window opens.

10. Select Software (XON/XOFF) from the **Use flow** control field.

11. Select **OK** to close the window. The **Modem Properties** window reopens.

12. Select the **Options** tab. The **Options** window opens.

13. Select **Bring up terminal window** after dialing from the **Connection Control** field. Select **OK**. The **Modem Properties** window closes and the wizard continues.

14. Select **Next** from the wizard window and enter the telephone number you are calling.

15. Select **Next**. The wizard reports that you have successfully configured a modem.

16. Select **Finish**. The wizard closes and the newly configured connection displays in your Dial-Up networking program folder.

17. Select your new connection right click the mouse.

18. Select **Properties**. The **Properties** window opens.

19. Select the **Server Types** tab and de-select all advanced options except TCP/IP.

20. Select **TCP/IP** Settings. The **TCP/IP Settings** window opens.

21. Select Specify an IP Address and enter the IP address fro the serial port interface.

22. Select **Server** assigned name server addresses.

23. De-select **Use IP header Compression** and **Use Default Gateway** on remote network.

24. Select **OK** to close the TCP/IP window.

25. Select **OK** to close the **New Connections Properties** window.

Using Dial-Up Networking with a PPP Serial Port Console

To use TCP/IP applications (Telnet, HTTP, and SNMP) over your PPP serial port interface:

1. Open **My Computer** on your PC.

2. Double-click **Dial-Up Networking**. The **Dial-Up Networking** program folder opens.

   ***Note:** You must have dial-up networking installed on your PC.

3. Double-click the PPP modem you previously created. The **Connect To** window opens.

4. Enter your password and select **Connect**. A **Pre-Dial Terminal** screen opens.

   When the modem has successfully connected, a **Post-Dial Terminal** screen opens.

5. Login in the **Post-Dial Terminal** screen using your CLI user name and password. It may be necessary to enter several carriage returns to view the Login prompt.

6. At the CLI prompt, go to configuration mode.

7. Enter **`set console transfer PPP`**. ASCII characters display below the CLI prompt. This is typical while the switch attempts to connect via PPP.

8. Select **Continue** (F7) from the **Post-Dial Terminal** screen. PPP verification completes and the **Connected To** window displays a message that the modem connection has been successfully established.

Changing the TCP Ports for HTTP and Telnet

Overview

For security purposes, you may want to change the TCP port number for Telnet requests and HTTP requests. The default settings for these ports are the well-known TCP ports. HTTP requests use TCP port 80, and Telnet requests use TCP port 23.

Once you change the TCP port number for HTTP, only users who know the new port number can open the Web Agent. And once you change the TCP port for Telnet, only users who know the new port number can start Telnet sessions to the switch.

When you change the TCP port number for either of these protocols, the change takes effect immediately and all connections through the previous port number are disconnected. Any changes that you make to these TCP port numbers are retained if you reset the switch or if the primary supervisor module fails over to the standby supervisor.

This section contains procedures for the following tasks:

- Changing the TCP Port Number for HTTP and Telnet Requests
- Starting a Telnet Session
- Opening the Web Agent

You must have administrative privilege to view or change the TCP port for HTTP or Telnet.

Changing the TCP Port Number for HTTP and Telnet Requests

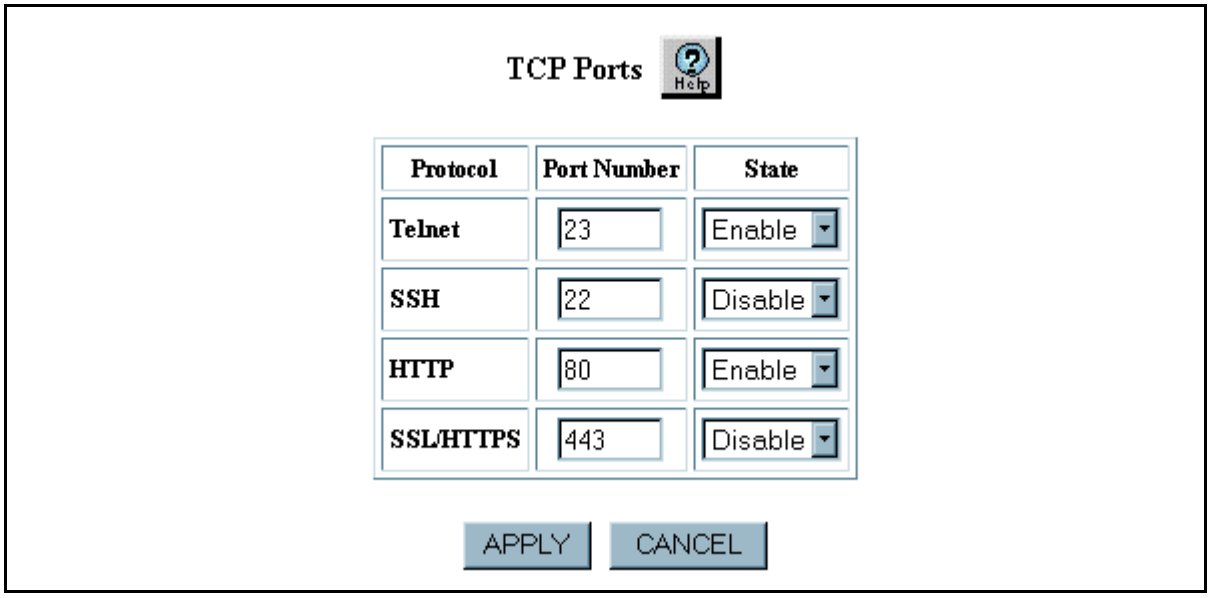
Web Agent Procedure

To change the TCP port for HTTP requests or Telnet requests by using the Web Agent:

1. In the navigation pane, expand the **System** > **Administration** folders.
2. Click **TCP ports**.

   The TCP Ports Web page is displayed in the content pane. See Figure 2-18.

Figure 2-18. TCP Ports Web page



3. To change the TCP port for Telnet requests, in the **Port Number** field for Telnet, enter the port number that you want to use.

   Valid port numbers are 23 or a port number from 9000 through 65355. The default port for Telnet is port 23.

4. To change the TCP port for SSH requests, in the **Port Number** field for SSH, enter the port number that you want to use.

   Valid port numbers are 22 or a port number from 9000 through 65535. The default port for SSH is port 22.

5. To change the TCP port for HTTP requests, in the **Port Number** field for HTTP, enter the port number that you want to use.

   Valid port numbers are 80 or a port number from 9000 through 65535. The default port for HTTP is port 80.

6. To change the TCP port for HTTPS requests, in the **Port Number** field for SSL/HTTPS, enter the port number that you want to use.

   Valid port numbers are 443 or a port number from 9000 through 65535. The default port for HTTPS is 443.

7. Click **Apply**.

CLI Commands

To change the TCP port for Telnet requests, use the following CLI command:

`(configure)#` **ip telnet port** <***tcp-telnet-port***>

To change the TCP port for HTTP requests, use the following CLI command:

`(configure)#` **ip http port** <***tcp-http-port***>

To view the current TCP port settings for Telnet and HTTP, use the following CLI command:

> show tcp configuration

*** Note:** You must have administrative privilege to enter these commands.

Starting a Telnet Session

After changing the TCP port for Telnet requests to a port number other than 23, you must specify the TCP port number in addition to the IP address or host name to start a Telnet session.

For example, if you change the TCP port from 23 to 9998 on switch 192.168.0.126, enter **telnet 192.168.0.126 9998** to start a Telnet session.

Opening the Web Agent

After changing the TCP port for HTTP requests to a port number other than 80, you must specify the TCP port number in addition to the IP address to open the Web Agent.

For example, if you change the TCP port from 80 to 9999 on switch 192.168.0.126, enter **http://192.168.0.126:9999** to open the Web Agent.

Managing Configuration Files

When you first install the switch, or upgrade from a previous installation, your configuration parameters are stored in a **startup.txt** file located in the switch's Non-Volatile Random Access Memory (NVRAM). When the switch is restarted, the startup.txt file runs and stores configuration parameters in volatile RAM as a running configuration.

Any changes you make to the switch configuration are automatically recorded in RAM, but not in NVRAM. When you want to retain your current configuration, you must manually save it to NVRAM through the Web Agent or the CLI.

For information about how to save your running configuration file to your startup configuration, see "Copying Configuration Files" later in this chapter. Always view and compare your running and startup configuration files to determine changes that you made to your running configuration.

***Important:** Do not copy the startup configuration to the running configuration. The switch does not support this activity. If you need to reapply the startup configuration, restart the switch.

*** Note:** If you initialize NVRAM (**nvram initialize** command), all switch settings except the following are reset to their default:

- Startup image
- Fabric mode
- 48-port mode

You can manage the files that contain the configuration data for your multiservice switch from either the Web Agent or the CLI. This section contains procedures for the following tasks:

- Viewing the Running Configuration
- Viewing the Startup Configuration
- Viewing the Script Execution Log File
- Copying Configuration Files
- Copying Files
- Viewing the Status of a TFTP Transfer

Viewing the Running Configuration

You must have administrator access to view the running configuration.

Web Agent Procedure

To view the running configuration:

1. In the navigation pane, expand the **System** > **Configuration** > **Configuration Files** folders.
2. Click **Running Config**. The Running Configuration Web page is displayed in the content pane.

CLI Command

To view the running configuration, use the following CLI command:

```
# show running-config
```

Viewing the Startup Configuration

You must have administrator access to view the startup configuration.

Web Agent Procedure

To view the startup configuration:

1. In the navigation pane, expand the **System** > **Configuration** > **Configuration Files** folders.
2. Click **Startup Config**. The Startup Configuration Web page is displayed in the content pane.

CLI Command

To view the startup configuration, use the following CLI command:

```
# show startup-config
```

Viewing the Script Execution Log File

Each time the startup.txt file or other script runs, a log file is generated. Log files contain the data returned from the script. You can view log file data from the Script Execution Log File using either the Web Agent or the CLI.

Web Agent Procedure

To view your Script Execution Log File from the Web Agent, select **Script Log File** from the **Configuration Management** folder in the navigation pane. The **Script Execution Log file** displays

CLI Command

To view your Script Execution Log File from the CLI, use the following CLI command:

```
# show file_name logfile.txt
```

Copying Configuration Files

You must have administrator access to copy the configuration files.

If you make any changes, you must save the running configuration as the startup configuration to save the changes. Before you do this, always copy your startup configuration to a file on the switch or on a TFTP server.

You can save the running configuration to the startup configuration and copy the startup configuration to a file using either the Web Agent or the CLI.

Configuration files are automatically saved as text files using the *.txt extension on your switch. If you save the new file to a TFTP server, you can edit the startup.txt file using a text editor of your choice and save copies of it with a **.txt** extension.

***Important:** Do not copy the startup configuration to the running configuration. The switch does not support this activity. If you need to reapply the startup configuration, restart the switch.

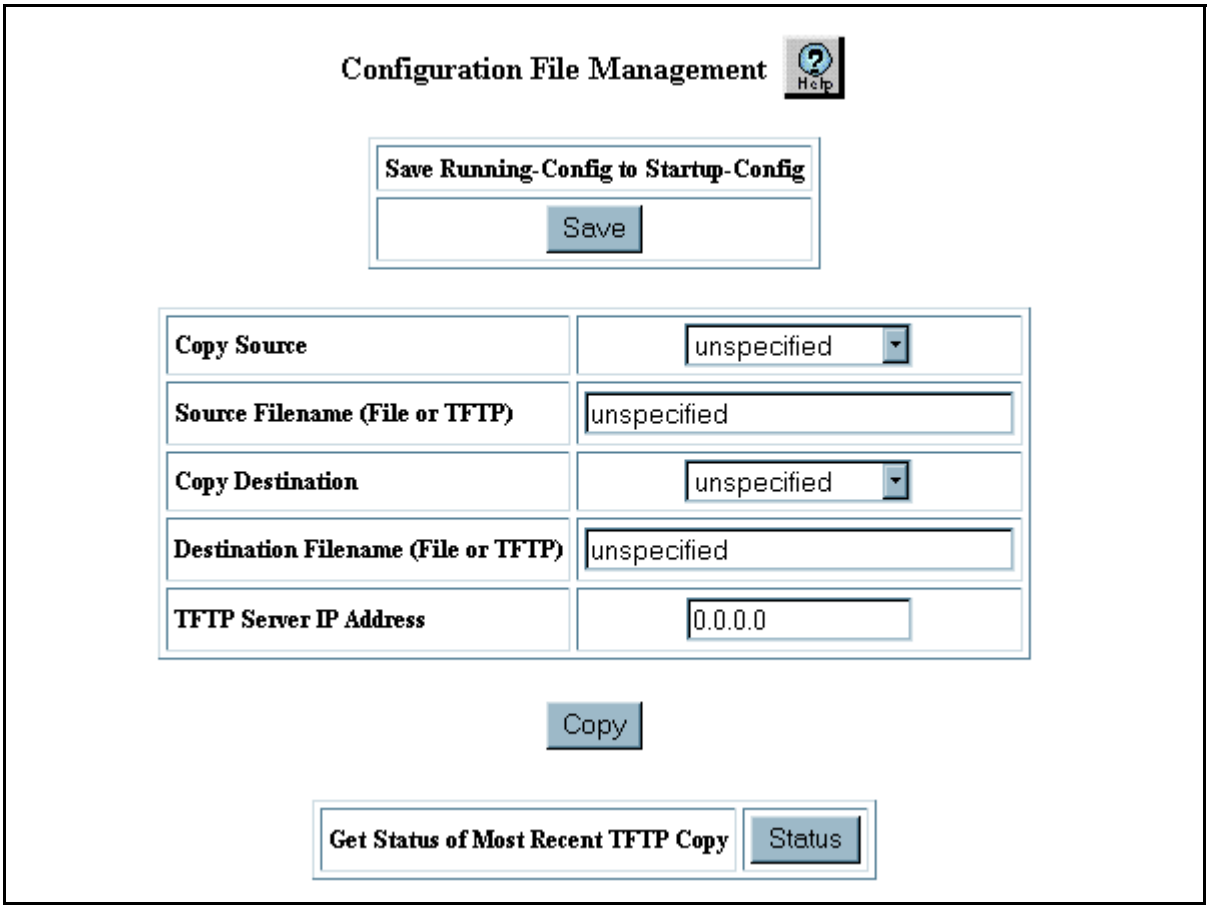
Web Agent Procedure

To save your running configuration as your startup configuration in NVRAM using the Web Agent:

1. In the navigation pane, expand the **System** > **Configuration** > **Configuration Files** folders, and then click **File Management**.

   The Configuration File Management Web page is displayed in the content pane. See Figure 2-19.

Figure 2-19. Configuration File Management Web Page



2. Click **Save**. The running configuration is saved as the startup configuration.

*** Note:** You can also save the running configuration to the startup configuration through the Configuration File Management Web page. See "Copying Files" for that procedure.

CLI Command

To save your running configuration as your startup configuration in NVRAM, use the following CLI command in **Enable** mode:

\# **copy running-config startup-config**

Copying Files

You must have administrator access to copy text files and BOOT, APP1, and APP2 images.

You can copy files to and from multiple locations. For example, if you modify the running configuration and you want to reinstate your startup configuration parameters, you can copy your startup configuration to your running configuration in volatile RAM. You can also upload or download configuration files by copying files from a TFTP server directory to the switch or to a startup or running configuration file. You can copy files using either the Web Agent or the CLI.

Web Agent Procedure

To copy files using the Web Agent:

1. In the navigation pane, expand the **System** > **Configuration** > **Configuration Files** folders, and then click **File Management**.

   The Configuration File Management Web page is displayed in the content pane. See Figure 2-19.

2. See Table 2-6 for an explanation of the Configuration File Management Web page parameters.

Table 2-6. Configuration File Management Web Page Parameters

Parameter	Description	Options
Save Running-Config to Startup-Config	Saves the running configuration to the startup configuration	N/A
Copy Source	Specifies the source file to be copied	• **Unspecified** - Specifies an initialized value. • **File** - Specifies a source file located on the switch in NVRAM. • **Running-Config** - Specifies a running configuration. • **Startup-Config** - Specifies a startup configuration. • **TFTP Server** - Specifies a source file located in a directory on a TFTP server.
Source Filename	Specifies the path and name of the source file	Source files can be ASCII files in NVRAM available for upload or files located on a TFTP server available for download
		1 of 2

Table 2-6. Configuration File Management Web Page Parameters

Parameter	Description	Options
Copy Destination	Specifies the location of the destination	• Unspecified- Specifies an initialized value. • **File** - Specifies that a source is copied or downloaded to NVRAM. • **Running-Config** - Specifies that a source is copied to the running configuration. • **Startup-Config** - Specifies that a source is copied to the startup configuration. • **TFTP Server** - Specifies that a source is copied to a TFTP server location.
Destination Filename	Specifies the path and name of the destination file	Configuration files and other files can be copied to NVRAM on the switch or to a TFTP server as a destination location.
TFTP Server IP Address	Specifies the IP address of a source or destination TFTP server	Copy (download) source files, located on a TFTP server, to your running configuration, your startup configuration, or a location on the switch. Or, copy (upload) your configuration files, or a file located on the switch, to a TFTP server.
		2 of 2

— File — To copy a file stored on the switch to your running or startup configuration, to a location on the switch, or to a location on a TFTP server.

— TFTP Server — To copy a file stored on a TFTP server to your running or startup configuration or to a file on the switch.

***Note:** If you select File or TFTP Server, you must also provide the path and filename of the source file in the Destination Filename field.

3. Select one of the following from the Copy Destination pull-down menu.

— **Running-Config** — Copies your startup configuration, or other file located on the switch or on a TFTP server, to your running configuration.

— **Startup-Config** — Copies your running configuration, or other file located on the switch or on a TFTP server, to your startup configuration.

— **File** — Copies your startup or running configuration, another file located on the switch, or a file located on a TFTP server to a file on the switch.

— **TFTP Server** — Copies your startup or running configuration, or another file located on the switch, to a location on a TFTP server.

***Note:** If you select File or TFTP Server, you must also provide the path and filename of the destination file in the Source Destination field.

4. Enter the IP address of the source or destination TFTP server in the TFTP Server IP Address field, if applicable.

5. Click **Copy**. The source configuration or file is copied to your specified destination.

***Note:** The Web Agent displays an `Invalid operation!` error message if you attempt to copy one of the following:

- The current running configuration to the running configuration.
- The startup configuration to the same startup configuration.
- The specified TFTP server to a TFTP server.

CLI Command

To copy files, use the following CLI command in Enable mode:

copy running-config +

*** Note:** Entering a **+** sign lists all of the applicable options for the **copy running-config** command

Viewing the Status of a TFTP Transfer

After you have copied the startup configuration or other files to a TFTP server, you can use either the Web Agent or the CLI to check the status of the TFTP transfer to ensure that files copied correctly.

Web Agent Procedure

To view the status of a TFTP transfer using the Web Agent:

1. In the navigation pane, expand the **System** > **Configuration** > **Configuration Files** folders, and then click **File Management**.

   The Configuration File Management Web page is displayed in the content pane. See Figure 2-19.

2. Select **Status** from the **Get Status of Most Recent TFTP Copy** field. A status message displays in the Status pane.

CLI Command

While doing a TFTP transfer using the CLI, the status, either successfully transferring or an error condition, is immediately returned.

*** Note:** You cannot view the status of a TFTP transfer done using the CLI from the Get Status of Most Recent TFTP Copy field on the Web Agent.

> show startup-config

3 Configuring System Information

Overview

You can manage system information by using either the Web Agent or the CLI. The following procedures to manage your system information are provided in this chapter:

- Entering General System Information
- Enabling the Simple Network Time Protocol
- Setting Summer Time Hours
- Setting the System Clock
- Setting the Temperature System
- Displaying the Power System Statistics
- Displaying Cooling System Statistics
- Performing a Reset

*** Note:** The last step in each procedure tells you to click **Apply** to save the setup or changes that you made. This step saves the setup or any changes to the running configuration only. The startup configuration is not changed. Therefore, these and other changes will be lost if the switch goes down or if you turn it off.

To save any changes to the startup configuration, you must copy the running configuration to the startup configuration. For information on how to perform this task, see "Copying Configuration Files" in Chapter 2, "Setting Up the Switch."

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Entering General System Information

You can enter general information about your system by using either the Web Agent or the CLI.

Web Agent Procedure

You can enter general system information using the following Web Agent fields:

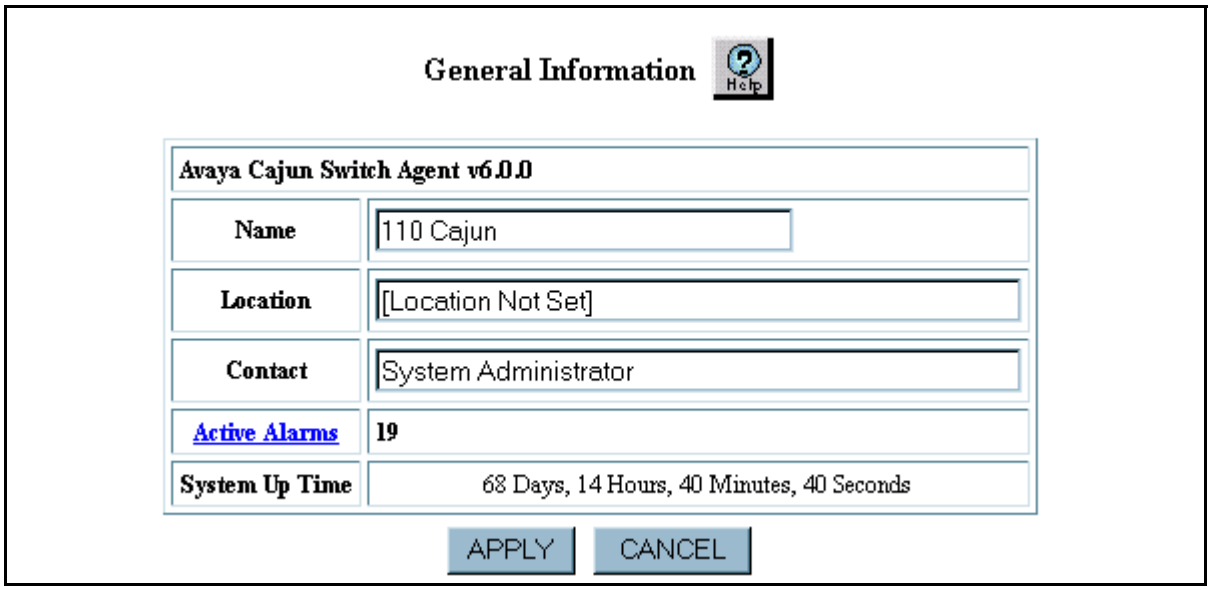
- Switch name
- Device location
- Device contact

To enter general system information from the Web Agent:

1. In the navigation pane, expand the **System** folder, and then click **General Information**.

   The General Information Web page is displayed in the content pane. See Figure 3-1.

Figure 3-1. General Information Web Page



2. Enter the switch name in the Name field.
3. Enter the location for the switch (for example, floor, closet in the Location field.

4. Enter the person who should be contacted in the event of a problem in the Contact field.

5. Click APPLY to save your changes, or CANCEL to clear your selection.

Enabling the Simple Network Time Protocol

You can enable Simple Network Time Protocol (SNTP) on your switch using either the Web Agent or the CLI. Enabling SNTP automatically synchronizes time on all computers, switches, and other devices connected to your switch.

When you enable SNTP, you are required to set your time zone and the rule or dates of Summer Time Hours for your location. For information about setting one-time summer hours, see "Setting One-Time Summer Time Hours" later in this chapter. For information about setting Summer Time Hours, see "Setting Summer Time Hours" later in this chapter.

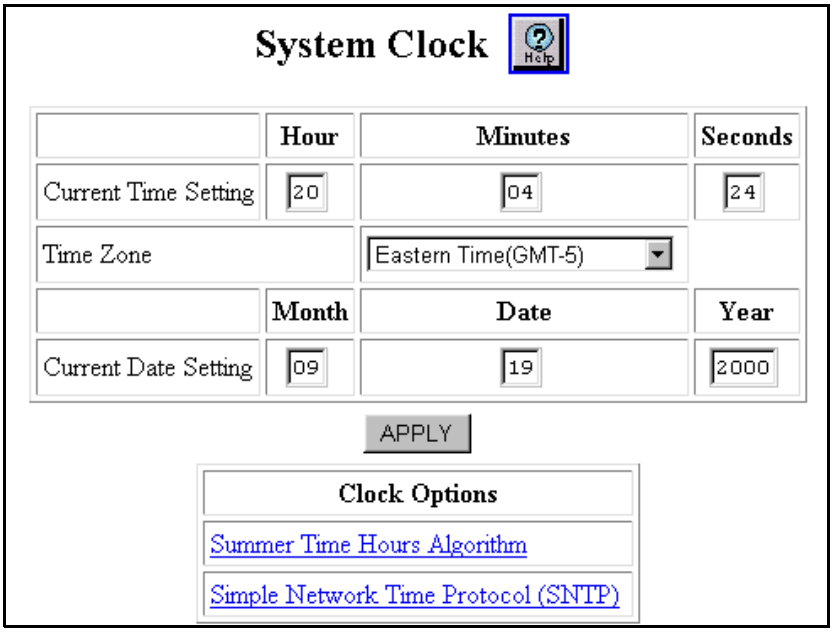
Web Agent Procedure

To enable SNTP on your switch using the Web Agent:

1. In the navigation pane, expand the **System** > **Configuration** folders, and then click **System Clock**.

   The System Clock Web page is displayed in the content pane. See Figure 3-2.

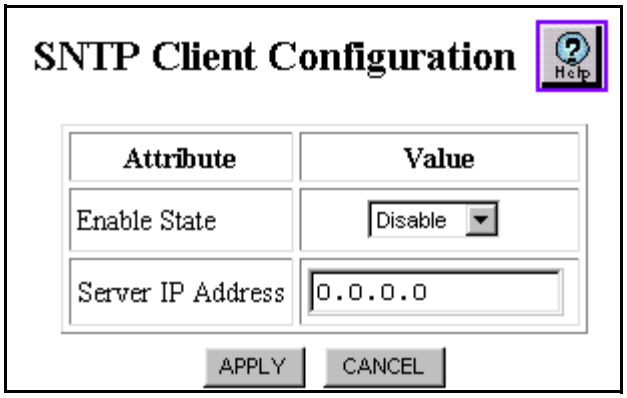
Figure 3-2. System Clock Web page.



2. Select Simple Network Time Protocol (SNTP) from the Clock Options box.

   The SNTP Client Configuration Web page is displayed in the content pane. See Figure 3-3.

Figure 3-3. SNTP Client Configuration Web Page



3. Select Enable from the Enable State pull-down menu. The default is Disable.

4. Enter the server IP address in the Server IP address field for the switch on which you want to enable SNTP.

5. Click APPLY to save your changes, or CANCEL to clear your selection.

CLI Command

To enable SNTP on your switch using the CLI, enter the following command in Configure mode:

`(configure)#` **sntp server** <***ip-address***>

Setting Summer Time Hours

Summer Time Hours, also referred to as Daylight Savings Time (DST), is the strategy of moving clocks ahead to provide greater amounts of daylight in the afternoon and to standardize time with other parts of the world. In many parts of the world, the Summer Time Hours algorithm is based on a standardized rule. For example, in the Western hemisphere, the rule used by most locations in Canada, Mexico, and the United States is to set clocks forward by one hour at 2:00 a.m. on the first Sunday in April and back an hour at 2:00 a.m. on the first Sunday in October annually. Many countries in Europe and Asia follow similar rules. The offset, or amount of time by which the clock is set forward or backward, varies from country to country.

Many parts of the world follow a one-time change of Summer Time Hours. When you configure the switch for these locations, you reset the clock by specifying a scheduled time and date.

This section provides the following procedures:

- Setting Recurring Summer Time Hours
- Setting One-Time Summer Time Hours

*** Note:** If you upgrade your switch from a previous version, the local time settings are saved as Greenwich Mean Time values. Always change the Summer Time Hours Algorithm before you set the clock.

Setting Recurring Summer Time Hours

You can set recurring summer time hours using either the Web Agent or the CLI.

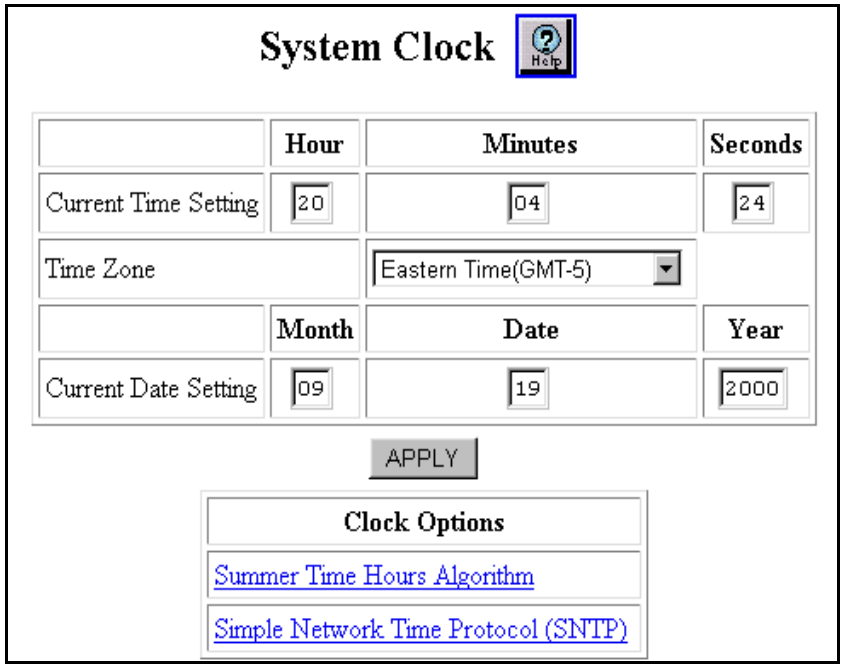
Web Agent Procedure

To set recurring Summer Time Hours using the Web Agent:

1. In the navigation pane, expand the **System** > **Configuration** folders, and then click **System Clock**.

   The System Clock Web page is displayed in the content pane. See Figure 3-4.

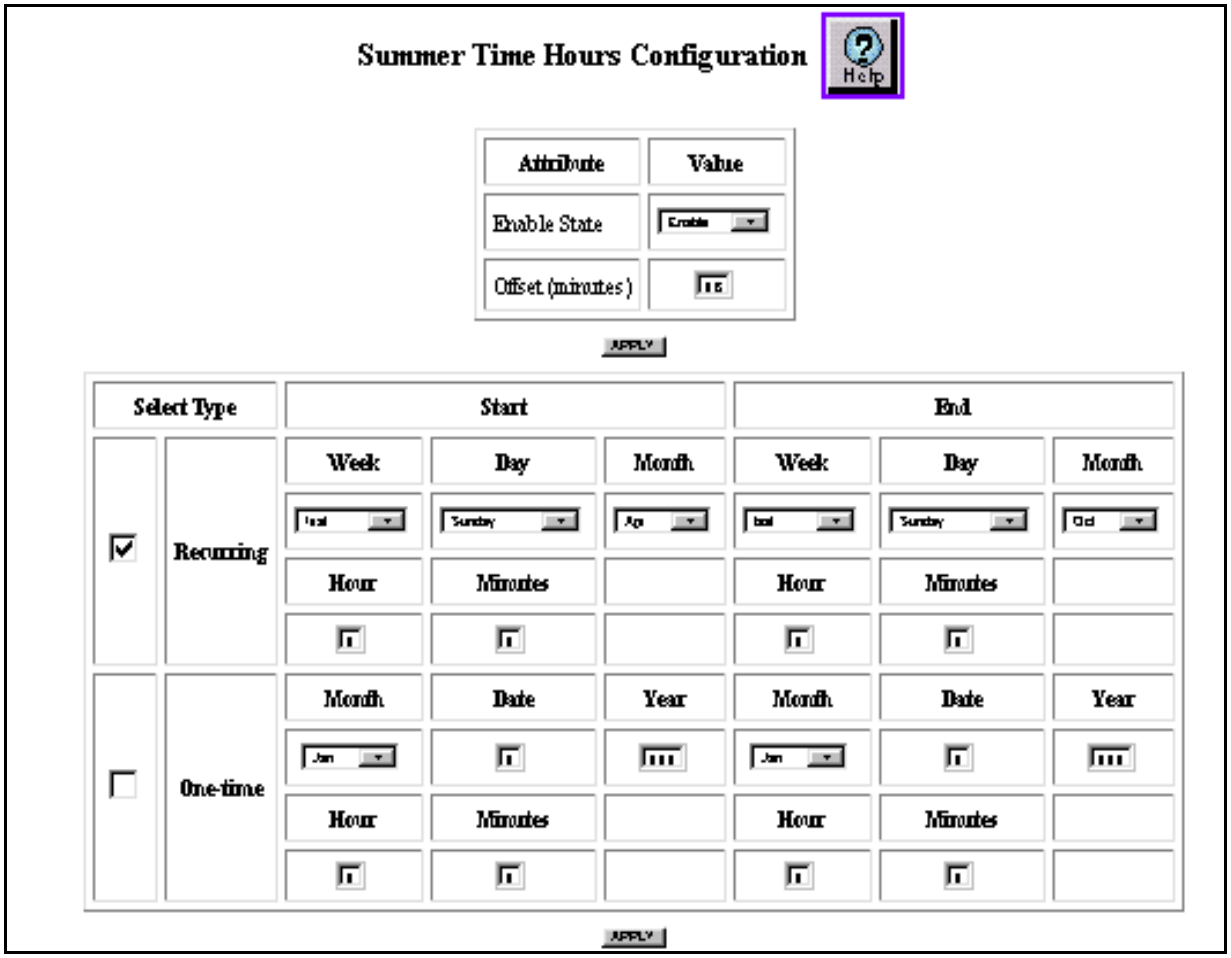
Figure 3-4. System Clock Web Page



2. Select **Summer Time Hours Algorithm** from the **Clock Options** field.

   The Summer Time Hours Configuration Web page is displayed in the content pane. See Figure 3-5.

Figure 3-5. Summer Time Hours Web Page



3. Select **Enable** from the **Enable State** pull-down menu.

4. Enter the reset value for the clock in minutes in the **Offset** field. For example, if you intend to reset the clock forward or backward by one hour, keep the default value of 60 minutes.

5. Set the Summer Time Hours that recur annually:

   a. Select the check box in the **Recurring** field.

   b. Select the values for the **Week**, **Day**, and **Month** when the Summer Time Hours are to start and end.

   c. Enter values for the **Hour** and **Minutes** when Summer Time Hours are to start and end.

6. Click **APPLY** to save your changes, or **CANCEL** to clear your selection.

See Table 3-1 for detailed information about the fields in the Recurring Summer Time Hours section of the Summer Time Hours Configuration Web page.

Table 3-1. Recurring Summer Time Hours Parameters

Parameter	Definition
Recurring	Select if the Summer Time Hours option is defined by a rule such as Daylight Savings Time (DST - all Start and End fields associated with Recurring Summer Time Hours provide the default values for DST). When you select the Recurring Summer Time Hours option, indicate the time, in hours and minutes, on a specified day, week, and month that you want the Summer Time Hours begin and end.
Start	Specifies the start of Summer Time Hours.
End	Specifies the end of Summer Time Hours.
Week	Select the week during which you want recurring Summer Time Hours to start or end. The selected week should reflect the day on which Summer Time Hours start. For example, if Summer Time Hours start on the first Sunday in April, select the First week. Options include: • **First -** First week of the month, the default Start value, when Daylight Savings Time starts in the Western hemisphere. • **Second -** Second week of the month. • **Third -** Third week of the month. • **Fourth -** Fourth week of the month. • **Last -** Remaining days of the month that form the last week of the month. Last is specified as the default End value, denoting when Daylight Savings Time ends in the Western hemisphere.
	1 of 2

Table 3-1. Recurring Summer Time Hours Parameters

Parameter	Definition
Day	Select the day of the week when you want recurring Summer Time Hours to start or end. Options are based on a seven-day week and include: • **Sunday -** the default Start and End values. In the Western hemisphere, DST starts on the first Sunday in April and ends on the last Sunday in October. • Monday • Tuesday • Wednesday • Thursday • Friday • Saturday
Month	Select the month when recurring Summer Time Hours start or end. The twelve months of the Gregorian calendar are provided. For recurring Summer Time Hours, the default Start value is April, the month during which DST starts in the Western hemisphere. The default End value is October, the month during which DST ends in the Western hemisphere.
Hour	Enter a value to represent the hour when Summer Time Hours start or end for Recurring settings. For Recurring Summer Time Hours, the default value is 02, meaning 2:00 a.m., for both Start and End hours.
Minutes	Enter a value to represent the number of minutes into the hour when Summer Time Hours start or end for Recurring Summer Time Hours. The default value is 00 for both Start and End minutes.
	2 of 2

CLI Command

To set recurring Summer Time Hours using the CLI, enter the following command from Configure mode:

`(configure)#` **clock summer-time recurring <*week*>**

Setting One-Time Summer Time Hours

You can set one-time summer time hours using either the Web Agent or the CLI.

Web Agent Procedure

To set Summer Time Hours that are not based on a standard rule using the Web Agent:

1. In the navigation pane, expand the **System** > **Configuration** folders, and then click **System Clock**.

   The System Clock Web page is displayed in the content pane. See Figure 3-4.

2. In the **Clock Options** field, click **Summer Time Hours Algorithm**. The Summer Time Hours Configuration Web page is displayed in the content pane. See Figure 3-5.

3. Select Enable from the Enable State pull-down menu. The default is Disable.

4. Enter the reset value for the clock in minutes in the Offset field. For example, if you intend to reset the clock forward or backward by one hour, keep the default value of 60 minutes.

5. Select the check box next to the One-Time field to set the date and time for Summer Time Hours on a one-time basis:

   a. Enter the specific Month, Day, and Year when the Summer Time Hours start and end.

   b. Enter the Hour and Minutes when the Summer Time Hours start and end.

6. Click APPLY to save your changes, or CANCEL to clear your selection.

See Table 3-2 for detailed information about the One-Time field parameters for the Summer Time Hours Configuration Web page.

Table 3-2. One-Time Summer Time Hours Configuration

Parameter	Definition
One-time	Select if Summer Time Hours change one time, such as on a specified date. When you select one-time Summer Time Hours, you indicate the time and date on which Summer Time Hours begin and end.
	1 of 2

Table 3-2. One-Time Summer Time Hours Configuration

Parameter	Definition
Start	Specifies the start of Summer Time Hours.
End	Specifies the end of Summer Time Hours.
	2 of 2

CLI Command

To set Summer Time Hours that are not based on a standard rule using the CLI, enter the following command from Configure mode:

```
(configure)# clock summer-time date
```

Setting the System Clock

The system clock is used for setting traps, alarms, and other events on the switch.

*** Note:** You must set SNTP and Summer Time Hours before you can set the system clock. See "Enabling the Simple Network Time Protocol" and "Setting Summer Time Hours" earlier in this chapter.

*** Note:** The system clock does not automatically change with Daylight Savings Time.

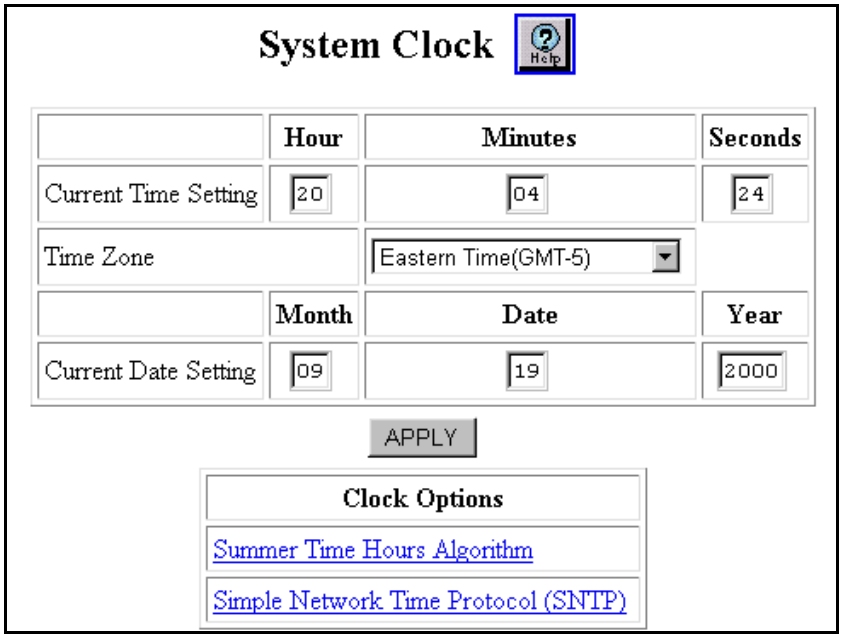
You can set the system clock from either the Web Agent or the CLI.

Web Agent Procedure

To set the system clock using the Web Agent:

1. In the navigation pane, expand the **System** > **Configuration** folders, and then click **System Clock**. The System Clock Web page is displayed in the content pane. See Figure 3-6.

Figure 3-6. System Clock Web Page



2. Enter the time in the **Current Time Setting** Hour, Minutes, and Seconds fields using 24-hour time format (for example, 10 p.m. is 22:00 00.

3. Select the time zone for your area from the **Time Zone** pull-down menu.

4. Enter the current Month, Date, and Year in the **Current Date Setting** fields.

5. Click **APPLY** to save your changes, or **CANCEL** to clear your selection.

CLI Command

To set the system clock using the CLI, enter the following command:

clock set <***time***> <***date***> <***year***>

Setting the Temperature System

You can set the upper and lower temperature warning systems for the switch backplane and slot 1. These warning systems cause the following to happen if the temperatures you set are reached:

- **Shutdown Temperature—**The switch shuts down if this temperature is reached.
- **Upper Warning Temperature—**The switch generates an alarm if this temperature is reached. If the **Shutdown Temperature** is reached, once the temperature drops below the **Upper Warning Temperature**, the switch restarts.
- **Lower Warning Temperature**—The switch generates an alarm if this temperature is reached. Once the temperature rises above this temperature, the alarm is cleared.
- **Low Limit Temperature**—The switch generates an alarm if this temperature is reached.

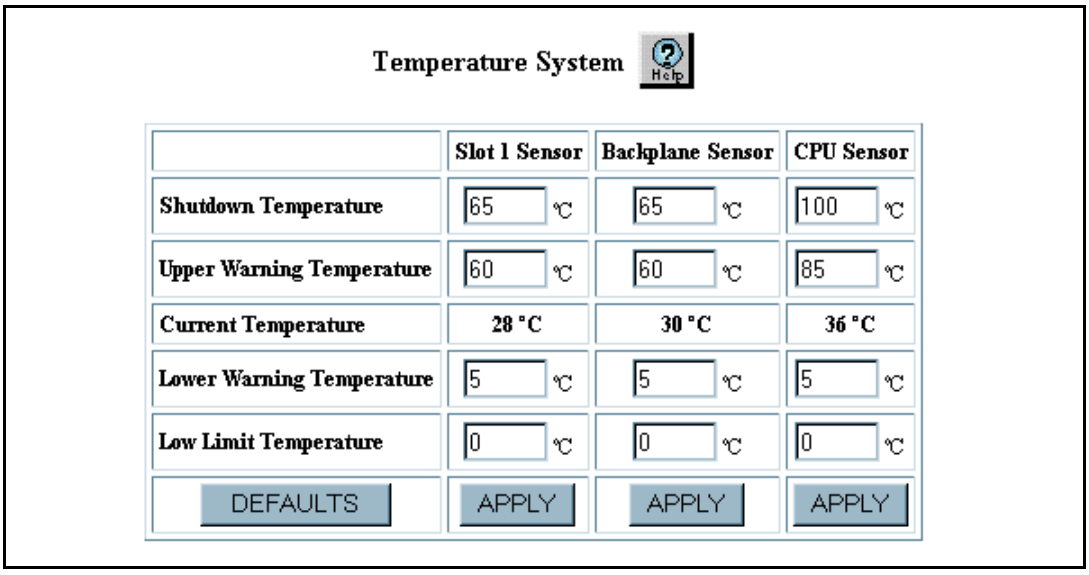
See Figure 3-7 for the default temperature system settings. You can set the temperature system using either the Web Agent or the CLI.

Web Agent Procedure

To configure the temperature warning systems:

1. In the navigation pane, expand the **System** > **Configuration** folders, and then click **Temperature System**. The Temperature System Web page is displayed in the content pane. See Figure 3-7.

Figure 3-7. Temperature System Web Page for P580



2. In the **Slot 1 Sensor** and **Backplane Sensor** fields, enter the desired temperature warnings.

*** Note:** The P882 does not display the current temperature of the backplane. In the **Current Temperature** field, the switch displays `--`.

*** Note:** If a redundant Supervisor module is installed, the **Slot 2 Sensor** column displays.

*** Note:** You cannot change settings for the backplane temperature sensors on the P882. If you attempt to change the settings, the switch displays the following error message:

```
Set Limits for Backplane Sensor failed -
failure writing to sensor
```

3. In the **CPU Sensor** fields, enter the desired temperature warnings.

4. Click **APPLY** to save your changes, or **DEFAULTS** to reset the fields to their default settings.

5. If you change the temperature ranges for the active supervisor, you must synchronize the active and standby supervisors to copy the temperature settings to the standby supervisor.

CLI Command

To check the temperature status, use the **show temperature** CLI command.

To set the temperature thresholds, use the following CLI commands:

- `(configure)#` **set temperature {supervisor-slot | backplane-sensor | cpu-sensor | probe} shutdown <*temperature*>**
- `(configure)#` **set temperature {supervisor-slot | backplane-sensor | cpu-sensor | probe} warning {upper | lower | low-limit} <*temperature*>**

Displaying the Power System Statistics

You can display the statistics for your switch's power system from the either the Web Agent and the CLI.

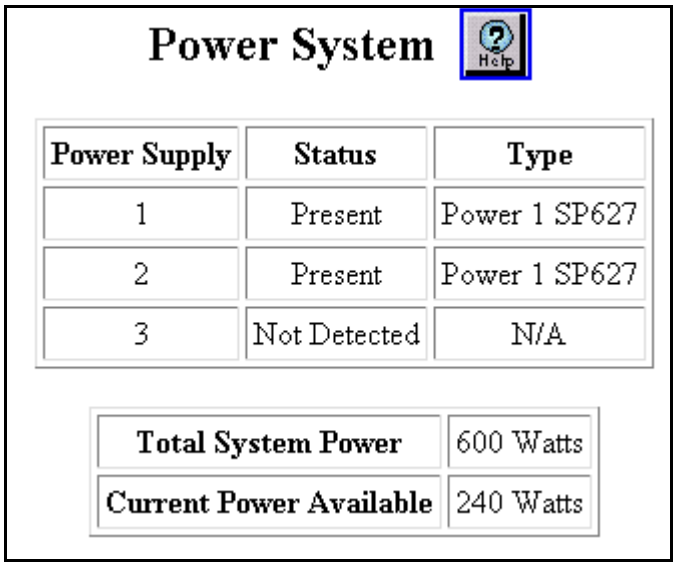
Web Agent Procedure

To display your switch's power system statistics using the Web Agent:

1. In the navigation pane, expand the **System > Configuration** folders, and then click **Power System**.

   The Power System Web page is displayed in the content pane and displays your switch's current power statistics. See Figure 3-8.

Figure 3-8. Power System Web Page



2. See Table 3-3 and review the Power System Web page parameters:

Table 3-3. Power System Web Page Parameters

Parameter	Definition
Power Supply	Identifies the power supply
Status	indicates whether the power supply is detected.
Type	Describes the type of power supply detected.
Total System Power	Displays the total system power in Watts.
Current Power Available	Displays the current power available.

*** Note:** The **Total System Power** field displays 1400 Watts if three power supplies are installed. The switch uses power from only

two of the three power supplies. The third power supply is a redundant power supply and is used only if one of the other power supplies fails.

CLI Command

To display your switch's power system statistics using the CLI, enter the following command from the User mode:

> **show system power**

Displaying Cooling System Statistics

You can display the statistics for your switch's cooling system from either the Web Agent or CLI.

Web Agent Procedure

To display your switch's cooling system statistics using the Web Agent:

1. In the navigation pane, expand the **System > Configuration** folders, and then click **Cooling System**.

   The Cooling System Web page is displayed in the content pane and displays the status of your switch's cooling system. See Figure 3-9 for the P580 switch, or Figure 3-10 for the P882 switch.

Figure 3-9. P580 Cooling System Status Web Page

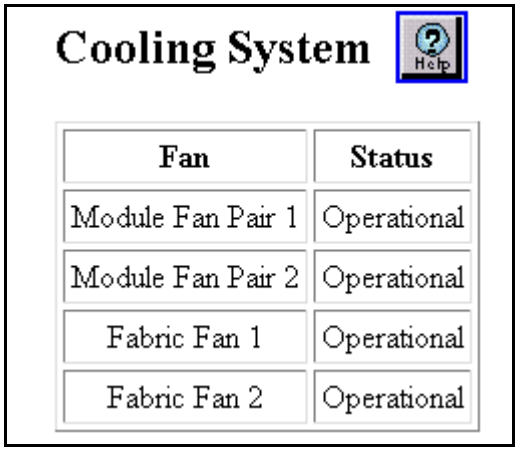
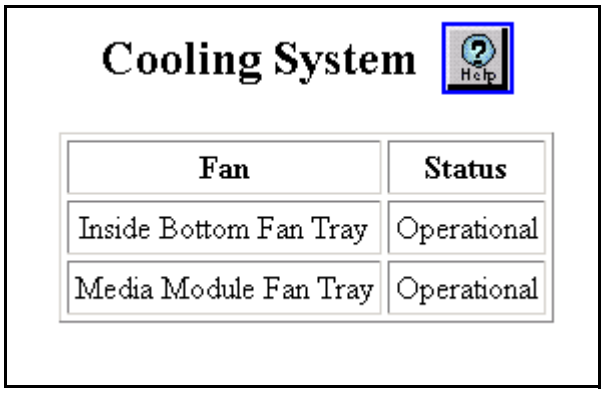


Figure 3-10. P882 Cooling System Status Web Page



2. Check the Status column to ensure that all the individual components are operational.

3. If a component's status is non-operational, power down the switch and contact a service representative to diagnose the failing unit.

CLI Command

To display your switch's power system statistics using the CLI, enter the following command from **Configuration** mode:

```
(configure)# show system fans
```

Performing a Reset

You can reset your switch from either the Web Agent or the CLI.

*** Note:** You must reset your switch after licensing an 80-Series modules.

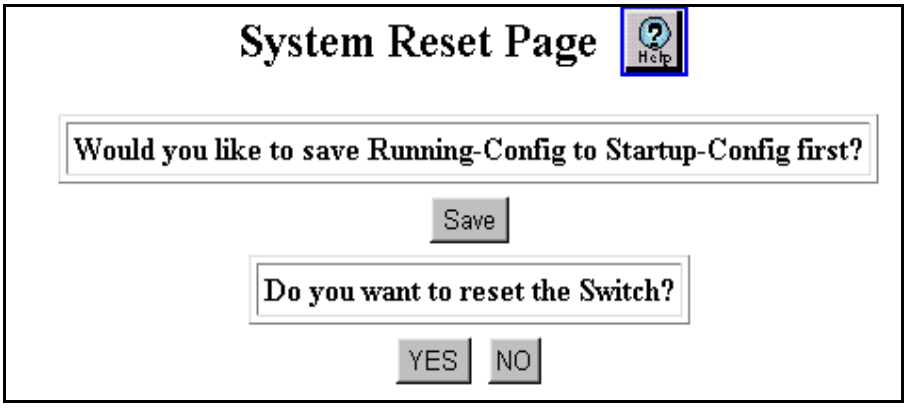
Web Agent Procedure

To reset your switch using the Web Agent:

*** Note:** You can also reset your switch using the **Module Reset** button located on the Supervisor module. See the Installation Sheet that came with the Supervisor Module for details.

1. In the navigation pane, expand the **System** folder, and then click **System Reset**. The System Reset Page Web page is displayed in the content pane. See Figure 3-11.

Figure 3-11. System Reset Page Web Page



2. Select **Save** to save your Running Configuration (Running-Config to the startup configuration (Startup-Config before performing a system reset.

   ***Note:** If you do not save your Running-Config to Startup-Config before you reset the switch, you will lose all of the modifications you made.

3. Click YES to reset the switch, or NO to cancel the operation.

CLI Command

To reset your switch using the CLI, enter the following command from Priv mode:

\# **reset**

4 Security

Overview

This chapter contains the following topics:

- Secure Mode
- SSHv2
- HTTPS Using SSLv3 or TLSv1
- RADIUS Client Support

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.

Secure Mode

Overview

***Important:** Avaya recommends that you enable secure mode on all switches that are running v6.0 and later application software.

Secure mode restricts management of the switch to the following secure protocols:

- HTTPS
- SSH
- SNMPv3

When secure mode is enabled:

- All non-secure protocols, such as Telnet, HTTP, and SNMPv1 and v2 are automatically disabled.
- You cannot use IP to manage the standby supervisor module. Any IP interfaces that you configured to access the redundant supervisor module are deleted.

Table 4-1 describes exactly what happens when you enable and disable secure mode.

Table 4-1. Secure Mode

When you enable secure mode...	When you disable secure mode...
SNMP v1 and v2 are automatically disabled.	SNMP v1 and v2 are automatically reenabled if SNMP v3 is enabled.
SNMPv3 remains in its current state (enabled or disabled).	SNMPv3 remains in its current state (enabled or disabled).
HTTP and Telnet are automatically disabled.	HTTP and Telnet remain disabled until you manually reenable them.
SSH and HTTPS remain in their current state (enabled or disabled). If you do *not* enable SSH before logging out, you can manage the switch only by using SNMP v3, if enabled, or the console port. The only way to disable secure mode at this point is by entering the **no secure mode** command at the console port. If you enable SSH before logging out, you can manage the switch (and disable secure mode) by using an SSH session.	SSH and HTTPS remain in their current state (enabled or disabled).
You cannot use Telnet or the Web Agent to manage the standby supervisor. Any IP interfaces that you configured to access the redundant supervisor module are deleted.	Any IP interfaces that were deleted when you enabled secure mode remain deleted.

This section contains the following information:

- Enabling Secure Mode
- Disabling Secure Mode
- Viewing the Secure Mode Setting

Enabling Secure Mode

To enable secure mode, use the following CLI command:

```
(configure)# secure-mode
```

Disabling Secure Mode

To disable secure mode, use the following CLI command:

```
(configure)# no secure-mode
```

Viewing the Secure Mode Setting

To view the secure mode setting, use the following CLI command:

```
> show secure-mode
```

SSHv2

Overview

Purpose of SSH

The Avaya Multiservice switch supports Secure Shell (SSH) version 2 for clients and servers. SSH is a protocol for secure remote login and other secure services. SSH provides a secure service that is similar to Telnet.

SSH is necessary because Telnet transmits unencrypted text TCP/IP packets that anyone on the same network can intercept. SSH encrypts the data being transmitted and allows for several methods of client/server and user authentication, connection integrity, and client/server verification. SSH runs on top of a TCP/IP connection.

SSH Server and Client

You can use the Avaya Multiservice switch as a server for SSH connections. The SSH server protocol relies on a public/private key pair that is generated on the server. The private key is kept on the server and cannot be viewed. The public key can be displayed and is used by remote clients to connect to the server. When a client tries to connect to the server, the server provides the public key to the client. Depending on the configuration of the client, it may use this key directly, or verify it against a locally stored copy.

The Avaya Multiservice switch can also be used as a client in SSH connections.

Encryption Ciphers

The Avaya Multiservice switch supports Blowfish and 3DES encryption ciphers. When the client connects to the host, the client supplies a list of ciphers that it supports. The server selects the strongest common cipher. You can also configure the Avaya Multiservice switch to force use of a single cipher.

User Authentication

The Avaya Multiservice switch supports only password authentication. Each login attempt requires a username and password authentication for logging onto the switch.You can also use a RADIUS server for remote password authentication over a network. User authentication occurs after an SSH session is successfully established.

The switch supports a maximum of seven client and server SSH sessions running simultaneously. All connections require password authentication.

Procedures

This section provides the following procedures:

- Enabling SSH on a TCP Port
- Disabling SSH
- Generating an SSH Server Key
- Displaying the Public SSH Key
- Configuring SSH Server
- Displaying SSH Connections
- Using the SSH Client

Enabling SSH on a TCP Port

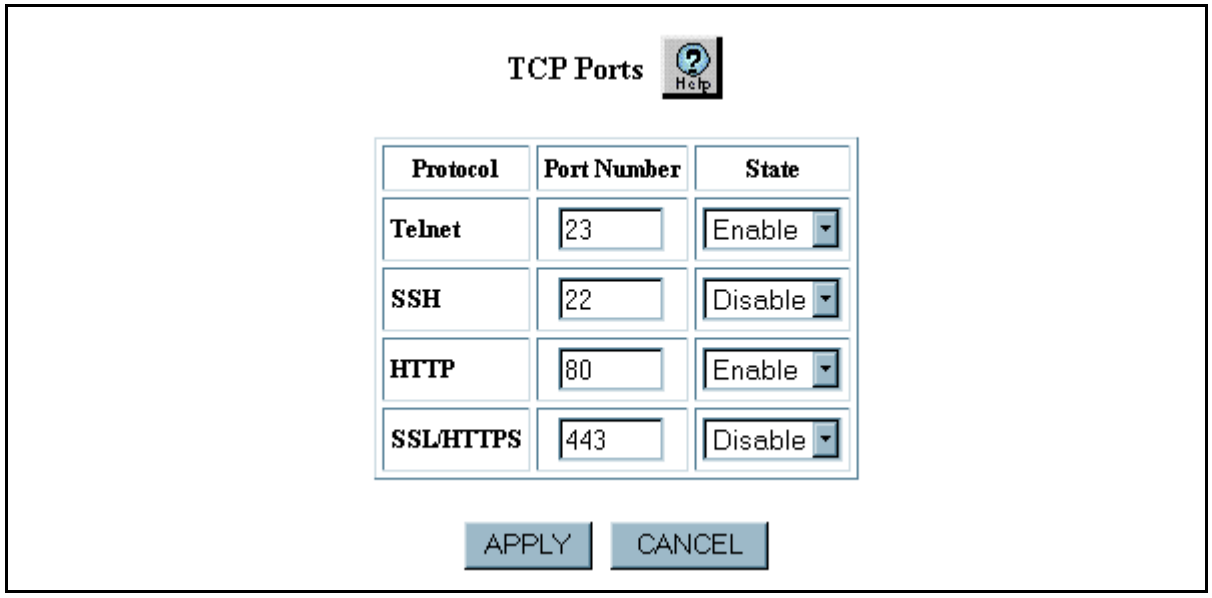
Before you can enable SSH, you must disable the Telnet port. You can, however, change the SSH port without disabling SSH. Valid SSH ports are 22 and 9000 to 65,535.

Web Agent Procedure

To enable SSH:

1. In the navigation pane, expand the **System** > **Administration** folders.
2. Click **TCP Ports**. The TCP Ports Web page is displayed in the content pane.

Figure 4-1. TCP Ports Web Page



3. In the **State** field for Telnet, select **Disable**.
4. In the **Port Number** field for SSH, enter the port number you want to use for SSH. Valid SSH ports are 22 and 9000 to 65,535. The default port for SSH is port 22.
5. In the **State** field for SSH, select **Enable**. SSH is disabled by default.
6. Click **Apply**. SSH is enabled on the specified port.

CLI Command

To enable SSH on a TCP port, use the following CLI command:

`(configure)#` **ip ssh {port [<*tcp-new-port*>] [enable] | [enable]}**

Disabling SSH

You must disable SSH before you enable Telnet.

Web Agent Procedure

To disable SSH:

1. In the navigation pane, expand the **System** > **Administration** folders.
2. Click **TCP Ports**. The TCP Ports Web page is displayed in the content pane. See Figure 4-1.
3. In the **State** field for SSH, select **Disable**.
4. Click **Apply**. SSH is disabled.

CLI Command

To disable SSH, use the following CLI command:

```
(configure)# no ip ssh
```

Generating an SSH Server Key

The Avaya Multiservice switch currently supports RSA and DSA key types. You can specify the key length of 768, 1024, or 2048 bytes. A key of length 1024 bytes provides more robust security and is the default key type. The key pair is saved to the local server.

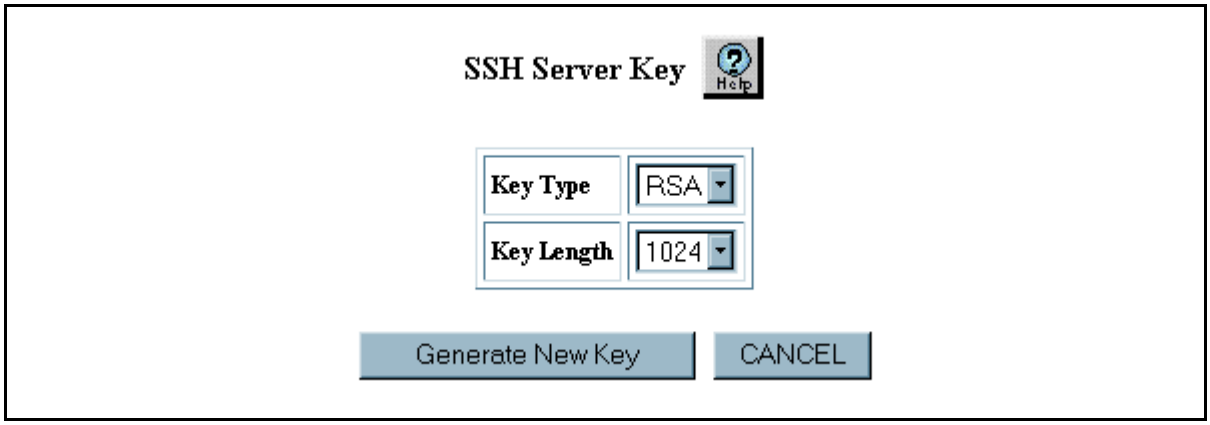
***Important:** If SSH is enabled and you regenerate the SSH server key, you must disable and then reenable SSH for the change to take effect. For information on how to disable and reenable SSH, see "Enabling SSH on a TCP Port."

Web Agent Procedure

To generate a server key:

1. In the navigation pane, expand the **System** > **Administration** > **Security** > **SSH** folders.
2. Click **Server Key**. The SSH Server Key Web page is displayed in the content pane. See Figure 4-2.

Figure 4-2. SSH Server Key Web Page



3. In the **Key Type** field, select RSA or DSA.
4. In the **Key Length** field, select 768, 1024, or 2048 bytes. The default is 1024 bytes.
5. Click **Generate New Key**.
6. Click **Apply**.

CLI Command

To generate SSH key pairs and save them on the local server, use the following CLI command:

`(configure)#` **ssh keygen [{rsa | dsa}] [key-size {768 | 1024 | 2048}]**

Displaying the Public SSH Key

The SSH Server Key is sent to the client in an SSH connection.

To display the public server key, use the following CLI command:

`(configure)#` **show ssh public-key**

This functionality is available only in the CLI. You cannot use the Web Agent to display the public SSH key.

Configuring SSH Server

Web Agent Procedure

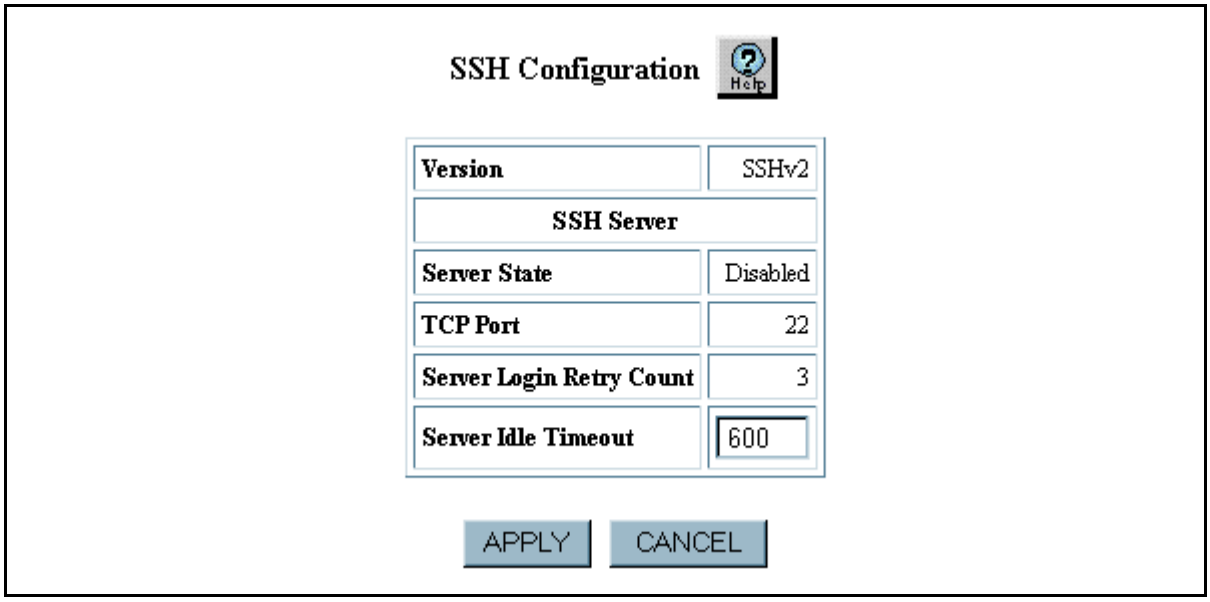
The SSH Configuration Web page displays the following:

- **Version**—The version of SSH running on the port. The Avaya Multiservice switch supports only SSH, version 2.
- **Server State**—The state of the server. Enabled or disabled.
- **TCP Port**—The TCP port on which SSH runs. Port 22 is the default port.

To configure SSH:

1. In the navigation pane, expand the **System** > **Administration** > **Security** > **SSH** folders.
2. Click **Configuration**. The SSH Configuration Web page is displayed in the content pane. See Figure 4-3.

Figure 4-3. SSH Configuration Web Page



3. In the **Server Idle Timeout** field, specify the timeout on a connection. Valid range is 0 to 1800 seconds. The default is 600 seconds.

4. Click **Apply**.

CLI Commands

To configure SSH, use the following CLI command:

- To set the server idle timeout, `(configure)#` **ssh timeout [<*seconds*>]**

Displaying SSH Connections

CLI Command

To display current SSH sessions, use the following CLI command:

```
# show ssh sessions
```

To end an SSH session, use the following CLI command:

```
(configure)# clear ssh <session-id>
```

This functionality is available only in the CLI. You cannot use the Web Agent to display all current SSH sessions.

Sample Output

Sample output of the **show ssh sessions** command is as follows:

```
SessionId User       RemoteIp:Port

---------- ---------- --------------

0          jsmith     10.10.6.100:1760

1          sjensen    10.10.8.110:1770

2          gschroeder 10.10.7.130:1771

3          tblair     10.10.6.100:1777
```

Using the SSH Client

You can use the SSH client to connect to an SSH server that is running on another machine.

To establish an SSH connection to a remote host, use the following CLI command:

```
(configure)# ssh [cipher {3des-cbc | blowfish-cbc}] [port <tcp-port>] [user <username>] {<ip-addr> | <hostname>}
```

If you do not specify a cipher, the client can use either cipher. Normally, if the remote host supports 3DES, that is the cipher that is used. If you do not specify a TCP port, the client uses port 22. Valid ports are 22 and 9000 to 65,535.

HTTPS Using SSLv3 or TLSv1

Overview

The Avaya P580 and P882 Multiservice switches support Secure Socket Layer, version 3 (SSLv3), and Transport Layer Security, version 1.0 (TLSv1).

SSL and TLS are protocols that provide data security between application protocols (such as HTTP, Telnet, NNTP, FTP) and TCP/IP. SSL and TLS are used to:

- Transmit encrypted data over TCP/IP networks, and
- Authenticate sites and clients with certificates

HTTPS is HyperText Transfer Protocol that is running either SSL or TLS. TLS is based on the SSLv3 protocol specification published by Netscape Communications Corporation. The Internet Engineering Task Force (IETF) defines TLS in RFC 2246 as the successor of SSL v3.

SSL and TLS use certificates and public and private keys to secure data. SSL server certificates prove the identity of the server to clients. The application software provides its own internal Certificate Authority (CA) for self-signing certificates. Certificates provide the following information:

- Name of the server's CA.
- Name of the entity to which the certificate was issued.
- Public key of the entity.
- Expiration date of the certificate.

Public and private keys are created from algorithms, called ciphers, that are used to encrypt and decrypt data. The public key is shared. The private key should never be shared. You cannot access the private key on the P580 or P882 switch.

Public-private key pairs work together: Data that is encrypted with the public key can be decrypted only with the private key and vice versa. The Avaya Multiservice switches use only RSA SSLv3 and TLSv1 cipher suites.

This section contains the following information and procedures:

- Supported Cipher Suites
- Viewing Cipher Suites
- Viewing the SSL Configuration
- Creating a Self-Signed SSL Server Certificate
- Viewing the Server Certificate
- Enabling SSL/HTTPS
- Restarting SSL/HTTPS
- Reverting to a Backup Certificate

Supported Cipher Suites

The P580 and P882 Multiservice switches support the following cipher suites:

- SSLv3 cipher suites:
  - SSL_RSA_WITH_DES_CBC_SHA
  - SSL_RSA_WITH_3DES_EDE_CBC_SHA
- TLSv1 cipher suites
  - TLS_RSA_WITH_DES_CBC_SHA
  - TLS_RSA_WITH_3DES_EDE_CBC_SHA

Viewing Cipher Suites

Web Agent Procedure

To view available cipher suites:

1. In the navigation pane, expand the **System** > **Administration** > **Security** > **SSL** folders.
2. Click **Ciphers**. The SSL Ciphers Web page is displayed in the content pane. This page lists the available cipher suites.

CLI Command

To view the available cipher suites, use the following CLI command:

> show ssl ciphers

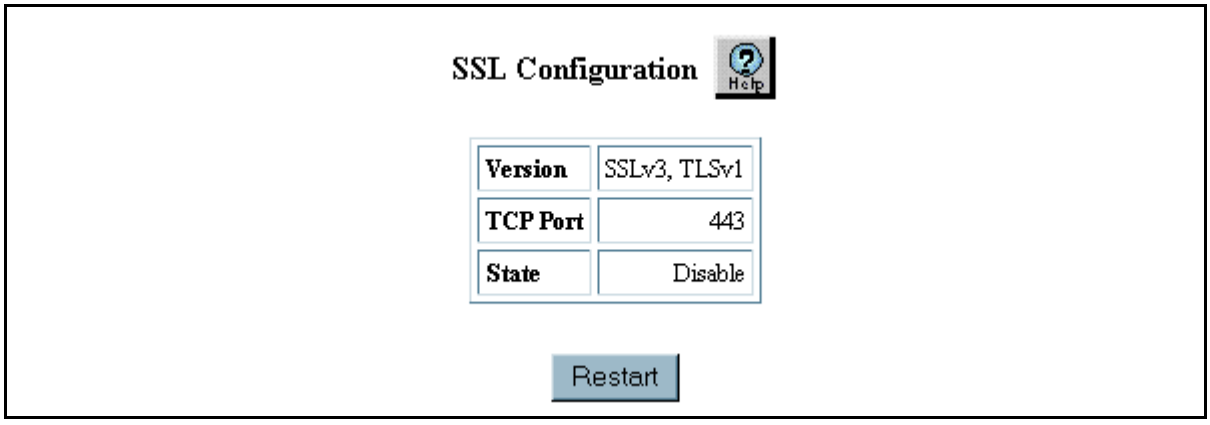
Viewing the SSL Configuration

Web Agent Procedure

To view the SSL configuration:

1. In the navigation pane, expand the **System** > **Administration** > **Security** > **SSL** folders.
2. Click **Configuration**. The SSL Configuration Web page is displayed in the content pane. See Figure 4-4.

Figure 4-4. SSL Configuration Web Page



CLI Command

To view the state of SSL (whether it is enabled or disabled), use the following CLI command:

> show ssl config

Creating a Self-Signed SSL Server Certificate

Overview

A *self-signed certificate* is a certificate for which the issuer is the same as the subject (the entity whose public key is being authenticated by the certificate). To create a self-signed certificate, you must first complete a *certificate signing request* (CSR) and then you can self-sign it. Once self-signed, the certificate is saved to a temporary file in the nonvolatile RAM (NVRAM). You must then restart SSL for the certificate to take effect. For information on how to restart SSL, see "Restarting SSL/HTTPS" later in this chapter.

To create your X.500 distinguished name, which is unique across the internet, you need the following information:

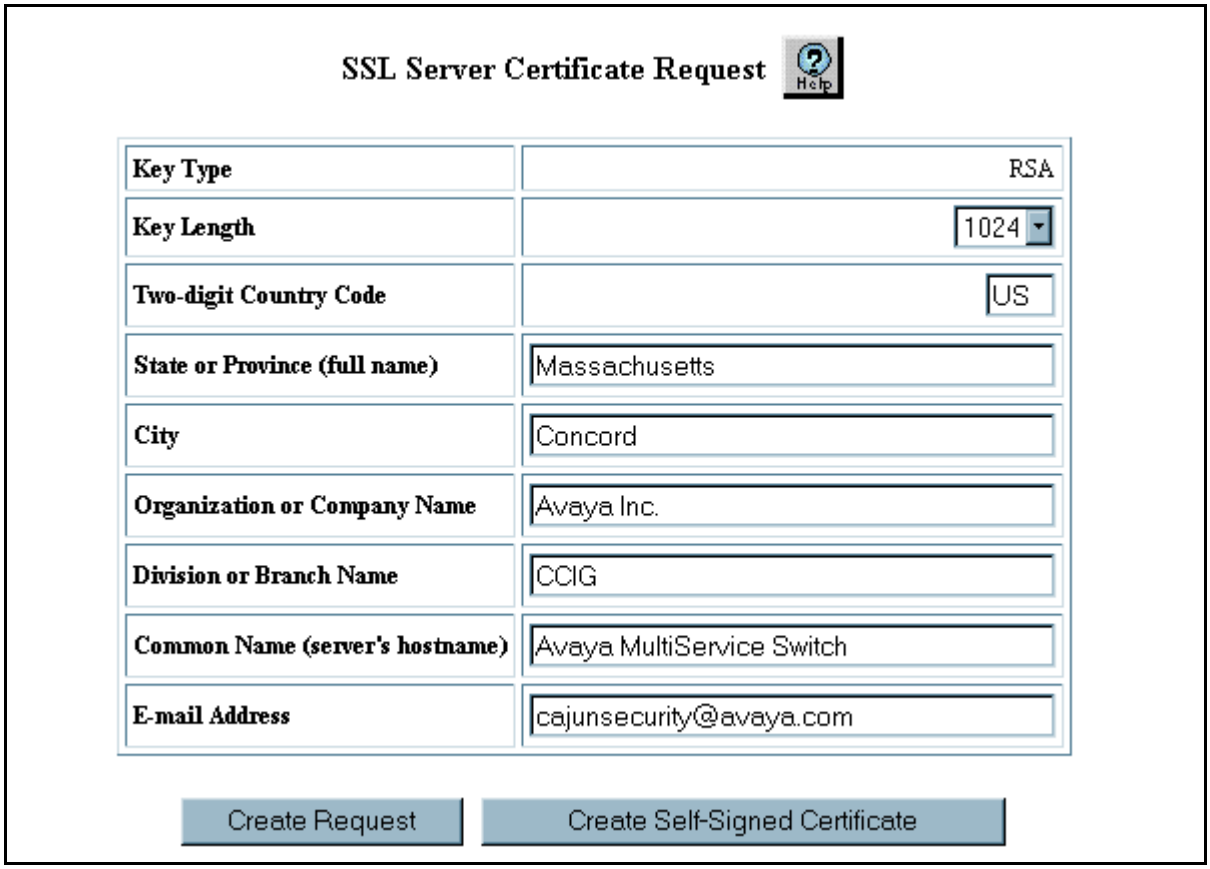
- Two-digit country code
- State or province (full name)
- City
- Organization or company name
- Division or branch name
- Common name (host name of the server)
- E-mail address

Web Agent Procedure

To create a self-signed certificate:

1. In the navigation pane, expand the **System** > **Administration** > **Security** > **SSL** folders.
2. Click **Certificates**. The SSL Server Certificate Web page is displayed in the content pane. See Figure 4-6.
3. Click **Certificate Request**. The SSL Server Certificate Request Web page is displayed in the content pane. See Figure 4-5.

Figure 4-5. SSL Server Certificate Request Web Page



4. Enter the appropriate information in the following fields:

   - **Key Type**—RSA only.
   - **Key Length**—Select the length of the key. Options are 512 or 1024 bits.
   - **Two-digit Country Code**—Enter your country code.
   - **State or Province (full name)**—Enter your state or province.
   - **City**—Enter the name of your city.
   - **Organization or Company Name**—Enter your organization or company name.
   - **Division or Branch Name**—Enter the name of your company's division or branch.
   - **Common Name (server's hostname)**—Enter a common name for the server. Use either the fully qualified domain name (FQDN) or the IP address of the switch.
   - **Email Address**—Your e-mail address.

5. Click **Create Self-Signed Certificate**.

6. Restart SSL for the new certificate to take effect. For information on how to restart SSL, see "Restarting SSL/HTTPS" later in this chapter.

CLI Command

To create a self-signed certificate, use the following commands:

- To create a CSR, `(configure)#` **ssl certreq [{512 | 1024}]**
- To self-sign the CSR, `(configure)#` **ssl selfcert**

Viewing the Server Certificate

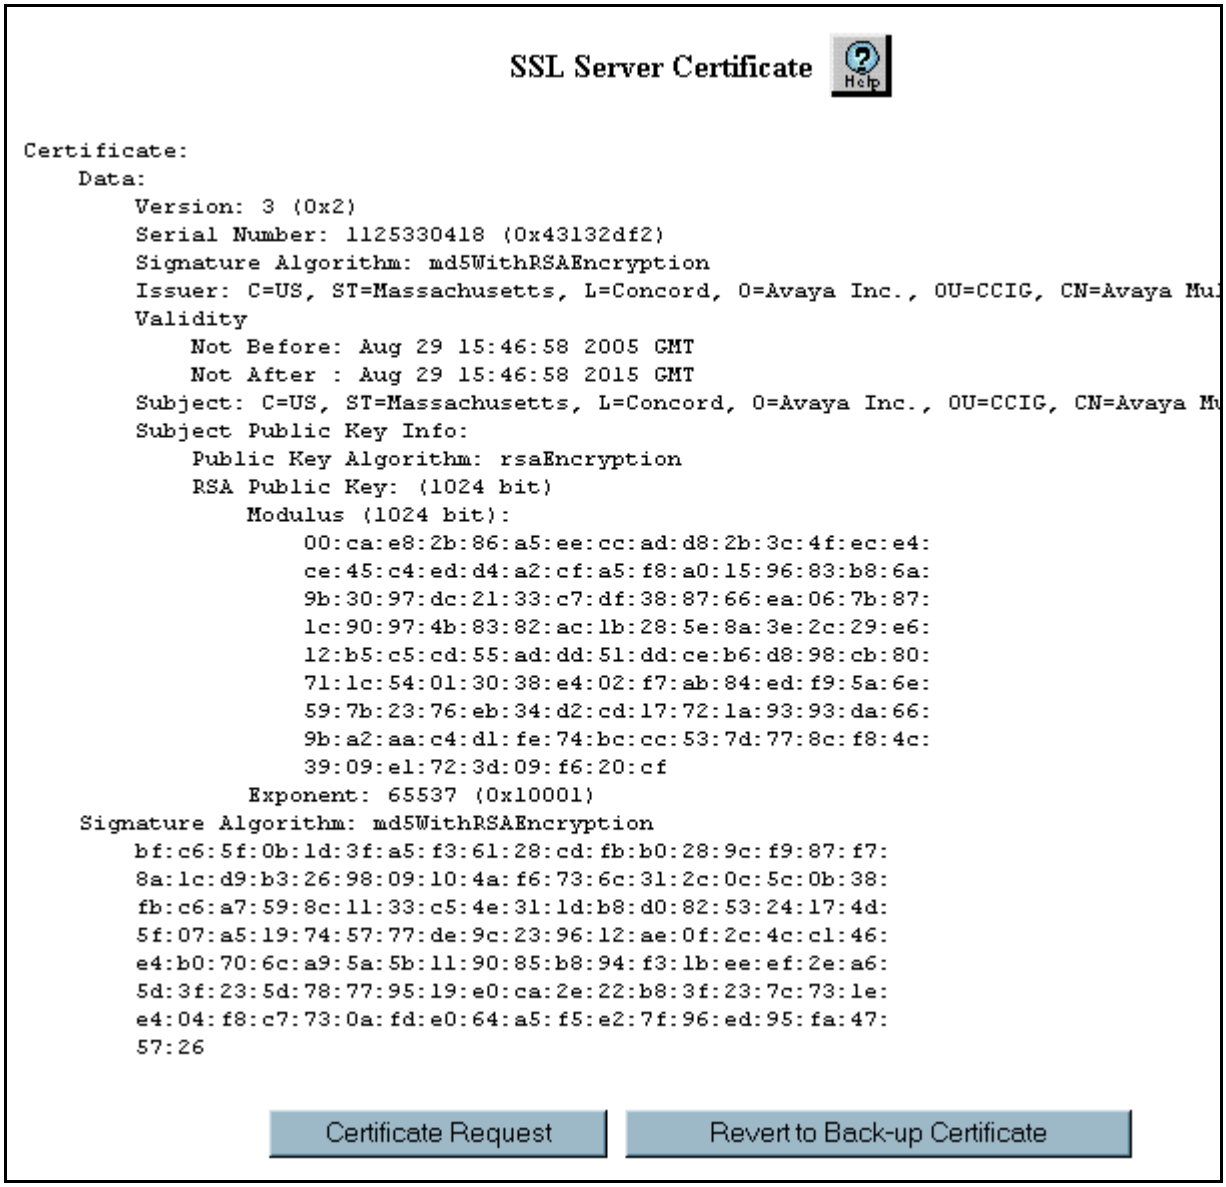
Web Agent Procedure

To view the SSL server certificate:

1. In the navigation pane, expand the **System** > **Administration** > **Security** > **SSL** folders.

2. Click **Certificates**. The SSL Server Certificate Web page is displayed in the content pane. See Figure 4-6.

Figure 4-6. SSL Server Certificate Web Page



CLI Command

To view the SSL certificate, use the following CLI command:

> show ssl cert

Enabling SSL/HTTPS

Before you can enable SSL/HTTPS on the switch, you must disable HTTP.

Web Agent Procedure

To enable SSL/HTTPS:

1. In the navigation pane, expand the **System** > **Administration** folders.
2. Click **TCP Ports**. The TCP Ports Web page is displayed in the content pane. See Figure 4-1.
3. In the **State** field for HTTP, select **Disable**.
4. In the **Port Number** field for SSL/HTTPS, enter the port number that you want SSL/HTTPS to use. Valid ports are 443 or 9000 to 65,535. The default port for SSL/HTTPS is 443.
5. In the **State** field for SSL/HTTPS, select **Enable**. SSL/HTTPS is disabled by default.
6. Click **Apply**. SSL/HTTPS is enabled on the specified port.

CLI Command

To enable SSL/HTTPS, use the following CLI command:

`(configure)#` **ip https {port [<*tcp-new-port*>] [enable] | [enable]}**

Disabling SSL/HTTPS

You must disable SSL/HTTPS before you enable HTTP.

Web Agent Procedure

To disable SSL/HTTPS:

1. In the navigation pane, expand the **System** > **Administration** folders.
2. Click **TCP Ports**. The TCP Ports Web page is displayed in the content pane. See Figure 4-1.
3. In the **State** field for SSL/HTTPS, select **Disable**.
4. Click **Apply**. SSL/HTTPS is disabled.

CLI Command

To disable SSL/HTTPS, use the following CLI command:

`(configure)#` **no ip https**

Restarting SSL/HTTPS

You must restart SSL after updating the certificate information.

Web Agent Procedure

To restart SSL:

1. In the navigation pane, expand the **System** > **Administration** > **Security** > **SSL** folders.
2. Click **Configuration**. The SSL Configuration Web page is displayed in the content pane. See Figure 4-4.
3. Click **Restart**.

CLI Command

To restart SSL/HTTPS, use the following CLI command:

```
(configure)# ssl restart
```

Reverting to a Backup Certificate

Overview

You can revert back to a backup version of the SSL server certificate. If you revert to a backup certificate, the current certificate is renamed and made the backup for later reuse.

Web Agent Procedure

To revert to a backup certificate:

1. In the navigation pane, expand the **System** > **Administration** > **Security** > **SSL** folders.
2. Click **Certificates**. The SSL Server Certificate Web page is displayed in the content pane. See Figure 4-6.
3. Click **Revert to Backup Certificate**. The backup certificate becomes the current certificate.

CLI Command

To revert to a backup certificate, use the following CLI command:

```
(configure)# ssl backcert
```

RADIUS Client Support

Overview

Purpose of RADIUS

In a network with many Avaya switches, configuring user accounts on each of the switches can be time-consuming.You can centralize the user accounts by using a Remote Authentication Dial-In User Service (RADIUS) server.

RADIUS is a service that authenticates users when they attempt to log in to a Network Access Device (NAD) such as an Avaya switch. RADIUS typically runs on a Windows or Linux server; however, it can run on other platforms as well depending on the vendor.

*** Note:** RADIUS supports a maximum of 27 characters for user names. If you use a RADIUS server to authenticate users, their switch user names must not exceed 27 characters, regardless of the 31-character maximum of the P580 and P882.

Authentication Process

RADIUS is a client/server architecture where each device that uses the RADIUS server is a RADIUS client. The client sends Access-Request messages to the RADIUS server. These messages include the user name, the password encrypted, and optional parameters depending on configuration.

***Important:** The RADIUS Client and Server *must* be configured with the exact same parameters.

Once the RADIUS server receives the Access-Request message, it searches its database for the user account. If the server finds the account, the password is correct, and the optional parameters match, the server sends an Access-Accept message to the RADIUS client. The Access-Accept message indicates that the user account exists, the password is correct, and the user has a certain access type (for example, administrative or read-only). If the RADIUS server does not find the account or the password is incorrect, then the server sends an Access-Reject message to the RADIUS client.

*** Note:** Due to an interoperability issue, the P580 and P882 RADIUS client does not accept Access-Accept messages from Windows 2000 RADIUS servers, which generate the Generate-Class-Attribute. To resolve this issue, obtain Windows 2000 service pack 3 or later. After installing the latest service pack, set the **Generate-Class-Attribute** field to **FALSE**.

This interoperability issue occurs because Microsoft RADIUS server includes a class attribute in Access-Accept messages that the P580 and P882 RADIUS client does not support. With service pack 3, you can disable generation of a class attribute. For more information on this issue, see http://support.microsoft.com/default.aspx?scid=kb;EN-US;Q297317.

Section Contents

This section contains the following topics:

- Realms and Groups
- Login Order of Operations
- Avaya VSAs
- RADIUS Server Files
- Configuring a RADIUS Client

Realms and Groups

Overview

Realms and groups provide two separate functions. A realm provides a way of organizing user accounts on the RADIUS server. Groups provide a way of organizing NADs that a user can log in to as well as delivering vendor-specific parameters that you configure.

For example, you might use a realm called *AvayaRealm* to organize all user accounts that can log into Avaya switches in a campus environment. In this campus, you organize network administrators in to two teams, one team for the north campus and one for the south campus. Each team needs read-write access to switches in their half of the campus and read-only access to switches in the other half of the campus.

You would then assign all of the north switches to a group named *NorthSwitches* and the south switches to a group named *SouthSwitches*.

For each user, you would create two user accounts in the AvayaRealm: one with a group name of *NorthSwitches* and one with *SouthSwitches*. Each account would have the appropriate permissions for the two switch types.

When a user from the north team logs into a switch in the north campus, the switch sends an Access-Request message with *@AvayaRealm* appended to the user name and a group name of *NorthSwitches*. The RADIUS server will send an Access-Accept message indicating that the user has read-write permission.

Similarly, when the same user logs in to a switch on the South campus, the message will append *@AvayaRealm* and a group name of *SouthSwitches*. The RADIUS server will send an Access-Accept message indicating that the user has read-only permission.

Realms

A realm provides a mechanism by which a RADIUS manager can organize user accounts. Consult the RADIUS vendor documentation for information on how to create realms on the server. Once created, user accounts are placed in the realms. The realm name is also configured on the NADs and when the NADs send Access-Request messages, the user name is appended with an ampersand (@) and the realm name.

For example: User *Bob* in *AvayaRealm* logs in to the switch as *Bob*. The Avaya switch sends an Access-Request message for user *Bob@AvayaRealm*. The RADIUS server, upon receiving the request, searches for *Bob* in the *AvayaRealm*.

Groups and VSAs

To provide user accounts the same granularity of privileges that local authentication provides, you can configure vendor-specific attributes (VSAs) on the RADIUS server and a group name on the switch. After you set the group name, the switch includes it in Access-Request messages that it sends to the RADIUS server.

If the user name, password, and group name match that of the user account, the RADIUS server sends an Access-Accept message to the client. VSAs that identify the privileges the user has are included in the Access-Accept message.

*** Note:** If a user has a RADIUS account that does not contain a group name, the RADIUS server still responds with an Access-Accept message; but the message does not contain a group name or VSAs. This absence of a group name presents a potential security risk. For more information, see “Configuring a RADIUS Client” later in this chapter.

Login Order of Operations

When a user attempts to log in to the Avaya switch, the switch first checks the local user accounts for the user name and password. If found, the user is logged in using the local settings for that account.

If no local account is found and RADIUS is enabled and configured, the switch sends an Access-Request message to the primary RADIUS server in an attempt to authenticate the user remotely. If the user login is found and correct, then the RADIUS server responds with an Access-Accept message that includes the user privileges. If the user account has the appropriate management type (for example, Web if he or she is trying to log in to the Web Agent), the user is granted access.

If the user login is incorrect or does not exist, then the RADIUS server sends an Access-Reject message to the switch and the user is denied access to the switch.

If the primary RADIUS server does not respond to the Access-Request message, the switch attempts to use the secondary server if it is configured. On subsequent retries, the switch alternates between the primary and secondary servers. If no secondary server is configured, the switch continues to try the primary server.

The switch waits the number of seconds specified in the retry interval and tries as many times as specified in the retry time.

Avaya VSAs

Switch Service Types

The switch service type VSA (attribute 1) specifies the access permission that a user has. The following types are supported:

- Read-only (can view the Avaya switch configuration). Setting of **1**.
- Read-write (Can configure the Avaya switch. To have read-write service type, the user account must be assigned to a group.) Setting of **2**.
- Administrative (Can create user accounts and configure the Avaya switch). Setting of **3**.
- Custom access type. For information on custom access types, see Chapter 2, "Setting Up the Switch." Setting of **4**.

Management Types

The management type VSA (attribute 2) specifies the method that a user can use to manage the switch. The following four types are supported:

- All. Setting of **1**.
- Local CLI (serial port on the supervisor). Setting of **2**.
- Remote CLI (Telnet or SSH session). Setting of **3**.
- Web Agent. Setting of **4**.

CAT Access Mask

The CAT access mask specifies the accessible features for a custom access type.

The mask is a 32-bit hexadecimal value. Each feature is assigned to a unique bit in the mask. To control the accessible features for a user, you set the appropriate bits. Table 4-2 shows the bit position of each feature.

Table 4-2. Bit Positions of CAT Access Mask

Feature	Bit Position	Binary Position	Hexadecimal Position
System Configuration	Bit 12	0000 0000 0000 0000 0000 1000 0000 0000	00000800
Modules and Ports	Bit 16	0000 0000 0000 0000 1000 0000 0000 0000	00008000
Event Management	Bit 20	0000 0000 0000 1000 0000 0000 0000 0000	00080000
Layer 2 Switching	Bit 24	0000 0000 1000 0000 0000 0000 0000 0000	00800000
Routing	Bit 28	0000 1000 0000 0000 0000 0000 0000 0000	08000000

For example, to give a user access to modules and ports, event management, layer 2 switching, and routing, you would enter a hexadecimal value of 08888000 (0000 1000 1000 1000 1000 0000 0000 0000 binary).

*** Note:** Only bits 12, 16, 20, 24, and 28 are used in the CAT access mask. Set all other bits to 0. If you inadvertently set other bits in the mask, the switch ignores them. The remaining bits in the mask are reserved for future use.

If you enter a mask of all zeros (00000000), the user can view general system information but cannot configure any settings.

CAT Read-Only Mask

The CAT read-only mask specifies the read-only accessible features for a custom access type.

The mask is a 32-bit hexadecimal value. To specify read-only access for a particular feature, you set the appropriate bit. Each feature is assigned to the same bit as the access mask. See Table 4-2.

The corresponding bit in the CAT access mask must also be set. If the access mask bit is set, but the same bit in the read-only mask is set to 0, the user has read-write access to the particular feature.

For example, to give a user read-write access to modules and ports and event management, and read-only access to layer 2 switching and routing, you would enter the following values:

- CAT access mask—08888000 hexadecimal (0000 1000 1000 1000 1000 0000 0000 0000 binary).
- CAT read-only mask—08800000 hexadecimal (0000 1000 1000 0000 0000 0000 0000 0000 binary).

*** Note:** Only bits 12, 16, 20, 24, and 28 are used in the read-only mask. Set all other bits to 0. If you inadvertently set other bits in the mask, the switch ignores them. The remaining bits in the mask are reserved for future use.

RADIUS Server Files

Each RADIUS vendor may have a different method for configuring client and user files. For information on configuring the client and user files on your RADIUS server, see the RADIUS server documentation. The following tables list the Avaya-specific attributes and values for each attribute.

Table 4-3. Avaya-Specific Attributes

Attribute	Attribute Number	Value Description	Value Number
Switch Service Type	1	Read-Only	1
		Read-Write	2
		Administrative	3
		Custom Access Type-User	4
Management Type	2	All	1
		Local CLI	2
		Remote CLI	3
		Web	4
Group	3	*<group name>*	Any alpha-numeric string, 22 characters or less.
CAT Access Mask		Bit Mask	For custom access types. See "CAT Access Mask" in the previous section for information on setting this value.
CAT Read-Only Mask		Bit Mask	For custom access types. See "CAT Read-Only Mask" in the previous section for information on setting this value.

Sample User File

The following is a sample user file. Each RADIUS vendor may have a different method for configuring user files. In this sample, text values are defined that represent the actual values used by the Avaya switch using the directives ATTRIBUTE and VALUE.

Two user accounts are shown: a use account that is not assigned to a group and a user account that is assigned to a group. The first account is *Bob* with password *BooBoo*. Bob is allowed to log in only to the NAD at 199.87.201.2. Bob is also granted administrative privileges.

The second account is *Ann* with password *Pokey*. The NAD she logs in to must be configured to use group *AvayaSwitches*. She is granted read-only permission to the CLI (either Telnet or serial cable to the supervisor console port).

```
# define the Avaya Vendor Specific Attributes
ATTRIBUTE Avaya-Service-Type
ATTRIBUTE Avaya-Mgt-Type
ATTRIBUTE Avaya-Realm
ATTRIBUTE Avaya-Group
# Note: NAS-IP-Address is a Standard RADIUS Attribute

# define the Avaya-Service-Types
VALUE Avaya-Service-Type Avaya-Administrative 3
VALUE Avaya-Service-Type Avaya-Read-Write 2
VALUE Avaya-Service-Type Avaya-Read-Only 1

# define the Avaya Management Types
VALUE Avaya-Mgt-Type Avaya-Mgt-All 1
VALUE Avaya-Mgt-Type Avaya-Console-CLI 2
VALUE Avaya-Mgt-Type Avaya-Remote-CLI 3
VALUE Avaya-Mgt-Type Avaya-Web 4

# define User Accounts
Bob   Password = "BooBoo", NAS-IP-Address = "199.87.201.2"
      Service-Type = Administrative
Ann   Password = "Pokey", Avaya-Group = "AvayaSwitches"
      Avaya-Service-Type = Avaya-Read-Only
      Avaya-Management-Type = Avaya-Local-CLI
      Avaya-Management-Type = Avaya-Remote-CLI
```

Sample Client File

The following is a sample client file. Client files hold the IP address(es) of the NADs and their associated Shared Secrets. Client files may vary from vendor to vendor. Consult the vendor documentation on how to configure Client files.

```
#Client Name                    Shared Secret
#----------------------------   ----------------------
199.87.201.2                    W3ftrFF4
10.30.44.1                      Gruuf66
```

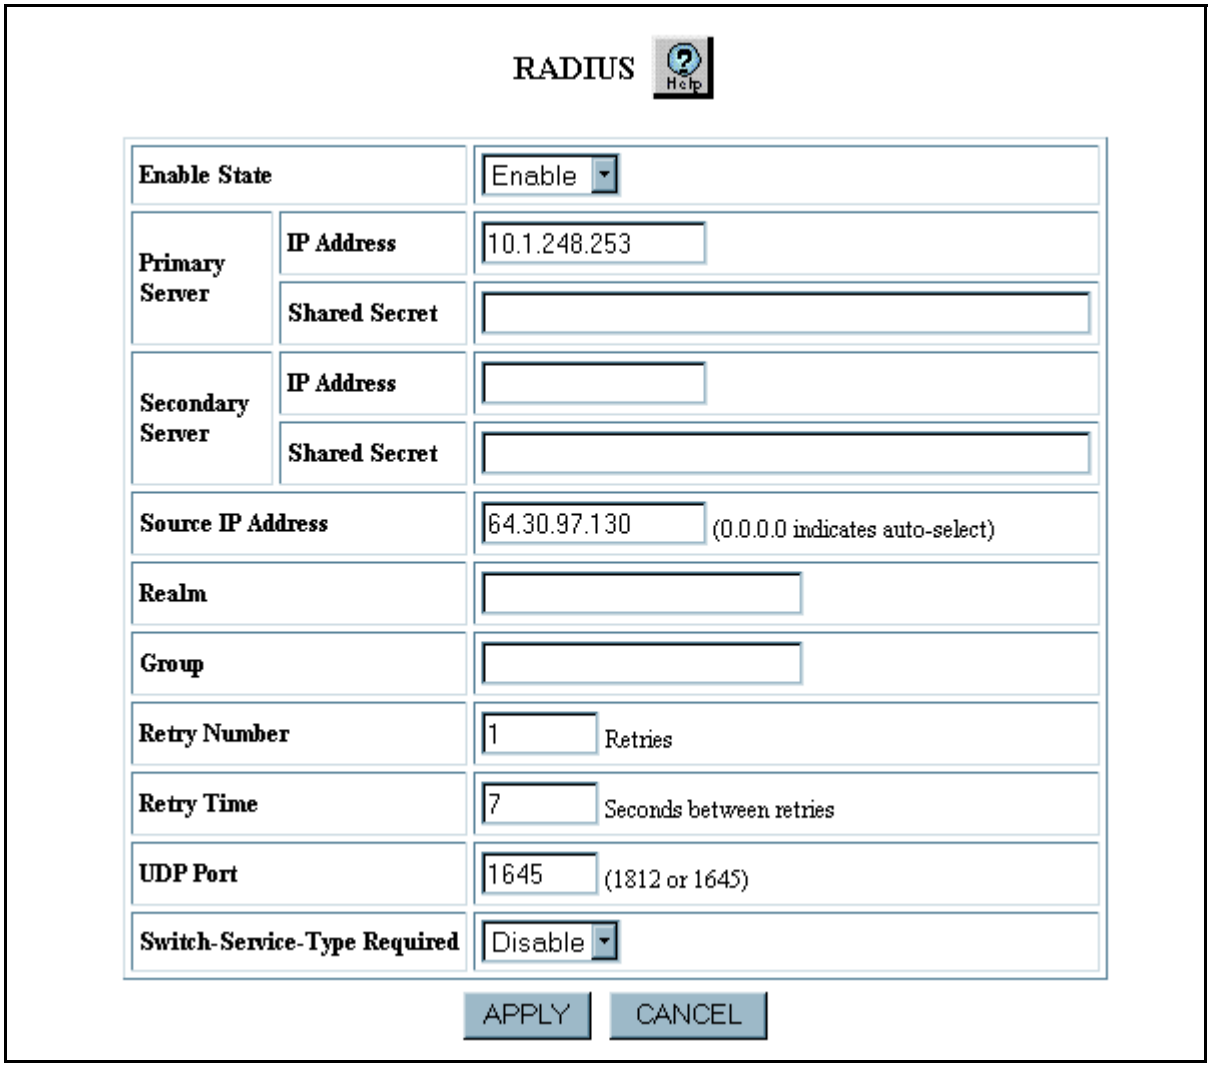
Configuring a RADIUS Client

Web Agent Procedure

To configure a RADIUS client:

1. In the navigation pane, expand the **System** > **Administration** folders, and then click **RADIUS**. The **RADIUS** Web page is displayed in the content pane (Figure 4-7).

Figure 4-7. RADIUS Web Page



2. In the **Enable State** field, select **Enable**.

3. Configure the RADIUS client as appropriate. Table 4-4 provides explanations of each field.

Table 4-4. RADIUS Web Page Configuration Parameters

Parameter	Definition
Enable State	Enable or disable RADIUS on the switch.
Primary Server	• **IP Address** - Enter the IP address for the primary RADIUS server. • **Shared Secret** - Enter the shared secret the switch will use for encrypting and decrypting passwords. Make sure the primary server is configured with the exact same characters (case sensitive). This value is itself encrypted and will not be displayed anywhere (Web Agent or CLI) once set. It can be changed by simply entering in a new shared secret.
Secondary Server	• **IP Address** - Enter IP address for the secondary RADIUS server. • **Shared Secret** - Enter the shared secret the switch will use for encrypting and decrypting passwords. Make sure the secondary server is configured with the exact same characters (case sensitive). This value is itself encrypted and will not be displayed anywhere (Web Agent or CLI) once set. It can be changed by simply entering in a new shared secret.
Source IP Address	Enter an IP interface address the switch will use as the source IP address in the Access-Request messages. This value must be an IP interface address on the switch. If set, and the IP interface becomes disabled, RADIUS will not function because the switch will not be able to send or receive RADIUS messages. If left 0.0.0.0 (the default), the switch automatically selects a source IP address from one of its active interfaces. If you use this setting, you must add each of the switch IP addresses to the Client file on the RADIUS server since you are not manually setting the source IP address.
Realm	Set this parameter only if realms are used on the RADIUS server for organizing user accounts. If so, enter the realm name for the user accounts that are authorized to log in to the Avaya switch. All user accounts that are authorized to log in to this switch must be assigned to the same realm.
Group	Enter the group name.The group name will be included in the Access-Request message sent to the RADIUS server. If you specify a group name, all user accounts must be assigned a group name on the RADIUS server and VSAs must be set for the user accounts.
	1 of 2

Table 4-4. RADIUS Web Page Configuration Parameters

Parameter	Definition
Retry Number	Enter the number of times to resend the Access-Request message if the RADIUS server does not respond.
Retry Time	Enter the time (in seconds) to wait before resending an Access-Request message.
UDP Port	Enter the UDP port number that you want the switch to use for RADIUS authentication. The default value is **1812**. Valid options are 1812 or 1645 only.
Switch-Service-Type Required	If this setting is *enabled*, the switch recognizes only Access-Accept messages that have the correct switch service type VSA (attribute 1). If this setting is *disabled*, the switch recognizes both the switch service type VSA and the service type standard radius attribute (attribute 6). For the standard radius attribute, the switch recognizes only two values: • Administrative (value 6) • NAS-Prompt (7),which the switch recognizes as read-only access.
	2 of 2

CLI Command

Use the following CLI commands to configure the RADIUS client on the switch:

- To enable or disable RADIUS client, `(configure)#` **set radius authentication [{enabled | disabled}]**
- To assign the switch to a group, `(configure)#` **set radius authentication group <*group*>**
- To set the user account realm that is authorized to log in to this switch, `(configure)#` **set radius authentication realm <*realm*>**.
- To set the maximum number of Access-Request messages to send if the server does not reply, `(configure)#` **set radius authentication retry-number <*retry-number*>**
- To set the time to wait before attempting to reauthenticate a login, `(configure)#` **set radius authentication retry-number <*retry-time-in-seconds*>**
- To set the primary or secondary RADIUS server, `(configure)#` **set radius authentication server <*ip-addr*> <*shared-secret*> [encrypted-type1] [{primary | secondary}]**

- To set the IP address used as the source IP address for Access-Request messages, `(configure)#` **set radius authentication source-ip <*ip-addr*>**
- To set whether the switch recognizes only Access-Accept messages that have the correct group name included, `(configure)#` **set radius authentication switch-service-type-required [{enabled | disabled}]**
- To set the UDP port number to use for RADIUS authentication, `(configure)#` **set radius authentication udp-port <*1812-or-1645*>**

5 Configuring SNMP

Overview

Supported versions of SNMP

The P580 and P882 support SNMPv1, v2, and v3.

SNMPv3 Security Features

SNMP version 3 provides the following security features that SNMPv1 and v2 do not provide:

- Encryption of protocol data units (PDUs) to prevent unauthorized users from viewing the PDU contents. SNMPv3 uses CBC-DES for its encryption protocol.
- Authentication of the user who sent the PDU. User authentication is provided by either the HMAC-SHA or HMAC-MD5 authentication protocol.
- Timeliness checks of the PDU to ensure that it has not been delayed or replayed.
- Ability to define which MIB objects and table rows that specific users can access and whether they have read-only, read-write, or notify access. This functionality is achieved by creating *views*, associating user *groups* with views, and then assigning users to groups.

*** Note:** For security reasons, you no longer can view the snmp configuration in the startup configuration file.

Contents

This chapter contains the following sections:

- Authentication and Encryption
- Timeliness Checks
- Views
- User Groups
- ATM Uplink Module
- Initial SNMPv3 User and SNMPv3 Administrator
- Changing the Engine ID of the Switch

- Creating an SNMPv3 Administrator
- Configuring Views
- Configuring Groups
- Configuring an SNMPv3 User
- Changing a User Password
- Configuring SNMPv1 or v2 Community Strings
- Setting the Administrative Contact
- Setting the Physical Location of the Switch
- Disabling or Reenabling SNMP
- Viewing the SNMP Status

***Important:** You must have administrator access to the switch to use the CLI commands that are discussed in this chapter.

RFC Standards

For more information on SNMPv3, see:

- RFC 2576, “Coexistence between Version 1, Version 2, and Version 3 of the Internet-Standard Network Management Framework.”
- RFC 3411, STD 62, “An Architecture for Describing SNMP Management Frameworks.”
- RFC3412, STD 62 “Message Processing and Dispatching for the Simple Network Management Protocol.”
- RFC 3413, STD 62, “SNMP Applications.”
- RFC 3414, STD 62, “User-Based Security Model (USM) for Version 3 of the Simple Network Management Protocol (SNMPv3).”
- RFC 3415, STD 62, “View-Based Access Control (VACM) for the Simple Network Management Protocol (SNMP).”
- RFC3416, STD 62 “Version 2 of the Protocol Operations for the Simple Network Management Protocol.”
- RFC3417, STD 62, “Transport Mappings for the Simple Network Management Protocol.”
- RFC3418, STD 62, “Management Information Base (MIB) for the Simple Network Management Protocol.”

Authentication and Encryption

Localized Keys

To perform authentication and encryption, the switch and NMS (network management system) share *localized keys*. When sending a PDU to the switch, the NMS (network management system) generates the localized key and places it in the PDU. When the switch receives the PDU, it compares the localized key in the PDU to the localized key stored in the switch memory. If the two versions match, the PDU is authenticated or decrypted.

To generate a localized key, the switch and NMS use HMAC-MD5 or HMAC-SHA to:

1. Hash the user password. The hashed user password is called the *non-localized key.*
2. Hash a combination of the non-localized key and the engine ID of the switch. This hashed combination is the *localized key.*

The NMS stores the non-localized key and generates the localized key only before sending a PDU to the switch. Each time you create a new SNMP user, the switch generates and stores the localized key for that user.

If authentication is enabled for a user, he or she must have an authentication password. And if encryption is enabled for a user, he or she must have an encryption password. For information on setting these passwords, see "Configuring an SNMPv3 User."

Engine ID

To perform authentication or encryption, the switch must have an engine ID. By default the engine ID is based on the IP address of the ethernet console port. You can, however, change the engine ID of the switch. For information on how to change the engine ID of the switch, see "Changing the Engine ID of the Switch."

If the switch is using the default engine ID and you change the IP address of the ethernet console port, the engine ID is also changed. All user accounts are invalid if the engine ID changes, and you must reconfigure them.

Timeliness Checks

The switch performs timeliness checks to ensure that PDUs are not reordered, delayed or replayed. These timeliness checks include the following information:

- Number of reboots since the switch was upgraded from v5.*x* application software to v6.*x* or since the last initialization of NVRAM.
- Number of seconds since the last switch reboot.

When the NMS starts up, it sends the switch a request message for this information. The switch responds with a report PDU that contains the information. Once the NMS and switch have synchronized this information, the NMS inserts it in all PDUs that it sends to the switch.

The switch discards any PDUs that contain:

- The incorrect number of switch reboots
- A 150-second or greater discrepancy in the number of seconds since the last switch reboot.

Views

Overview

A view defines the specific object identifiers (OIDs) that either can or cannot be accessed by the user groups to which the view is assigned.

View information is stored in the View-Based Access Control Model (VACM) MIB. Only users that have administrator access to the switch can use CLI commands to modify this MIB. Users that have administrator access can, however, give specific user groups read-write access to the MIB. Users or community strings assigned to those user groups can then use SNMP to view or modify the MIB.

For information on how to configure a view, see “Configuring Views.”

Predefined Views

When you upgrade the switch from v5.*x* application software to v6.*x*, the following predefined views are created:

- **normal**—Includes the MIB objects that the normal security level included. The view includes all the supported MIBs except those objects that are administrator only (snmpTargetMIB, snmpNotificationMIB, snmpProxyMIB, snmpUsmMIB, snmpVacmMIB, snmpCommunityMIB, genlic.mib, load.mib, and the promChassisSystemReset object).
- **admin**—Includes the MIB objects that the admin security level included. The view includes all the supported MIBs including the administrator-only objects. The USM and VACM MIBs are not included in the predefined admin view.
- **restricted**—Created for the initial SNMPv3 user. The view includes System Table and SNMP Table of RFC 1213 MIB, snmpEngine Table of snmpFrameworkMIB, snmpMPDStats Table of snmpMPDMIB, and usmStats Table of snmpUsmMIB.
- **internet**—Includes all MIB objects, including the USM and VACM MIBs. The view is created for administrative SNMPv3 users and is assigned to the predefined **internet** group, which provides read-write access to the view.

User Groups

Overview

Groups associate views with specific users or community strings and determine the access that the users or community strings have to the views. Access to a view can be read, write, or notify. Users and community strings are able to view, modify, and receive traps only from the object identifiers (OIDs) that are included in the views associated with their group.

Group information is stored in the User-Based Security Model (USM) MIB. Only users that have administrator access to the switch can use CLI commands to modify this MIB. Users that have administrator access can, however, give specific user groups read-write access to the MIB. Users or community strings assigned to those groups can then use SNMP to view or modify the MIB.

This section contains the following information:

- Predefined Groups
- Migration of Existing Community Strings

For information on how to configure groups, see "Configuring Groups."

Predefined Groups

When you upgrade the switch from v5.*x* application software to v6.*x*, the following predefined groups are created:

- **normalRO**—Provides read-only access to the **normal** view.
- **normalRW**—Provides read-write access to the **normal** view.
- **adminRO**—Provides read-only access to the **admin** view.
- **adminRW**—Provides read-write access to the **admin** view.
- **initial**—Provides read-only access to the **restricted** view. The predefined user **initial** is assigned to this group.
- **internet**—Provides read-write access to the internet view, which is created for administrative SNMPv3 users.

  The internet group is accessible only by SNMPv3. You cannot use CLI commands to access the group.

  Do not assign the internet group to a community string. The group requires both authentication and encryption, which community strings do not support.
- **noAccess**—Provides only notify access to the internet view. Assign community strings or users to this group when you want them to receive trap messages, but not have read or write access.

CAUTION:

Avaya recommends that you *not* modify these predefined groups. When you upgrade the switch from v5.*x* application software to v6.*x*, the existing community strings are assigned to these predefined groups. If you modify them, the community strings may lose their access to the switch. For more information on the migration of existing community strings, see "Migration of Existing Community Strings."

None of the predefined groups, except the internet group, require authentication or encryption. Thus these groups are accessible by SNMPv1, v2, or v3. The internet group, however, requires both authentication and encryption and is accessible only by SNMPv3.

Table 5-1 lists all the predefined views, groups, and users.

Table 5-1. Predefined Views, Groups, and Users

View	Group	User	Community String
normal—Includes the MIB objects that the normal security level included. The view includes all the supported MIBs except those objects that are administrator only.	**normalRO**—Provides read-only access to the **normal** view.		public
	normalRW—Provides read-write access to the **normal** view.		
admin—Includes the MIB objects that the admin security level included. The view includes all the supported MIBs including the administrator-only objects. The USM and VACM MIBs are not included in the predefined admin view.	**adminRO**—Provides read-only access to the **admin** view.		
	adminRW—Provides read-write access to the **admin** view.		
restricted	**initial**—Provides read-only access to the **restricted** view.	initial	
internet—Includes all MIB objects, including the USM and VACM MIBs. The view is created for administrative SNMPv3 users and is assigned to the predefined **internet** group, which provides read-write access to the view.	**internet** The internet group is accessible only by SNMPv3. You cannot use CLI commands to access the group. Do not assign the internet group to a community string. The group requires both authentication and encryption, which community strings do not support.		

Migration of Existing Community Strings

When you upgrade the switch from v5.*x* application software to v6.*x*, the existing community strings are assigned to either a predefined group or, if previously assigned to a custom access type, a group of the same name. All community strings that had a security level of normal are assigned to either the normalRO or normalRW group. All community strings that had a security level of admin are assigned to either the adminRO or adminRW group. All community strings that were assigned to a custom access type, are assigned to a group of the same name; however, that group is not automatically created. You must manually create the group.

Table 5-2 lists the groups to which existing community strings are assigned when you upgrade the switch from v5.*x* application software to v6.*x*.

Table 5-2. Migration of Existing Community Strings

Access of Pre-6.0 Community String	Security Level of Pre- 6.0 Community String	6.*x* Group
Read-only	Normal	normalRO
Read-write	Normal	normalRW
Read-only	Admin	adminRO
Read-write	Admin	adminRW
Custom access type		Group of same name
None		noAccess

ATM Uplink Module

Overview

The ATM Uplink module supports only SNMPv1 and v2. It does not support SNMPv3. To make secure SNMPv3 access to the ATM Uplink module possible, the application software uses a *proxy forwarder.*

The proxy forwarder converts SNMPv3 requests to SNMPv1 requests and then internally forwards them to the ATM Uplink module. When the ATM Uplink module responds to requests, the proxy forwarder converts the SNMPv1 responses to SNMPv3 responses and forwards them to the originator of the initial request.

To identify a PDU as destined for the ATM Uplink module, the NMS inserts the engine ID and context name of the ATM Uplink module. The engine ID and context name of the ATM Uplink modules are as follows:

- Engine ID—Engine ID of the switch, where the slot number of the ATM Uplink module is added to the last octet.

  For example, if the engine ID of the switch is 00:00:1a:e9:01:7f:00:00:01:00:00:00 and an ATM Uplink module is in slot 11, the engine ID of the ATM Uplink module is 00:00:1a:e9:01:7f:00:00:01:00:00:0B. (B is the hexidecimal value for 11.)

- Context name—**ATM<*mod-num*>**, where <*mod-num*> is the slot number of the module.

You must configure the engine ID and context name of the ATM Uplink module in the NMS.

Process

The proxy forwarder process is as follows:

1. An NMS prepares an SNMPv3 request PDU. This PDU contains the switch security credentials, including the engine ID and context name of the ATM Uplink module.
2. The supervisor modules receives the PDU from the NMS. It verifies that the security information matches the credentials of the switch. When the supervisor detects an engine ID and context name that differ from those of the switch, it checks the snmpProxyTable for a match.
3. The switch uses the information in the snmpProxyTable (and, because the indicated target is an SNMPv1 agent, the snmpCommunityTable) to convert the PDU to an SNMPv1 PDU and then send it to the ATM Uplink module.
4. The ATM Uplink module receives and processes the PDU and sends its response to the supervisor module.
5. The supervisor module matches the response to its pending forwarded requests, converts the SNMPv1 PDU to an SNMPv3 PDU, and forwards it to the NMS.

Accessing the ATM-Uplink Module

If an SNMPv3 user is assigned to a group that has read-only access to a view, he or she can view the ATM-Uplink module OIDs. However, for an SNMPv3 user to modify the ATM-Uplink module OIDs:

- The user must be assigned to a group that provides write access to a view.

  AND

- The view must include the promModuleTable.

For an SNMPv1 or v2 community string to access the ATM-Uplink module, you must create a string specifically for that module and assign it read-only or read-write access. For information on how to create a community string for the ATM-Uplink module, see “Creating or Modifying a Community String for the ATM-Uplink Module”.

Initial SNMPv3 User and SNMPv3 Administrator

Initial User

When you upgrade the switch from v5.*x* application software to v6.*x*, an initial SNMPv3 user is automatically created. The user name is **initial**, and it is assigned to the predefined group **initial**. The **initial** group has read-only access to the **restricted** view.

This **initial** user is created so that you can create new users by using SNMP. To create new users, you can clone the **initial** user. For more information on cloning a user, see RFC 3414.

SNMPv3 Administrator

After either manually assigning the switch an engine ID or assigning the console port an IP address, you can create an SNMPv3 administrator. This administrator will be able to create and modify SNMPv3 users by using SNMP.

You must have administrator access to the switch to create the SNMPv3 administrator. The administrator can, however, give specific user groups read-write access to the USM and VACM MIBs. Users assigned to those user groups can then use SNMP to create or modify SNMPv3 users.

CAUTION:

To prevent SNMPv1 or v2 access to the USM or VACM MIBs, assign views that include these MIBS only to groups that require authentication and encryption.

Changing the Engine ID of the Switch

You can override the default engine ID, which is based on the IP address of the switch. After changing the engine ID, however, you must change all SNMPv3 user passwords. For information on changing SNMPv3 user passwords, see "Changing a User Password."

***Important:** Each switch must have a unique engine ID. If you change the default engine ID of a switch, make sure that no duplicates exist in your network.

To change the engine ID of the switch, use the following CLI command:

`(configure)#` **snmp-server engineid [<*engine-Id*>]**

Table 5-3 provides an explanation of the command variable.

Table 5-3. snmp-server engineid Command

Parameter	Definition
<*engine-Id*>	A 12-byte hexidecimal value. Separate each byte with a colon. Do not use all 00s or all FFs. Example: 00:00:00:09:0a:fe:ff:12:97:33:45:00. **Important:** The last byte of the engine ID must be 00. This restriction makes SNMP access to the ATM Uplink module possible. The engine ID of an ATM Uplink module is the engine ID of the switch, where the slot number of the ATM Uplink module is added to the last byte.

Viewing the Engine ID

To view the currently configured engine ID of the switch, use the following CLI command:

`(configure)#` **show snmp engineid**

Creating an SNMPv3 Administrator

After either manually assigning the switch an engine ID or assigning the console port an IP address, you can create an SNMPv3 administrator. This administrator will be able to create and modify SNMPv3 users by using SNMP.

You must have administrator access to the switch to create the SNMPv3 administrator. The administrator can, however, give specific user groups read-write access to the USM and VACM MIBs. Users assigned to those user groups can then use SNMPv3 to create or modify SNMPv3 users.

CAUTION:

To prevent SNMPv1 or v2 access to the USM or VACM MIBs, assign views that include these MIBS only to groups that require authentication and encryption.

To create an SNMPv3 administrator, use the following CLI command:

`(configure)#` **snmp-server user <*username*> group internet auth {sha | md5} <*auth-password*> priv <*priv-password*>**

Table 5-4 defines the command keywords and variables.

Table 5-4. snmp-server user Command

Parameter	Definition
<username>	User name for the SNMPv3 user. The user name can range from 1 to 32 alphanumeric characters.
sha	Authenticates the user by means of HMAC-SHA.
md5	Authenticates the user by means of HMAC-MD5.
	1 of 2

Table 5-4. snmp-server user Command

Parameter	Definition
<*auth-password*>	The authentication password for the user: • Text passwords can range from 8 to 64 characters. • Localized HMAC-SHA-hashed passwords must be 20 bytes. • Localized HMAC-MD5-hashed passwords must be 16 bytes. Enter all localized passwords in the format of nn:nn:nn....
<*priv-password*>	The encryption password for the user. • Text passwords can range from 8 to 64 characters. • Localized, HMAC-SHA- or HMAC-MD5-hashed encryption passwords must be 16 bytes. Enter all localized passwords in the format of nn:nn:nn....
	2 of 2

Configuring Views

This section contains procedures for the following tasks:

- Creating or Modifying a View
- Viewing Configured Views
- Deleting a View

Creating or Modifying a View

To create or modify a MIB view, use the following CLI command:

`(configure)#` **snmp-server view <*viewname*> <*OIDST*> [{included | excluded}]**

Table 5-5 provides an explanation of the command keywords and variables.

Table 5-5. snmp-server view Command

Parameter	Definition
<*viewname*>	The name of the view that you want to create or modify. The view name can range from 1 to 32 alphanumeric characters.
<*OIDST*>	The object identifier (OID) for the object that you want to either include or exclude from the view. You must enter the numeric OID. Use the wildcard character ***** to specify a sub-tree family.
{included | excluded}	***Specifies whether the object is included or excluded from the view.***

Viewing Configured Views

To view the currently configured views, use the following CLI command:

show snmp view [<*viewname*>]

Deleting a View

To delete a view or remove an OID from a view, use the following CLI command:

`(configure)#` **no snmp-server view** <***viewname***> **[**<***OIDST***>**]**

Configuring Groups

This section contains procedures for the following tasks:

- Creating or Modifying a Group
- Viewing Configured Groups
- Deleting a Group

Creating or Modifying a Group

To create or modify a group, use the following CLI command:

`(configure)#` **snmp-server group <*groupname*> {noAuth | auth | priv} [read <*readview*>] [write <*writeview*>] [notify <*notifyview*>]**

Table 5-6 provides an explanation of the command keywords and variables.

CAUTION:

Avaya recommends that you *not* modify the predefined groups. When you upgrade the switch from v5.*x* application software to v6.*x*, the existing community strings are assigned to these predefined groups. If you modify them, the community strings may lose their access to the switch. For more information on the migration of existing community strings, see "Migration of Existing Community Strings."

Table 5-6. snmp-server group Command

Parameter	Definition
<*groupname*>	The name of the group that you want to create or modify. The group name can range from 1 to 32 alphanumeric characters.
noAuth	Requires neither authentication or encryption of PDUs.
auth	Requires authentication but not encryption of PDUs.
priv	Requires authentication and encryption of PDUs.
<*readview*>	The MIB view to which you want the group to have read access.
<*writeview*>	The MIB view to which you want the group to have write access.
<*notifyview*>	The MIB view for which you want the group to receive trap messages.

Viewing Configured Groups

To view the currently configured groups, use the following CLI command:

show snmp group [<*groupname*>]

Deleting a Group

To delete a specific group or all groups of a specific group name (if multiple groups have the same group name), use the following CLI command:

`(configure)#` **no snmp-server group <*groupname*> {noAuth | auth | priv}**

If multiple groups have the same group name, you must enter the appropriate security keyword (noAuth, auth, or priv) to delete one of the groups. If you do not enter a security keyword, all groups of the group name that you enter are deleted.

Configuring an SNMPv3 User

This section contains procedures for the following tasks:

- Creating or Modifying a User
- Adding a User to a Group
- Removing a User from a Group
- Viewing Configured Users
- Deleting a User

Creating or Modifying a User

To create or modify an SNMPv3 user, use the following CLI command:

`(configure)#` **snmp-server user <*username*> [group <*groupname*>] [[localized] auth {sha | md5} <*auth-password*> [priv <*priv-password*>]]**

Table 5-7 defines the command keywords and variables.

Table 5-7. snmp-server user Command

Parameter	Definition
<*username*>	User name for the SNMPv3 user. The user name can range from 1 to 32 alphanumeric characters. **Important:** Do not assign a community string and SNMPv3 user the same name.
<*groupname*>	Name of the group to which you are assigning the user.
localized	Use this keyword if you want to enter the authentication password and privacy password in their localized form instead of text. Localized passwords consist of the engine ID plus the password and are then hashed by either HMAC-SHA or HMAC-MD5.
sha	Authenticates the user by means of HMAC-SHA.
md5	Authenticates the user by means of HMAC-MD5.
<*auth-password*>	The authentication password for the user: • Text passwords can range from 8 to 64 characters. • Localized HMAC-SHA-hashed passwords must be 20 bytes. • Localized HMAC-MD5-hashed passwords must be 16 bytes. Enter all localized passwords in the format of nn:nn:nn....
<*priv-password*>	The encryption password for the user. • Text passwords can range from 8 to 64 characters. • Localized, HMAC-SHA- or HMAC-MD5-hashed encryption passwords must be 16 bytes. Enter all localized passwords in the format of nn:nn:nn....

Adding a User to a Group

To add an SNMPv3 user to a group, use the following CLI command:

`(configure)#` **snmp-server user** <***username***> **group** <***groupname***>

Removing a User from a Group

To remove an SNMPv3 user from a group, use the following CLI command:

`(configure)#` **no snmp-server user** <***username***> **group** <***groupname***>

Viewing Configured Users

To view the currently configured SNMPv3 users, use the following CLI command:

show snmp user [<***username***>**]**

This command displays the following information:

- User name
- Group to which the user belongs
- Authentication protocol that is used for the authentication and encryption passwords: `MD5`, `SHA`, or `None` if authentication is not enabled.
- Localized authentication key of the user if authentication is enabled. For information on localized keys, see "Authentication and Encryption."
- State of encryption, `Yes` if enabled, or `No` if disabled.
- Localized encryption key of the user if encryption is enabled. For information on localized keys, see "Authentication and Encryption."

Deleting a User

To delete an SNMPv3 user, use the following CLI command:

`(configure)#` **no snmp-server user** <***username***>

Changing a User Password

You must change user passwords when the engine ID changes.

To change a user password, use the following CLI command:

`(configure)#` **snmp-server password <*username*>**

The switch prompts you to enter the new password or passwords. The passwords are case-sensitive and can range from 8 to 64 characters. For security reasons, the CLI does not display the passwords when you enter them.

Configuring SNMPv1 or v2 Community Strings

With the introduction of SNMPv3, configuration of SNMPv1 and v2 community strings has changed. Community strings are now associated with groups instead of custom access types.

The ATM Uplink module, however, does not support SNMPv3 or groups. To access the ATM-Uplink module OIDs by community string, you must create a string specifically for that module and assign it read-only or read-write access.

This section contains procedures for the following tasks:

- Creating or Modifying a Community String for the Switch
- Creating or Modifying a Community String for the ATM-Uplink Module
- Viewing Configured Community Strings
- Setting the Trap Receiver
- Deleting a Community String

Creating or Modifying a Community String for the Switch

To create or modify a community string to access the switch, use the following CLI command:

`(configure)#` **snmp-server community <*community-string*> group <*groupname*> [<*ip-addr*> [notify]]**

Table 5-8 defines the command keywords and variables.

Table 5-8. snmp-server community

Field	Definition
<*community-string*>	The name of the community string. The community string can range from 1 to 26 characters. **Important:** Do not assign a community string and SNMPv3 user the same name.
<*groupname*>	Name of the group to which you are assigning the community string. **Important:** Do not assign the community string to a group that requires authentication or encryption. Community strings do not support authentication or encryption.
<*ip-addr*>	The IP address from which the community string is valid. Trap messages are sent to this IP address if you enter the **notify** option.
[notify]	Sends trap messages to the IP address that you specify.

Creating or Modifying a Community String for the ATM-Uplink Module

To create or modify a community string to access the ATM-Uplink module, use the following CLI command:

`(configure)#` **snmp-server atm-community <*community-string*> <*slot*> {ro | rw} [<*ip-addr*>]**

Table 5-9 defines the command keywords and variables.

Table 5-9. snmp-server community

Field	Definition
<*community-string*>	The name of the community string. The community string can range from 1 to 26 characters. **Important:** Do not assign a community string and SNMPv3 user the same name.
<*slot*>	The slot number of the ATM Uplink module.
	1 of 2

Table 5-9. snmp-server community

Field	Definition
ro	Assigns read-only access to the community string.
rw	Assigns read-write access to the community string.
[<*ip-addr*>]	The IP address from which the community string is valid.
	2 of 2

Viewing Configured Community Strings

To view the currently configured community strings, use the following CLI command:

show snmp community [<*community-string*>]

Setting the Trap Receiver

To set the trap receiver for a community string, use the following CLI command:

`(configure)#` **snmp-server notify <*ip-addr*> <*community-string*>**

Table 5-10 provides an explanation of the command variables.

Table 5-10. snmp-server notify Command

Parameter	Definition
<*ip-addr*>	The IP address to which you want trap messages sent.
<*community-string*>	The existing community string for which you are setting the trap receiver.

Deleting a Community String

Switch Community String

To delete a community string for the switch, use the following CLI command:

`(configure)#` **no snmp-server community <*community-string*> [<*ip-address*>]**

Table 5-11 provides an explanation of the command variables.

Table 5-11. no snmp-server community Command

Parameter	Definition
<*community-string*>	The community string that you want to delete.
[<*ip-addr*>]	The IP address from which you want to remove access to the community string.

ATM Module Community String

To delete a community string for the ATM-Uplink module, use the following command:

`(configure)#` **no snmp-server atm-community <*community-string*> [<*ip-addr*>]**

Table 5-12 provides an explanation of the command variables.

Table 5-12. no snmp-server community Command

Parameter	Definition
<*community-string*>	The community string that you want to delete.
[<*ip-addr*>]	The IP address from which you want to remove access to the community string.

Setting the Administrative Contact

CLI Command

To set the administrative contact for the switch, use the following CLI command:

`(configure)#` **snmp-server contact <*contact-name*>**

The switch displays the administrative contact when you enter the **show snmp** command. For information on this command, see "Viewing the SNMP Status."

Table 5-13 provides an explanation of the command variable.

Table 5-13. snmp contact Command

Parameter	Definition
<*contact-name*>	The name of the administrative contact for the switch. The contact name can range from 1 to 127 characters.

SNMP

To set the administrative contact for the switch, use the sysContact object, OID 1.3.6.1.2.1.1.4.

The complete path to this object is:

iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1).system(1).
sysContact(4)

Setting the Physical Location of the Switch

CLI Command

To set the physical location of the switch, use the following CLI command:

`(configure)#` **snmp-server location** <***server-location***>

The switch displays the physical location of the switch when you enter the **show snmp** command. For information on this command, see "Viewing the SNMP Status."

Table 5-14 provides an explanation of the command variable.

Table 5-14. snmp location Command

Parameter	Definition
<*server-location*>	The physical location of the switch. The location can range from 1 to 127 characters.

SNMP

To set the physical location of the switch, use the sysLocation object, OID 1.3.6.1.2.1.1.6.

The complete path to this object is:

iso(1).org(3).dod(6).internet(1).mgmt (2).mib-2(1).system(1).
sysLocation(6)

Disabling or Reenabling SNMP

Disabling SNMP

To disable SNMP, use the following CLI command:

```
(configure)# no snmp-server
```

This command overrides secure mode, which disables SNMPv1 and v2 and enables SNMPv3. For information on secure mode, see “Secure Mode.”

Reenabling SNMP

To reenable SNMP, use the following CLI command:

```
(configure)# snmp-server enable
```

This command enables the three versions of SNMP: SNMPv1, v2, and v3.

Viewing the SNMP Status

To view the status of SNMP, use the following CLI command:

```
(configure)# show snmp
```

This command displays the status of SNMP (enabled or disabled) and the administrative contact and physical location of the switch, if set.

6 Using VLANs, Hunt Groups, and VTP Snooping

Overview

The following information and procedures are provided in this chapter and pertain to layer 2 and layer 3 module configurations:

- VLAN Introduction
- Creating and Implementing VLANs
- Using Hunt Groups to Aggregate Bandwidth
- Configuring VTP Snooping

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.

VLAN Introduction

This section introduces Virtual Local Area Networks (VLANs) and explains how they operate.

Need for VLANs

If a traditional bridge receives a frame with a broadcast, multicast, or unknown destination address, it forwards the data to all bridge ports except the port on which it was received. This process is referred to as bridge *flooding*. As networks grow and the amount and types of traffic increase, bridge *flooding* may create unnecessary traffic problems that can clog the Local Area Networks (LAN).

To help control the flow of traffic through a switch and meet the demands of growing networks, vendors have responded by using:

- Customized packet filtering to control which packets are forwarded.
- More routers as broadcast firewalls to divide the network into broadcast domains.
- Spanning Tree Protocol to control the flow of traffic among LANs (for redundant links).

Advantages with VLANs:

- Segment traffic and usage patterns in a manner similar to creating subnets and segments in traditional networks.
- Reduce the cost of equipment moves, upgrades, and other changes.
- Simplify network administration.
- Create logical work groups for users who share the same system resources.
- Users not required to share the same physical location.
- Reduce the need for routing to achieve higher network performance and reduced costs.
- Control or filter communication among broadcast domains.

What is a VLAN?

The Avaya Multiservice Switch has the ability to create separate logical LANs on the same physical device. These logical segments are referred to as Virtual LANs (VLANs).VLANs are typically groups of users with similar job functionality (i.e. sales, marketing, engineering etc.) and share common resources. VLANs are not constrained by their physical location and can communicate as if they were on a common LAN. VLAN members can reside on single or multiple ports on one or more media modules on a switch, or on different switches. VLANs are limited broadcast domains, meaning all members of a VLAN receive every broadcast packet seen by members of the same VLAN, but not packets sent by members of a different VLAN. A router is required when communicating between different VLANs.

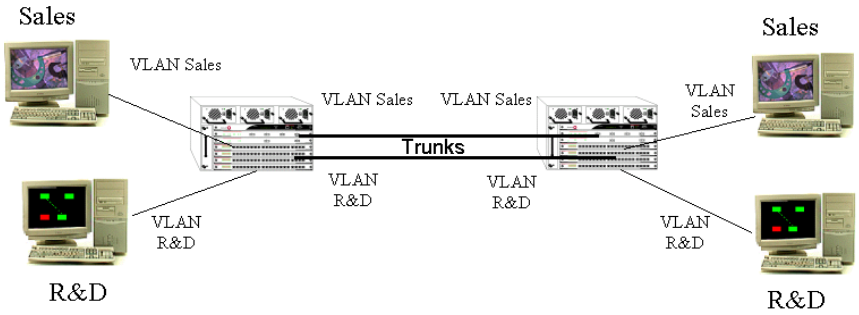
Port Based VLANS

The Avaya Multiservice switch is a Port Based VLAN architecture. VLAN membership is defined by groups of switch ports. When a VLAN is created three types of information is configured for that VLAN, the VLAN name, the VLAN Identifier or VLAN ID, and the switch ports assigned to that VLAN. VLAN assignment of a packet is based on a global VLAN ID. Regardless of any name you assign to a VLAN, the switch looks only at the VLAN ID number to determine a packet’s VLAN destination.

For example, ports 1, 2, and 3 on a module are members of VLAN A, ports 4, 5, and 6 on another module are also members of VLAN A. Traffic is forwarded through the switch to all ports that are members of VLAN A.

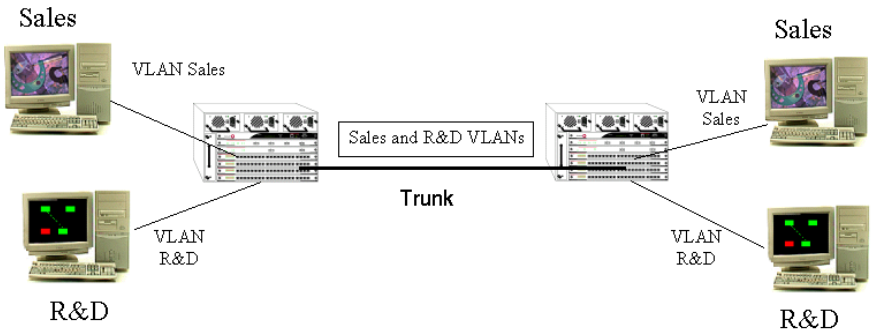
A switch that is strictly port based needs additional information to separate traffic passing from one VLAN to another. Since each port is dedicated to a particular VLAN, there is no need to analyze the traffic arriving on a port to determine its VLAN membership. As an example, (Figure 6-1) two switches with two VLANs, Sales and R&D, would need a dedicated switch to switch connections (Trunk) for both VLANs between switches. Therefore all traffic arriving on that port must belong to the VLAN assigned to that port since, in this case, no unique identifiers are sent with the frames.

Figure 6-1. Vlans No Tagging



The Avaya Multiservice switch is able to separate VLAN traffic between switches across a single Trunk port. To accomplish this, the switch implements VLAN tagging and trunking. VLAN tagging is enabled on a switch port by selecting a Trunk mode for that port; clear, IEEE 802.1Q or Cisco-Multi Layer mode. A trunk port can send frames in clear mode, with no VLAN ID, or the VLAN ID, over the same trunk. A frame is classified as belonging to a particular VLAN based on the value of the VLAN Identifier (VID) that is included in the Tag Header. Therefore using our example, and implementing VLAN tagging, we need only one connection (trunk) between the two switches to carry the traffic from both VLANs.(Figure 6-2)

Figure 6-2. VLANs with 802.1Q tagging



IEEE 802.1Q VLAN Tagging

The Avaya Multiservice switch is compliant with the IEEE 802.1Q standard for VLANs and defines a Tag Header. Two Tag formats are defined as an Ethernet Encoded (4 bytes) for 802.3 and Ethernet V2 and SNAP (Service Network Access Point) for Token Ring and Fiber Distributed Data Interface (FDDI).

The Ethernet version of the Tag Header consists of 4 bytes, two bytes for Tag Protocol ID and two bytes for Tag Control. The Tag Protocol ID bytes contain an Ethernet Type value of 81-00 which identifies the frame as a tagged frame. The Tag Control specifies tag formats that are used to embed explicit VLAN membership information within each frame in a 12-bit VID that provides 4094 possible VLAN ID's. IEEE 802.1Q defines the bridging rules for VLANs (ingress and egress rules which are described in detail in "VLAN Operation" later in this chapter).

VLAN Operation

VLAN operation is based on three sets of rules:

- Ingress Rules
- Forwarding Rules
- Egress Rules

Ingress Rules

Every frame received by the switch is classified to one VLAN. There are two ways in which frames are classified to VLANs:

- Untagged frames are classified to the VLAN associated with the port on which the frame is received (Port-based VLANs).
- Tagged frames are classified to the VLAN identified by the VLAN tag in the tag header of the frame.

***Note:** The switch supports a feature called **Automatic VLAN Creation** for tagged frames. For more information, see "VLAN Considerations" for more information.

Forwarding Rules

These rules determine the set of ports on the switch through which members of the VLAN can be reached. This is called binding a port to a VLAN. A port may be bound to a VLAN using the Web Agent in the following three ways:

1. Setting the **Port VLAN** attribute in the Switch Port Configuration Web page.This identifies the VLAN to which all untagged frames received on the port are forwarded. **Static Binding**, the port is bound to the VLAN selected in the Port VLAN parameter.

*** Note:** A port has one **Port VLAN**. Changing this to a new VLAN removes the port from the old VLAN.

2. Setting the **VLAN Binding** attribute in the Switch Port Configuration Web page to **Bind to All** should be done on links connecting two layer2 switches, where multiple VLANs span across both switches, such that members of each VLAN are found on both sides of the link. Bind-to-all should not be used when the switches on both ends of the link act as routers, such that each IP subnet and each VLAN are confined to one side of the link only and do not have members connected to the switch at the other end. In such routing cases, the link is never used for intra-VLAN traffic but rather is used only for traffic routed from one router to the other. Thus, there is no need for the link to belong to multiple VLANs, and should not be configured to bind-to-all. It should be bound to a single VLAN that is dedicated to the connection between the two routers. Bind-to-all in this case is not only unnecessary, but also undesired as a lot of irrelevant broadcast/multicast traffic of other VLANs will be sent onto this link and into the switch on the other end, unnecessarily increasing the control-plane load on the supervisor and increasing the chance for harmful layer3 configuration errors.

3. Setting the **VLAN Binding** attribute in the Switch Port Configuration Web page to **Bind to Received**. This causes the port to be bound to all VLANs (as identified by the VLAN tag in tagged frames) received on this port. Consequently, ports are bound to those VLANs that actually have members that are reachable through the port.

***Note:** When an untagged frame arrives on a port that is set to **Bind to ALL,** it forwards the frame to the “port VLAN”. When a tagged 802.1Q frame arrives on a port that is set to **Bind to All** and the VLAN doesn’t exist on the switch the frame is dropped.

Ingress: Untagged frames are classified to the VLAN associated with the port on which the frame is received. Tagged frames are classified to the VLAN identified by the VLAN tag in the tag header of the frame.

Forwarding: Only forward frames to the port for the assigned VLAN.

Egress: All frames transmitted out of the port to be tagged using the IEEE 802.1Q/Multi-Layer tag header format. The tagged used will be that assigned to the port.

Binding a Port to more than one VLAN using the CLI

A port can be statically configured to Bind to more than one VLAN. This causes the port to become a member of each specified VLAN. This feature is configured from the CLI. This feature cannot be configured from the Web Agent.

To configure from the CLI use the following command:

`(configure)#` **set vlan {**<***vlan-id***> | **name** <***vlan-name***>**}** <***mod-swport-range***>**[...,**<***mod-swport-range***>**]**



CAUTION:

This configuration should only be used under special circumstances and with the assistance of customer support.

*** Note:** The Web Agent does not allow static binding of multiple VLANs to a single port. If you use the CLI to statically bind multiple VLANs to a single port, do not attempt to use the Web Agent to statically bind additional VLANs to the port or remove existing VLANs from the port.

Egress Rules

For a given port/VLAN combination, egress rules determine whether frames transmitted from the port on the VLAN are tagged or untagged. This is accomplished by setting the Trunking Mode attribute in the Switch Port Configuration Web page.

For example, choosing the IEEE 802.1Q trunk mode causes all frames transmitted out of the port to be tagged using the IEEE 802.1Q tag header format. Individual port/VLAN combinations may be changed to cause frames transmitted from the port to be untagged (or clear mode).

Creating and Implementing VLANs

Adding users to VLANs include:

- VLAN Considerations
- Creating a VLAN
- Configuring VLAN Parameters
- Assigning Ports to VLANs

VLAN Considerations

Be aware of the following issues when configuring VLANs:

- If you set **Trunk Mode** to **Clear**, you must set the **VLAN Binding Type** to **Static** (default).
- The switch supports a feature called **Automatic VLAN Creation** for tagged frames. When this feature is enabled, the switch creates new VLANs when it receives packets from previously unknown VLANs. Vlan's can be created automatically without manually creating the VLAN on each switch.
- **Automatic VLAN Creation** does not work on an 80-series supervisor module (M8000R-SUP) if the **Port default VLAN** is set to **discard**.
- If you enable **Automatic VLAN Creation** AND set "VLAN Binding" type to **Bind to Received**, make sure that you set the binding type before enabling **Automatic VLAN Creation** or else the port may not be automatically added to the VLAN.
- Assigning one of the ports of a 48-port module to a VLAN for which there exists IP interface will force that interface to be in the UP state, even when there is no physical connection to the module.

Creating a VLAN

You can create a VLAN using either the Web Agent or the CLI.

CAUTION: Before creating or adding additional VLANS, it is strongly recommended that you review the contents of this chapter and Chapter 9, "Managing the Address Forwarding Table."

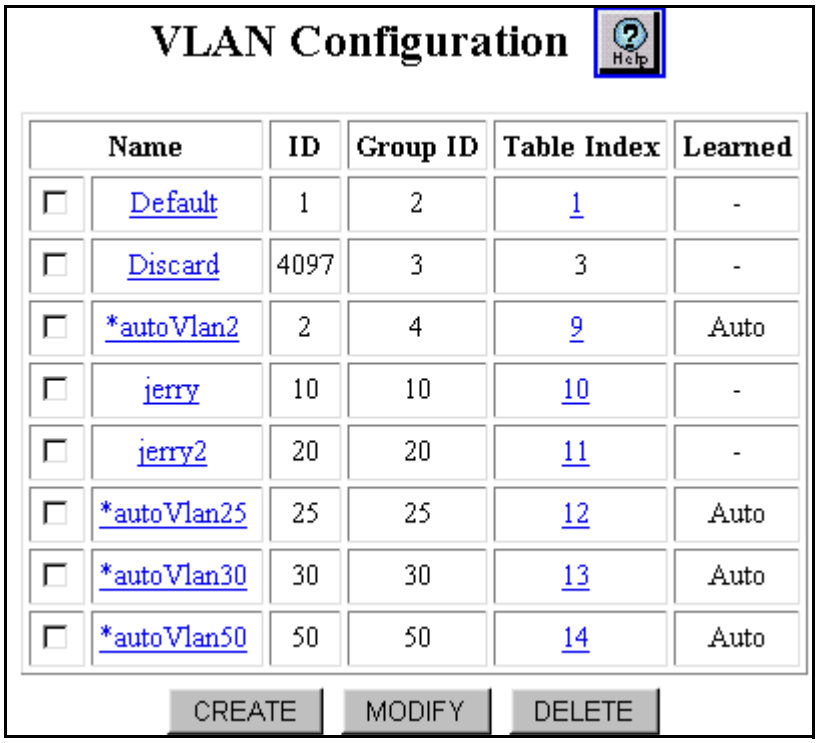
Web Agent Procedure

To create a VLAN:

1. In the navigation pane, expand the **L2 Switching > VLANs** folders, and then click **Configuration**.

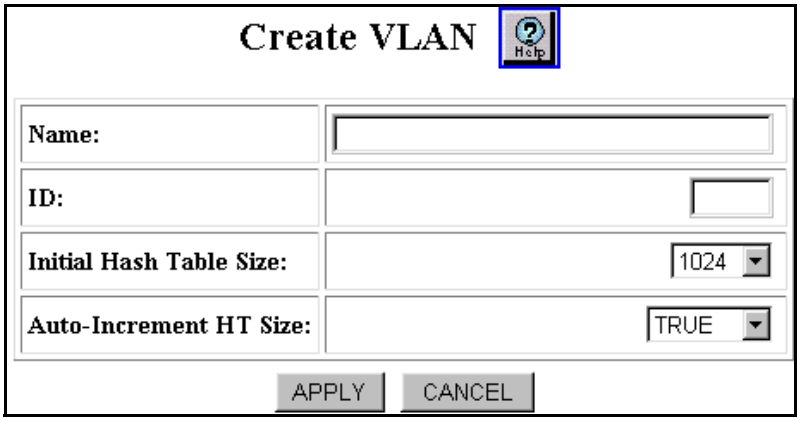
   The VLAN Configuration Web page is displayed in the content pane. See Figure 6-3.

Figure 6-3. VLAN Configuration Web Page



2. Click **CREATE**. The Create VLAN Web page is displayed in the content pane. See Figure 6-4.

Figure 6-4. Create VLAN Web Page



3. Enter a name for the VLAN in the **Name** field.

4. In the ID field, enter an unused VLAN ID value (between 2 to 4094).

VLAN IDs are global and must be consistent from switch to switch, even when switches are manufactured by different vendors.

5. See Table 6-1 to configure the Create VLAN Web page parameters:

Table 6-1. Create VLAN Web Page Parameters

Parameter	Definition
Name	Unique Name assigned to the VLAN in a switch. A maximum of 31 alphanumeric characters
ID	Identifier used throughout the network to identify this VLAN. If you want ports on more than one device to participate in a particular VLAN, you must use the same VLAN ID to identify the VLAN on every device.VLAN ID 1 is reserved for the Default VLAN. VLAN ID 4097 is reserved for the discard VLAN. **Note:** The switch supports up to 1000 VLANs and a maximum of 24,000 MAC Addresses. **Note:** In order to support the maximum number of VLANs, VLAN ID numbers should be chosen from the range of 1 to 1000.
	1 of 2

Table 6-1. Create VLAN Web Page Parameters

Parameter	Definition
Initial Hash Table Size	For every VLAN created, a hash table is allocated of the initial size. The hash table keeps track of the Layer 2 (MAC) addresses associated with a VLAN. The default setting for **Initial Hash Table Size**,1024, allows the user to configure approximately 20 simultaneous VLANs. However, the switch supports 1,000 VLANs. To support 1000 simultaneous VLANs, the Hash table size should be at 16. To increase the number of VLANs you can implement, simply decrease the initial hash table size for each new address table instance. **Note:** The number of addresses for a given hash table is approximately 4:1 (for example, if you have a hash table of 6 bytes, the VLAN can hold 64 addresses in its table instance.
Auto Increment HT Size	*Auto-Increment HT (Hash Table) Size* determines whether the size of the VLAN's Hash Table can grow if software determines it to be undersized for the number of known Layer 2 addresses of a VLAN. Select TRUE to allow the switch to increase the size of the Hash Table. Select FALSE to disable the feature. The default is TRUE.
	2 of 2

6. Click **APPLY** to create the new VLAN, or **CANCEL** to restore previous settings.

CAUTION:

Before configuring Hash Tables and Auto Increment Hash Table parameters, Avaya strongly recommends that you review the contents of Chapter 9, "Managing the Address Forwarding Table."

Configuring VLAN Parameters

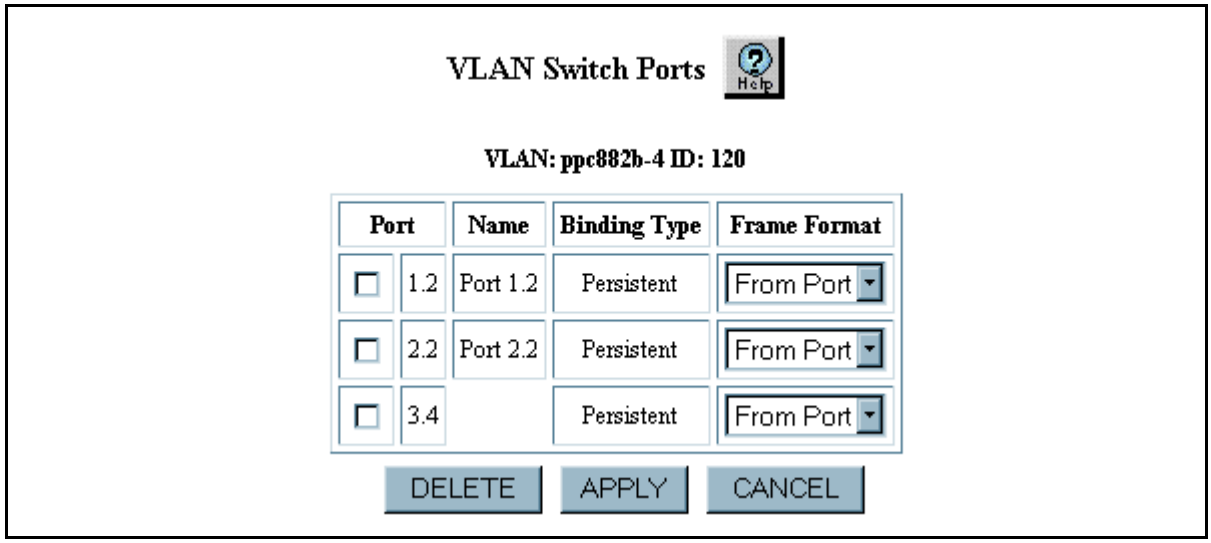
To configure all ports assigned to a VLAN:

1. In the navigation pane, expand the **L2 Switching > VLANs** folders, and then click **Configuration**. The VLAN Configuration Web page is displayed in the content pane. See Figure 6-3.

2. Click on the Name of the VLAN whose members you want to view.

   The VLAN Switch Ports Web page is displayed in the content pane. See Figure 6-5.

Figure 6-5. VLAN Switch Ports Web Page



3. See Table 6-2 to configure the VLAN Switch Ports Web page parameters:

Table 6-2. VLAN Switch Port Web Page Parameters

Parameter	Defines...
Port	The switch port associated with the VLAN.
Name	The switch default port name or the user assigned port name.
	1 of 2

Table 6-2. VLAN Switch Port Web Page Parameters

Parameter	Defines...
Binding Type	The binding types are set through the switch ports. See "Assigning VLANs to a Port and Associated Issues" in Chapter 8, "Configuring Ports." • **Static** - when switch ports are added manually and can be removed • **Persistent** - when switch ports are bound to VLANs automatically but can not be removed. automatically (i.e., if the binding for a switch port is set to 'Bind to All'). • **Dynamic** - when a switch port is assigned to VLAN using automatic VLAN binding,(i.e. if the binding for a switch port is set to "Bind to Received"). The VLAN may be deleted, but if the port VLAN binding is "Bind to Received", the VLAN may be re-added by receiving tagged traffic. This causes the port to again dynamically bind to the VLAN. See "Configuring Switch Ports Settings" in Chapter 8, "Configuring Ports," for more information.
Frame Format	• **From Port** - causes port to send frames using the frame format specified in the Trunk Mode attribute of the corresponding switch port. • **Clear** - causes port to send untagged frames on this port for this VLAN. **Note:** See "VLAN Introduction" earlier in this chapter.
	2 of 2

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Assigning Ports to VLANs

See "Configuring Switch Ports Settings," in Chapter 8, "Configuring Ports," for instructions on assigning ports to VLANs.

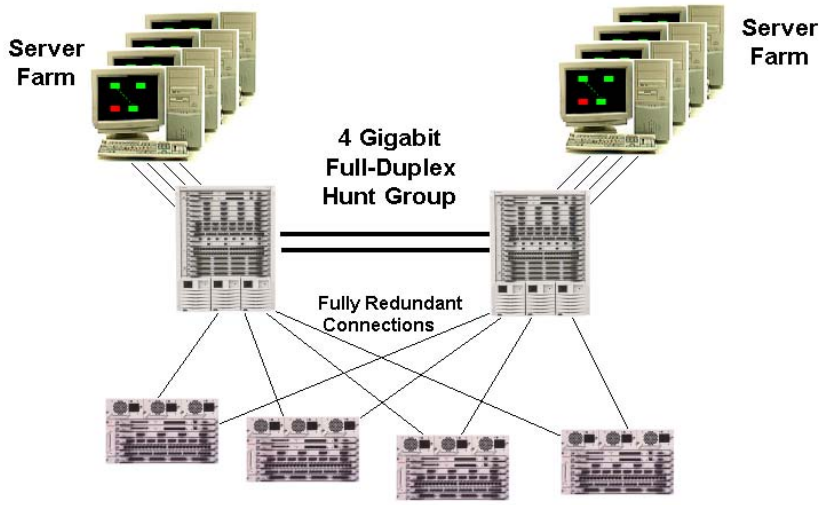
CAUTION: Before changing the VLAN that a port is assigned to, you must remove the port from the hunt group that the port is assigned to. Do *not* attempt to change the VLAN and remove the port from the hunt group simultaneously. If you do, the switch assigns all ports in the hunt group to the new VLAN.

Using Hunt Groups to Aggregate Bandwidth

Overview

Hunt groups make it possible for you to aggregate multiple switch ports to act as one switch port, effectively combining the bandwidth into a single connection. Frames sent to the hunt group will be transmitted by one of the ports in the hunt group. Hunt groups load share traffic across all member ports, this aggregates the bandwidth of the hunt group ports. Hunt groups also provide fault tolerance. If a port in a hunt group fails, the remaining ports in the hunt group will assume the traffic and continue forwarding. Figure 6-6 shows an example of a hunt group.

Figure 6-6. Hunt Group Example



Hunt groups expand the capacity of the core switched backbone. While there is no specific limit on the number of ports in a hunt group (see "Hunt Group Configuration Considerations"). In the example, the hunt group consists of two full-duplex Gigabit Ethernet links shared between two switches, for an aggregate capacity of 4 Gbps. Enterprise-level servers can be directly attached to the switching core using either 100 Mbps or Gigabit Ethernet connections. Fault-tolerant spanning tree links from the core to the next level of network spread the distribution and capacity to building or departmental switching centers. Each of these next-level switches may in turn support tens (or even hundreds) of work group switches and users.

This section contains the following information and procedures:

- Hunt Group Features
- Hunt Group Terminology
- How Hunt Groups Load Share
- Hunt Group Example
- Optimizing Hunt Group Throughput
- Hunt Group Configuration Considerations
- Hunt Groups and Multicast Traffic
- Creating Hunt Groups
- Adding Ports to the Hunt Group
- Viewing Hunt Group Members
- Removing Ports from Hunt Groups

Hunt Group Features

Hunt group features include:

- Shared traffic load.
- Fault tolerance — If a port in a group fails, the remaining ports in the group pick up the traffic load.
- Support for any number of same-speed same media type connections in a group — The group is not restricted to a single module in a switch.

- Quicker recovery from link failure — If a port in the group fails, the remaining ports carry the load. Recovery is not limited by Spanning Tree Protocol convergence time (convergence time is the time the network takes to resume steady-state forwarding after Spanning Tree Protocol reconfiguration).
- Supports up to 25 hunt groups per switch, if 48-port mode is disabled and up to 13 hunt groups if 48-port mode is enabled.

Hunt Group Terminology

Base Port/Flood port- When the hunt group is configured, there is one port designated *base port*. All ports in the hunt group assume the identity of the base port. The base port passes all flood frames, broadcast frames, destination unknown unicast, and multicast frames for VLANs associated with the hunt group. Spanning Tree treats all ports in the hunt group as one port. The base port sends and receives Bridge Protocol Data Units (BPDU's).

Member port - A port that is a member of the hunt group. Sometimes referred to as a "Participating port".

Non-member port - A port that is not a member of a hunt group. sometimes referred to as a "non-Participating port".

Forwarding Engine (FE) - A generic name for hardware that makes layer 2 and layer 3 forwarding decisions.

Participating (or Non-Participating) Forwarding Engine - A forwarding engine that has a port is part of the hunt group. Example, a M8024 (80-series module with 24 10/100 Mbps ports), the first 12 ports are serviced by a single Forwarding Engine. If any of those 12 ports participates in a hunt group, then that Forwarding Engine is considered a participating forwarding engine.

How Hunt Groups Load Share

Forwarding Engines and Fabric Ports

Each 80-series Media module has at least two Forwarding Engines, and as many as eight. These engines are the major hardware components for bridging and routing data traffic. The Forwarding Engines can service up to twelve 10/100 Mbps ports. Example, 24-port 10/100 Mbps 80-series media module has two Forwarding Engines that service twelve 10/100 Mbps ports each. An eight-port Gigabit module has eight forwarding engines one for each of the Gigabit ports per slot. See Table 6-3 for the relationship of media modules (50-series and 80-series) to Forwarding Engines.

80-Series Forwarding Engines

The Forwarding Engines are numbered starting at one and increased by one for each Forwarding Engine. The Supervisor module always has Forwarding Engines numbered one and two. For a 7-slot P580 switch, numbers one and two for the Supervisor module and up to eight per media module slot for a maximum total of 50 Forwarding Engines. For a 17-slot P882, numbers one and two for the Supervisor module and up to eight for each of the media modules, for a maximum total of 130. The numbering sequence will vary depending on the type and total number of media modules. If the switch has empty slots the numbering sequence is bypassed for the empty slots.

Fabric Ports

The P580/P882 has two backplane connections or ports for each media module. These ports connect the media modules to the backplane Cross Bar Switch Fabric, which in turn connects all modules to all other modules and ports. These ports are referred to as “Switch Fabric Ports”. There are two Switch Fabric ports per media module slot (slots 2 through 7, or slots 2-17) and one fabric port for the Supervisor slot (slot 1) (Figure 6-7).Each 80-series media module can use up to two fabric ports in the slot, the number of Forwarding Engines to fabric ports varies with the media module. There is anywhere from one to four Forwarding Engines per fabric port for each 80-series media module (Table 6-3).

The number of Forwarding Engines are distributed among the ports for 80-series modules. The 80-series 10/100 Mbps media modules are distributed with 12 10/100 ports for each Forwarding engine. For example, the 48 port 10/100 media module has four Forwarding Engines. The first 12 (1-12) ports are on the first Forwarding Engine, the second 12 (13-24) are on the second Forwarding Engine, which in turn use one fabric port, the third and fourth grouping of 12 ports each have a Forwarding Engine and use the second fabric port. The Forwarding Engines are distributed on a one to one basis for the 80-series Gigabit media modules. For example the 8 port 1000 TX, the first four ports have four Forwarding Engines, one for each Gigabit port (1-4), and uses one fabric port, the other four ports (5-8) also have four Forwarding Engines, which use the second fabric port.

50-series Forwarding Engines and Fabric Ports

On a 50-series layer 3 module, each fabric port has two Forwarding Engines, one for layer 3 traffic, and one for layer 2 traffic. All layer 3 traffic will be associated with one Forwarding Engine and all the layer 2 traffic will be associated with the other Forwarding Engine.

Example, the twelve-port 10/100 layer 3 50-series media module, layer 3 traffic is coming in the first port and the last port, (ports 1,12). Using the 80-series even distribution, the first six ports are associated with one FE and the second six ports with the second FE, and distributed across both fabric ports.

This is different for the 50-series modules. The twelve port layer 3 50-series media module, all twelve ports are associated with one Forwarding Engine for layer 3 traffic and one Forwarding Engine for layer 2 traffic. In addition only one fabric port is used for all twelve ports.

Load Share Function

Hunt groups load share by directing different traffic to different ports in the hunt group, when sending traffic to a particular user. Load sharing is done in a round-robin fashion across the ports in a hunt group. This is based upon **BOTH** the **Destination MAC Address** and the **Source Forwarding Engine**. The hunt group ports on which unicast packets traverse to reach the destination depends upon the source user's associated Forwarding Engine.

Hunt Group Example

See Figure 6-7. One port in the hunt group will be designated as the *base port* or *flood port*. All flood traffic for all VLANs is sent through this port only. All ports are members of all VLANs associated with the hunt group base port. There are 8 non-member Forwarding Engines. Load sharing is accomplished by using the combination of the Source Forwarding engine and the Destination MAC Address to assign a hunt group port. Users A and B are associated with FE#9 and therefore the first port in the hunt group will be used for unicast packets from A to Destination E and from B to E. Users C and D are associated with FE#10 and therefore the second port in the hunt group will be used for unicast packets from C to E and D to E.

When MAC Address E is learned, it is assigned to the first hunt group port for FE#1, the second hunt group port for FE#2, the first port for FE#3, the second port for FE#5, the first port for FE#7, the second port for FE#8, the first for FE#9, and the second port for FE#10. The second destination MAC Address is round-robin assigned in the same fashion and so on.

Figure 6-7. Slot to Fabric Port Relationship

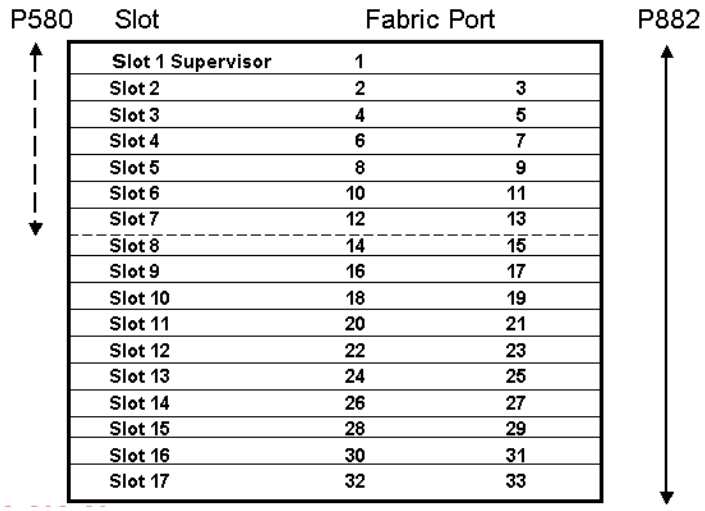


Table 6-3. Module Forwarding Engines and Fabric Ports

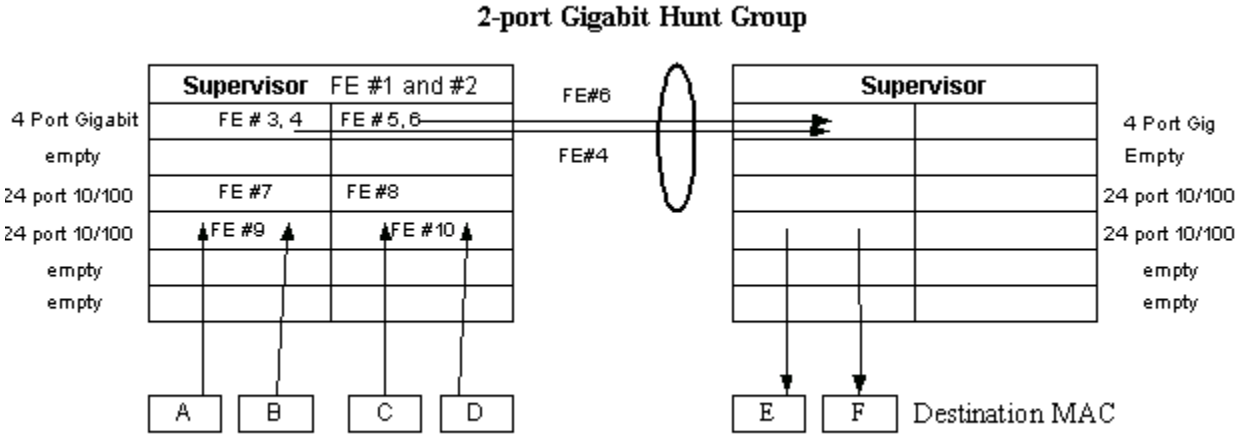
Module Type	Total Number of forwarding engines	Number of Fabric Ports	Description
80-series M8000R Supervisor	2	1	FORE port, CPU
80-series 24 port 10/100Mbps TX and 100Mbps FX	2	2	1 forwarding engine per fabric port
80-series 48 port 10/100Mbps	4	2	2 forwarding engines per fabric port
80-series 4 port Gigabit Fiber or TX	4	2	2 forwarding engines per fabric port
80-series 8 port Gigabit Fiber or TX	8	2	4 forwarding engines per fabric port
50-series 10 port 100Mbps FX(layer 2)	1	1	1 forwarding engine to 1 fabric port only
50-series 10 port 100Mbps FX (layer 3)	2	1	2 forwarding engines to 1 fabric port only
50-series 12 port 10/100 TX (layer 3)	2	1	2 forwarding engines to 1 fabric port
50-series 2 port Gigabit Fiber (layer 2)	2	2	1 forwarding engine to 1 fabric port
50-series 2 port Gigabit Fiber (layer 3)	4	2	1 forwarding engine to 1 fabric port
			1 of 2

Table 6-3. Module Forwarding Engines and Fabric Ports

Module Type	Total Number of forwarding engines	Number of Fabric Ports	Description
50-series 4 port Gigabit Fiber (layer 2)	4	2	2 forwarding engines per fabric port
50-series 20 port 10/100Mbps TX (layer 2)	2	2	1 forwarding engine per fabric port
			2 of 2

*** Note:** All 80-series media modules are L3 capable.

Figure 6-8. Load Sharing Example



Optimizing Hunt Group Throughput

To optimize hunt group throughput follow these considerations.

- Forwarding Engines in the hunt group ports should be separate from the ports involved in the modules traffic flows.
- The greater the number of non-member Forwarding Engines in the system, the greater the ability of the system to load-share.
- Distribute the hunt group across more than one Fabric Port per slot.
- Distribute the hunt group across media modules.

Hunt Group Configuration Considerations

Consider the following before creating hunt groups:

- Hunt group end to end physical connections must be configured with the hunt group ports connected to hunt group ports.
- All ports of the hunt group must be of the same bandwidth and all layer 2 or all layer 3 ports (but not a mix)
- Can have a mix of 80-series and 50-series as long as they are the same bandwidth and all layer 2 or all layer 3. (but not a mix)
- A hunt group cannot be distributed between more than two switches.
- You must disable or disconnect the ports in a hunt group until both ends of the link are configured.
- All ports in the hunt group take on the configuration of the base port
- Both sides of the hunt group have the same configuration
- Any change to a port in the hunt group will be reflected on all ports.
- When removing a port from a hunt group with the Web Agent, it must be removed before its VLAN assignment is changed.
- If 48-port mode is enabled on the switch, only the *first thirteen hunt groups* are retained and all others are discarded. The ports associated with the discarded hunts groups lose their hunt group bindings, but retain their VLAN bindings and are still active as VLAN bridges. Spanning tree may or may not, due to the topology or spanning tree settings on those ports, forward traffic over unwanted trunks.

Hunt Groups and Multicast Traffic

All multicast traffic is forwarded by the base port of the hunt group. If the base port changes, multicast traffic is interrupted until the sessions are relearned through the new base port.

To decrease or eliminate this multicast traffic interruption, you can:

- To *decrease* the multicast traffic interruption, decrease the IGMP query interval from 125 seconds to 5 seconds:

  a. Open the IGMP Interfaces Web page (**Routing** > **IGMP** > **Interfaces**).

  b. In the **Query Request Interval in (sec)** field, enter **5**, and then click **APPLY**.

- To *eliminate* the multicast traffic interruption, configure all ports in the hunt group as static router ports, and bind the router ports to all VLANs:

  a. Open the Router Port Display/Configuration Web page (**L2 Switching** > **Global Configuration** > **Display/Configure Router Ports**).

  b. In the **Port** field, enter a hunt group port number.

  c. In the **VLAN** field, select **All**, and then click **CREATE**.

  d. Repeat for all ports participating in the hunt group.

  For more information on configuring static router ports, see "Configuring Static Router Ports" in Chapter 20, "Managing Intelligent Multicasting."

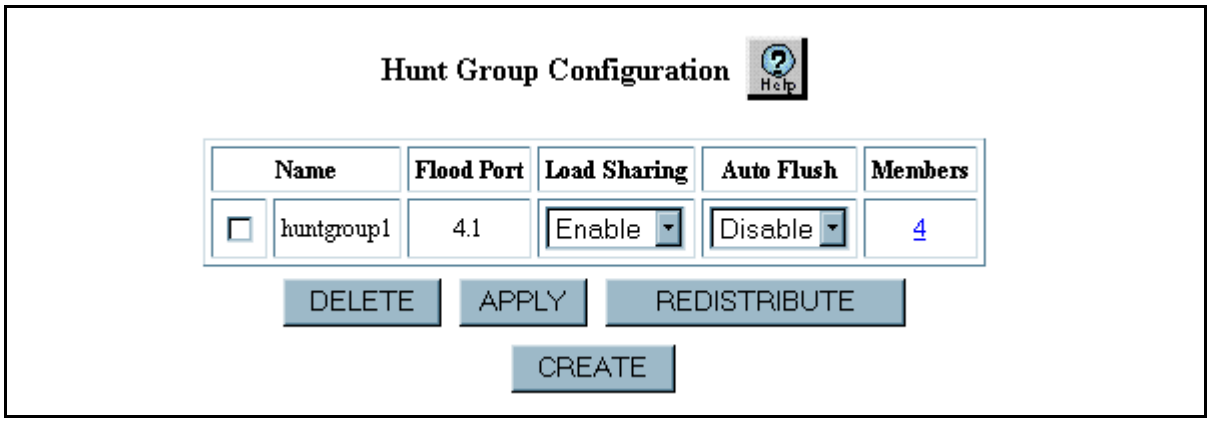
Creating Hunt Groups

Web Agent Procedure

To create a hunt group:

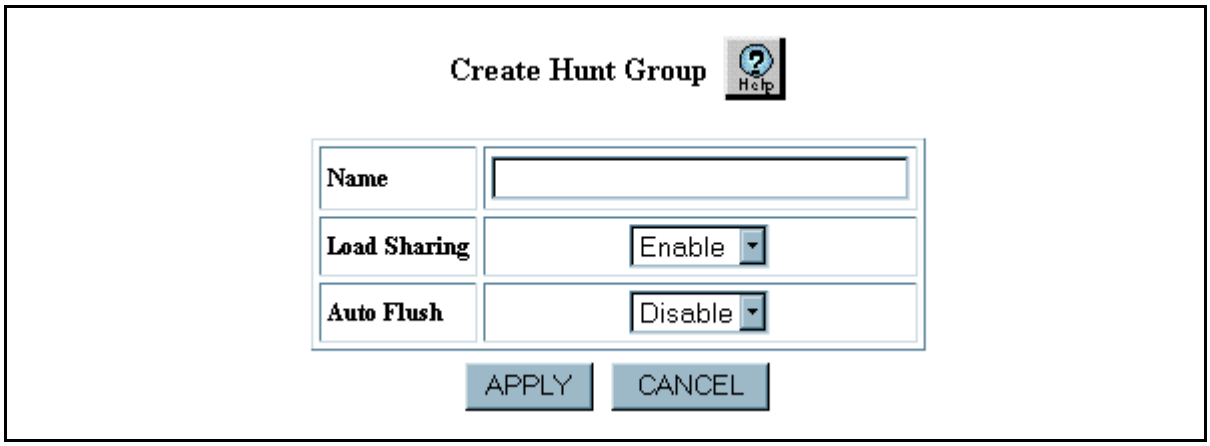
1. In the navigation pane, expand the **L2 Switching** folder, and then click **Hunt Groups**. The Hunt Group Configuration Web page is displayed in the content pane. See Figure 6-9.

Figure 6-9. Hunt Group Configuration Web Page



2. Click **CREATE**. The Create Hunt Group Web page is displayed in the content pane. See Figure 6-10.

Figure 6-10. Create Hunt Group Web Page



3. In the **Name** field, enter a name for the hunt group. This field accepts a maximum of 31 alphanumeric characters.

4. In the **Load Sharing** field, select **Enable**.

5. In the **Auto Flush** field, select **Enable** if you want to enable the auto flush feature for the ports participating in the hunt group. The default setting for this field is **Disable**.

   When you enable auto flush for a hunt group, all AFT entries that were learned on the hunt group are marked invalid if the links to all of the hunt group ports fail. Once the AFT entries are marked invalid, they can be learned on a redundant port. When auto flush is enabled, failover to a redundant port occurs much sooner.

6. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Commands

Use the following CLI commands to configure a hunt group:

- To create a hunt group, `(configure)#` **set huntgroup <*huntgroup-name*> [load-sharing {enable | disable}]**
- To enable or disable the auto flush feature for the ports participating in a hunt group, `(configure)#` **set huntgroup auto-flush <*huntgroup-name*> {enable | disable}**

Adding Ports to the Hunt Group

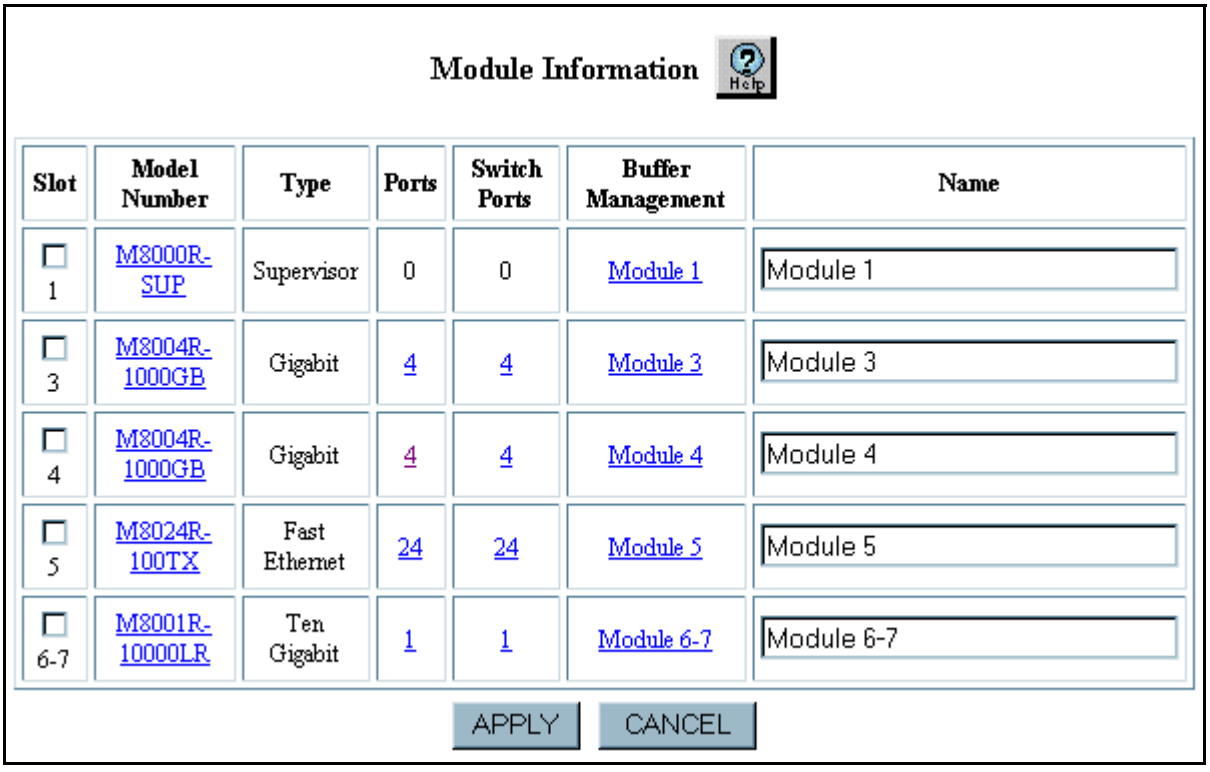
Once you have created and named a hunt group, you can add as many additional ports as needed to the hunt group. The next two hunt group configuration steps, adding physical ports, and switch ports to the hunt group, can be done using the Web Agent or the CLI.

Web Agent Procedure

To add ports to a hunt group using the Web Agent:

1. In the navigation pane, expand the **Modules and Ports** folder, and then click **Configuration**. The Module Information Web page is displayed in the content pane. See Figure 6-11.

Figure 6-11. Module Information Web Page



2. Select a port number from the **Ports** column. The Physical Port Configuration Web page is displayed in the content pane. See Figure 6-12.

Figure 6-12. Physical Port Configuration Web Page

Physical Port Configuration - Module 4 Help

Port	Name	Enable	Status	Type	Connector	Auto Negotiation Mode	Speed State	Duplex State	Flow Control State
4.1	Port 4.1	☑	No Link	Gigabit	GBIC-1000-SX	Enabled	1 Gb/s	Full Duplex	Disabled
4.2	Port 4.2	☑	No Link	Gigabit	GBIC-1000-SX	Enabled	1 Gb/s	Full Duplex	Auto-Negotiating
4.3	Port 4.3	☑	No Link	Gigabit	GBIC-1000-SX	Enabled	1 Gb/s	Full Duplex	Disabled
4.4	Port 4.4	☑	No Link	Gigabit	GBIC-1000-SX	Enabled	1 Gb/s	Full Duplex	Auto-Negotiating

APPLY CANCEL Prev Module Next Module Modules All Module Ports Configuration

3. Uncheck the checkbox from the **Enable** column for a specific port. The check mark disappears to disable the port.

*** Note:** Check the speed of the ports that you are configuring into the hunt group. The ports must be physically connected to each other (for example, in a four-port Gigabit hunt group, ensure that you have four fiber cables with switch ports connected at each end).Also Auto-negotiation should be disabled on the ports to further ensure against speed mismatch

— If this is a new hunt group, disable all of the ports you are adding to the hunt group.

— If you are adding ports to an existing hunt group, disable the ports you are adding.

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

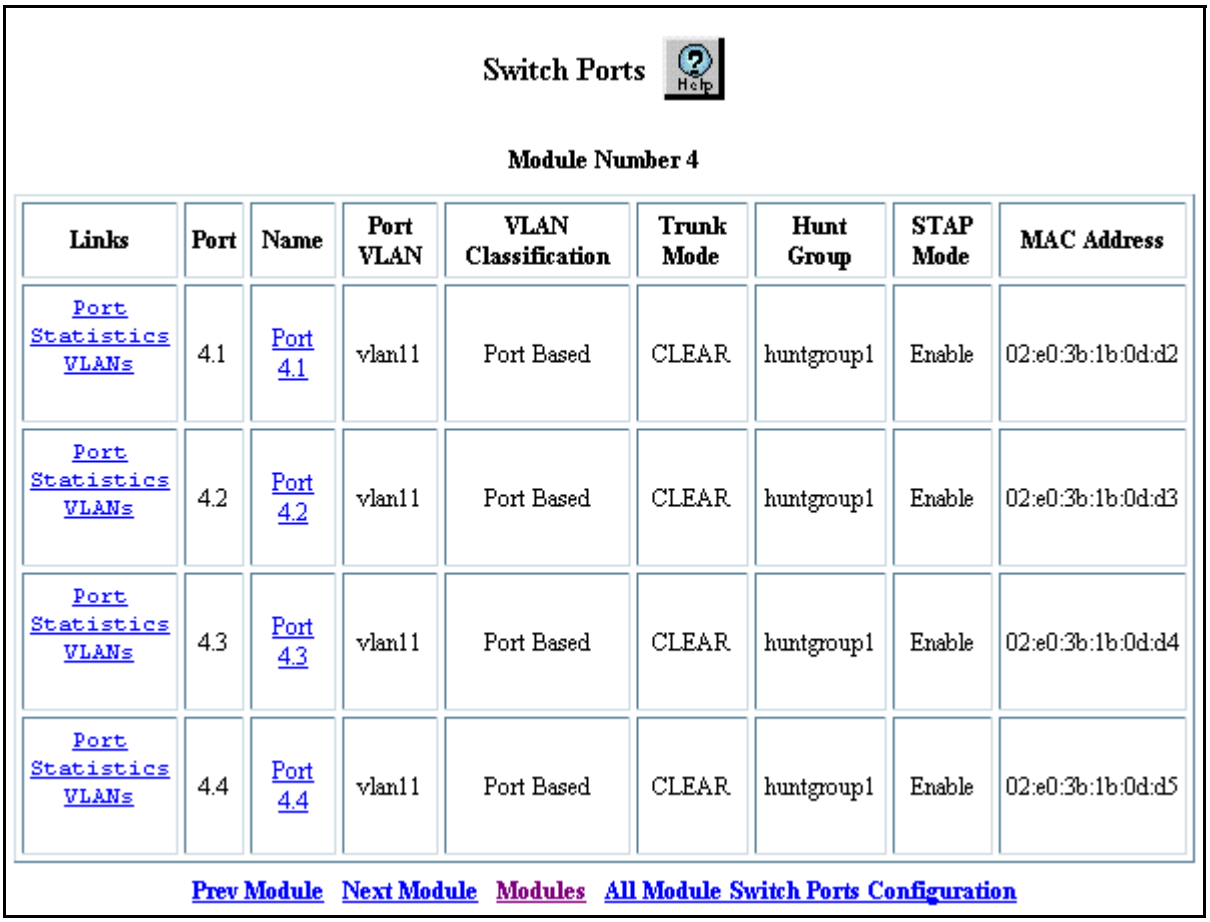
Adding Switch Ports to the Hunt Group

1. Select **Modules** at the bottom of the Physical Port Configuration Web page. The Module Information Web page is redisplayed. See Figure 6-11.

   As an alternate procedure, expand the **Modules & Ports** folder, and then click **Configuration**.

2. Select a number from the **Switch Ports** column, for the module whose port or ports you are adding to the hunt group. The Switch Ports Web page is displayed in the content pane. See Figure 6-13.

Figure 6-13. Switch Ports Web Page



Switch Ports

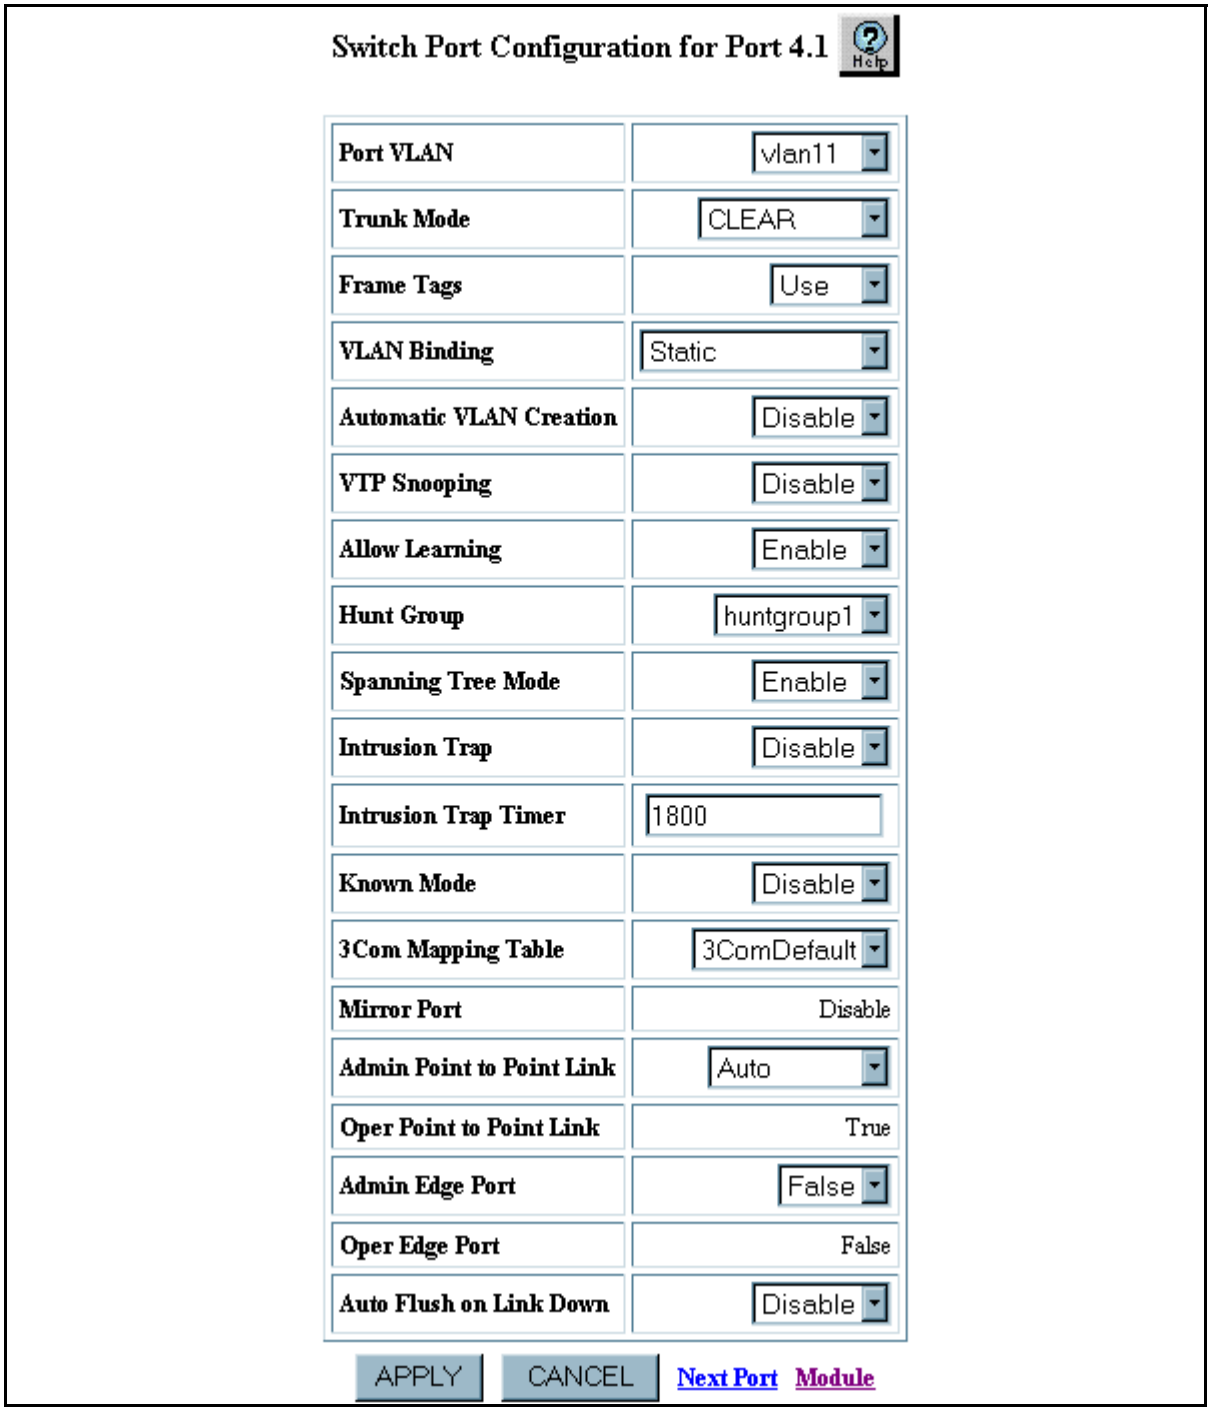
Module Number 4

Links	Port	Name	Port VLAN	VLAN Classification	Trunk Mode	Hunt Group	STAP Mode	MAC Address
Port Statistics VLANs	4.1	Port 4.1	vlan11	Port Based	CLEAR	huntgroup1	Enable	02:e0:3b:1b:0d:d2
Port Statistics VLANs	4.2	Port 4.2	vlan11	Port Based	CLEAR	huntgroup1	Enable	02:e0:3b:1b:0d:d3
Port Statistics VLANs	4.3	Port 4.3	vlan11	Port Based	CLEAR	huntgroup1	Enable	02:e0:3b:1b:0d:d4
Port Statistics VLANs	4.4	Port 4.4	vlan11	Port Based	CLEAR	huntgroup1	Enable	02:e0:3b:1b:0d:d5

Prev Module Next Module Modules All Module Switch Ports Configuration

3. Select the name of the port you want to configure from the **Name** column. The Switch Port Configuration Web page for that switch port is displayed in the content pane. See Figure 6-14.

Figure 6-14. Switch Port Configuration Web Page



4. Select the hunt group assignment from the **Hunt Group** pull-down menu.

5. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

6. Repeat Steps 1-9 for any additional ports that you want to add to this hunt group.

7. Repeat steps 1 through 10 on the switch at the other end of the hunt group connection.

8. In the navigation pane, expand the **Modules & Ports** folder, and then click **Configuration**. The Module Information Web page is redisplayed. See Figure 6-11.

9. Select a port number from the **Ports** column. The Physical Port Configuration Web page is displayed. See Figure 6-12.

10. Select a checkbox from the **Enable** column for that port. The group can now function as a load-sharing connection.

11. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

*** Note:** If thousands of addresses have been learned on a port and a link in the hunt group goes down, the switch-over of traffic between ports may take several seconds.

*** Note:** If you use static VLAN binding and hunt groups, make sure to first bind all the VLANs to all the ports that you will use in the hunt group and then apply the hunt group to those ports.

Viewing Hunt Group Members

You can view hunt group members from either the Web Agent or the CLI.

Web Agent Procedure

To view details about hunt group members using the Web Agent:

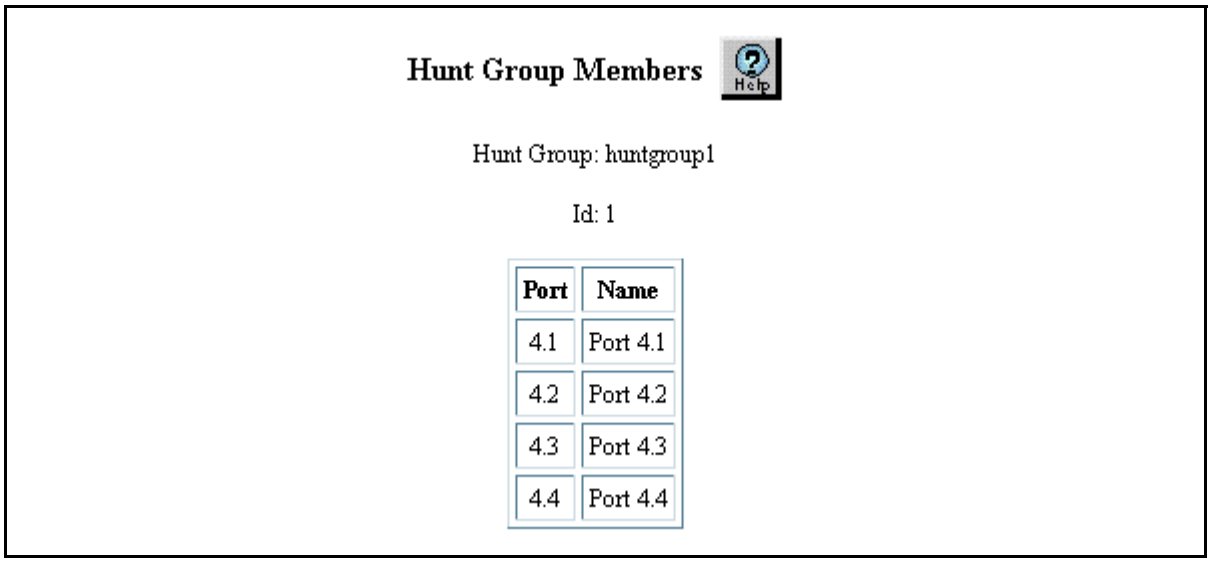
1. In the navigation pane, expand the **L2 Switching** folder, and then click **Hunt Groups**.

   The Hunt Group Configuration Web page is displayed in the content pane. See Figure 6-9.

2. Select a underlined number in the **Members** column.

   The Hunt Group Members Web page is displayed in the content pane and displays information about the members of that hunt group (Figure 6-15).

Figure 6-15. Hunt Group Members Web Page



CLI Commands

To view details about hunt group members using the CLI, enter the following command in configuration mode:

`(configure)#` **show huntgroup** <***huntgroup-name***>

To add ports to a hunt group using the CLI, enter the following command from Configure mode:

`(configure)#` **set port huntgroup {**<***mod-num***> **|** <***mod-swport-range***>**} [...,{**<***mod-num***> **|** <***mod-swport-range***>**}]** <***huntgroup-name***>

Removing Ports from Hunt Groups

You can remove a port from a hunt group by using either the Web Agent or the CLI.

Web Agent Procedure

To remove a port from a hunt group by using the Web Agent:

1. Open the Switch Port Configuration Web page for the port that you want to remove.
2. Click **None** in the **Hunt Group** field.
3. Click **Apply**.

CLI Command

To remove a port from a hunt group, use the following CLI command:

`(configure)#` **clear port huntgroup** <***mod-swport***>

Configuring VTP Snooping

Enabling VLAN Trunk Protocol (VTP®) Snooping on the Avaya Multiservice switch allows it to automatically synchronize its VLAN configuration with that of a Cisco VTP server switch. VTP is a Cisco® layer 2 protocol used to maintain VLAN configuration consistency among switches. Both switches must be on the same network, have at least VLAN 1 (Default) bound to a connecting link that is forwarding according to the Spanning Tree Protocol, and both ends of the link have the same Trunk Mode setting. Valid Trunk Mode options that work with VTP Snooping are Cisco Inter-Switch Link (Multi-layer) and IEEE 802.1Q.

VLAN additions, deletions, and name changes made on the network's Cisco VTP server are automatically updated on Avaya Multiservice switches that have VTP Snooping enabled and are connected to the Cisco VTP server with the same VTP Domain name. VLAN changes made on an Avaya Multiservice switch are not automatically updated on any other switch. Please note that VTP- learned VLANs may not be modified or deleted on an Avaya Multiservice switch while VTP is enabled. This restriction is in place to help maintain VLAN configuration consistency among VTP Snooping switches within the VTP Domain. Also the Avaya Multiservice switch handles VLAN name or VLAN ID conflicts between itself and a VTP server by giving locally configured VLANs precedence. For example if a VLAN ID or VLAN name is received in a VTP message that is also configured locally on an Avaya Multiservice switch, the VTP message's VLAN information is ignored for that VLAN and a message is entered in the switch's event log indicating the condition.

*** Note:** When VTP Snooping creates VLANs on the Avaya Multiservice switches, the hash table size for the VLANs is automatically set to 64. Depending on the number of VLANs that are created, you may need to manually change this hash table setting. For information on the relationship between hash table sizes and the number of VLANs that are configured on the switch, see “Relationship between VLANs, AFT and Hash Table Sizes,” in Chapter 9, “Managing the Address Forwarding Table.”

VTP Snooping is **Disabled** by default. You only need to change VTP Snooping port settings if you want to Enable its ability to learn VLAN changes from a Cisco VTP server.

You can configure VTP snooping from either the Web Agent or the CLI.

Web Agent Procedure

To configure VTP snooping using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Configuration**. The Module Information Web page is displayed in the content pane. See Figure 6-11.

2. Select the number in the **Switch Ports** column for the module you want to configure VTP. The Switch Ports Web page for that module is displayed. See Figure 6-13.

3. Select the **port** in the **Name** column on which you want to enable VTP Snooping throughout the switch. The Switch Port Configuration Web page is displayed for that port. See Figure 6-14.

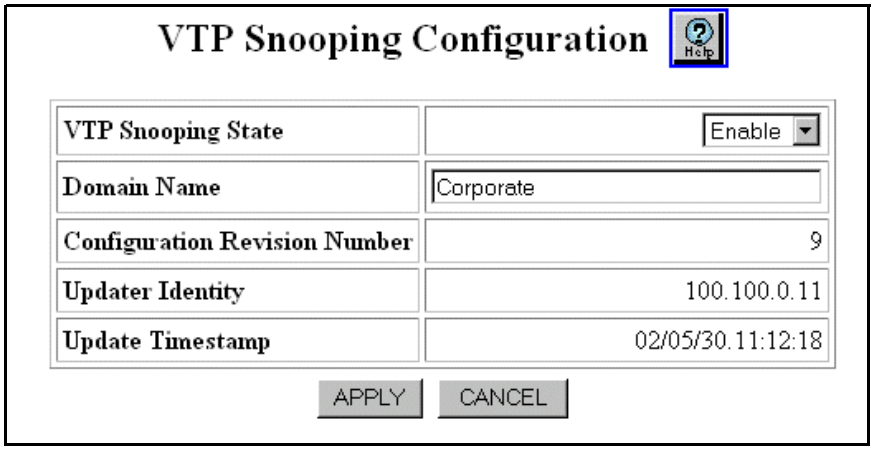
4. Select either **IEEE 802.1Q** or **Multi-layer** from the **Trunk Mode** pull-down menu to match the setting of the trunk mode for the switch port at the other end of the link.

5. Select **Enable** from the **VTP Snooping** pull-down menu. This enables VTP snooping on the switch port.

6. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

7. In the navigation pane, expand the **L2 Switching** > **VLANs** folders, and then click **VTP Snooping**. The VTP Snooping Configuration Web page is displayed in the content pane. See Figure 6-16.

Figure 6-16. VTP Snooping Configuration Web Page



8. See Table 6-4 and configure the **VTP Snooping** parameters.

Table 6-4. VTP Snooping Parameters

Parameter	Definition
VTP Snooping State	Select Enable to enable VTP snooping globally for the switch. The default value is **Disable**. **Note:** Enabling or disabling VTP Snooping does not clear any learned VTP information.
Domain Name	Enter the name associated with the Cisco VTP Domain. The default is Null (not set). Changing this parameter automatically clears the learned VTP information (the remaining parameters in this table). **Note:** The domain name is automatically learned from a Cisco VTP switch provided both the Domain Name is Null and the VTP Snooping State is enabled on the switch.
Configuration Revision Number	Displays the VTP snooping configuration revision number associated with the last successful VTP configuration update on the switch. **Note:** VLANs are only be learned by VTP Snooping when a received VTP message has a Configuration Revision Number greater than this learned value.
Updater Identity	Displays the IP address of the Cisco switch that initiated the VTP configuration update.
Update Timestamp	Displays the date and time that the Cisco switch initiated the VTP configuration update. The format of the timestamp is "yy/mm/dd.hh:mm:ss", where yy/mm/dd represents the year, month, and day and hh:mm:ss represents the hours, minutes, and seconds.

9. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To configure VTP snooping using the CLI, enter the following command in Configure mode:

```
(configure)# set vtp-snooping enable
```

7 Configuring Rapid Spanning Tree

Overview

Contents

This chapter contains the following sections:

- How RSTP Achieves Rapid Recovery
- Spanning Tree Configurations
- Configuring Global Spanning Tree Options
- Viewing Global Spanning Tree Information
- Configuring Spanning Tree Bridges
- Viewing Bridge Port Information
- Configuring a Bridge Port
- Configuring Spanning Tree Settings for Switch Ports

Purpose of Spanning Tree

STP prevents loops in layer 2 networks that have redundant paths. Loops in layer 2 networks can result in duplicate frames and switches learning the same MAC addresses on multiple interfaces.

To prevent loops, STP calculates a single path through the network and then blocks any redundant paths. STP first identifies a root switch and then identifies the most efficient path from the root switch to each switch in the network. Any redundant paths that exist enter a blocked state.

If a network segment in the spanning tree becomes unreachable or if you change the cost of a path, Spanning Tree recalculates a path through the network and then activates the new best path.

RSTP vs. STP

In application software v6.0 and later, the Rapid Spanning Tree Protocol (RSTP) replaces the 802.1D Spanning Tree Protocol (STP). STP is relatively slow at recovering from a failure in the network. RSTP was created to decrease this recovery time.

When a switch is running RSTP, a port can change from blocking to forwarding more quickly than when the switch is running STP. This decrease in transition time makes it possible for RSTP to recover more quickly from failures in the network.

In STP, a port that is blocking must change to listening and then learning before it can change to forwarding. Normally, a port remains in the listening state for 15 seconds and in the learning state for 15 seconds. With these two 15-second delays, a port normally takes 30 seconds to change from blocking to forwarding. This 30-second transition time results in a 30-second loss of traffic, which is not acceptable in many of today's networks.

Port States

STP has four different port states: listening, learning, blocking, forwarding, and disabled. RSTP has combined three of these states— listening, blocking, and disabled—into a single port state: discarding.

Interoperability with STP

RSTP is completely interoperable with STP. Switches running RSTP and switches running STP can be on the same network and work together to detect and break loops in the network.

Interoperability is achieved by the ability of RSTP to detect the presence of bridges running STP and to operate in common-spanning-tree mode. When an RSTP bridge is connected to an STP bridge and receives STP bridge protocol data units (BPDUs), the RSTP bridge sends only STP BPDUs out the port that is connected to the STP bridge. If a bridge running STP receives RSTP BPDUs, the STP bridge drops the RSTP BPDUs.

You can also manually set the P580 or P882 to run common Spanning Tree.

Upgrading to RSTP

To ensure that your network is not disrupted when you upgrade the switch from v5.*x* application software to v6.0 or later, RSTP has the following settings:

- The default setting for Spanning Tree version is common Spanning Tree.
- Switch ports for which fast start was enabled are set as edge ports. In v6.0 and later, the fast start feature is replaced with the ability to set switch ports as edge ports. For information on setting switch ports as edge ports, see "Configuring Spanning Tree Settings for Switch Ports" later in this chapter.
- All settings for port priorities and path costs are preserved.

How RSTP Achieves Rapid Recovery

Port Roles

In addition to port states, RSTP assigns and maintains port roles for all ports in a Spanning Tree domain. One of five possible roles can be assigned to a port: root, designated, alternate, backup, and disabled.

Root and designated ports are the only ports that actively participate in the spanning tree (by forwarding frames). Alternate and backup ports are blocked, but if a failure occurs in the network, they will will rapidly change to root or designated if necessary.

Alternate and backup ports are essential to RSTP's rapid recovery from port failures. An alternate or backup port can forward traffic immediately. If a port fails in an STP network, however, a 30-second loss of traffic occurs while STP recalculates the Spanning Tree topology.

Root ports provide the lowest cost path to the root bridge. Each bridge in the Spanning Tree domain has a root port that forwards frames to the root bridge. The root bridge does not have a root port.

Designated ports provide the lowest cost path from a network segment to the root bridge. Each network segment has one designated switch on which one port is designated. All traffic sent to and from the network segment passes through the designated port.

Alternate ports provide an alternate path in the direction of the root bridge. If the root port on the bridge fails, one of the alternate ports quickly changes to forwarding state.

Backup ports provide a backup path in the direction of the leaves of the spanning tree. If a designated port on a LAN segment fails, then one of the backup ports on that LAN segment quickly assumes the role of designated port for the segment.

Because RSTP maintains this information, it is able to more quickly activate a redundant path.

Rapid Transition to Forwarding

Two new RSTP port settings, **Edge Port** and **Point to Point Link**, make rapid transition to the forwarding state possible. A *point-to-point link* is connected to exactly one other bridge (normally with a direct cable between them). An *edge port* is not connected any other bridge.

Edge ports and ports that are connected to point-to-point links can change from discarding directly to forwarding.

IEEE Standards

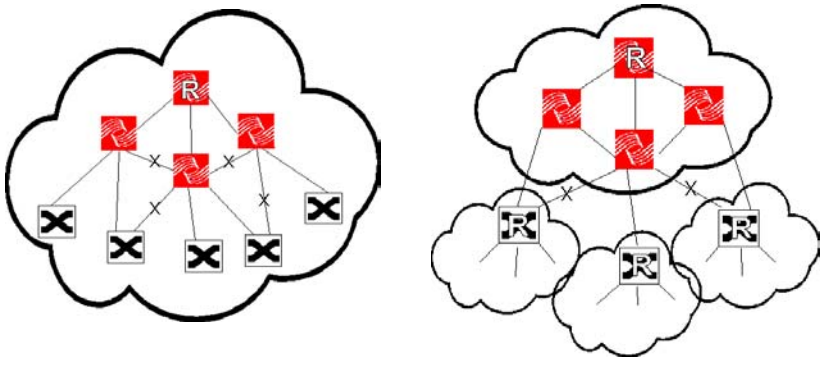
For more detailed information about the STP and RSTP, see the IEEE 802.1D standard for Media Access Control (MAC) bridges and IEEE802.1w amendment for rapid reconfiguration.

Spanning Tree Configurations

The Avaya Multiservice switches support the following Spanning Tree configurations:

- IEEE 802.1D Spanning Tree
- Per-VLAN Spanning Tree
- Dual-Layer Spanning Tree (Figure 7-1)
- Global Disable

Figure 7-1. Spanning Tree Models



Single 802.1D Spanning Tree
One Spanning Tree
Longer convergence
One path to and from root for all VLANs
Improper configuration
can shut down Trunk Links

Multi-Level Spanning Tree
Backbone terminates 802.1D STP
Smaller STP Domains
Quicker Convergence
VLAN Load Balancing
Interoperates w/ existing Bridge/Routers

IEEE802.1D Spanning Tree

All Avaya Multiservice switches participate in a Single Spanning Tree domain in the IEEE802.1D STP mode. All ports with STP configured belong to the same spanning tree domain and rules are as defined in IEEE802.1D. BPDUs are as defined by 802.1D and are sent out Clear on each link regardless of whether or not the link has a tagging method defined. As documented in the IEEE specification, 802.1 D Spanning Tree is intended for environments where only one VLAN is used in the network. If you are using 802.1 D Spanning Tree in the network and have multiple VLANs, you should set the P580 or P882 switches to run Dual-Layer Spanning Tree.

Per-VLAN Spanning Tree

Per-VLAN mode is the Default Spanning Tree setting on all multiservice switches. In this mode, the switch runs a separate spanning process for each VLAN. Each logical Spanning Tree has its own BPDUs which are tagged with the appropriate VLAN Identifier. In this configuration, the switch can participate in as many Spanning Tree domains as there are VLANs defined on the switch. This conforms to a “virtual” bridging model where the switch runs as if each VLAN is a separate logical bridge (separate Address Forwarding Tables, separate spanning trees, etc.). If spanning tree is not required on all VLANs, you can disable it for individual VLANs on a VLAN by VLAN basis. Also, different root bridges can be configured with different bridge priorities based upon VLAN. This will allow load sharing to occur based upon VLAN. Similarly link costs and priorities can be adjusted on a per-VLAN basis allowing further load sharing per VLAN.

Per VLAN Spanning Tree is the recommended method whenever many VLANs are passed over interswitch links, but not every VLAN resides on every interswitch link. If IEEE 802.1D Spanning Tree mode is used without care under these constraints, VLANs can be split due to blocking on inappropriate links. The only time that Per-VLAN Spanning Tree may become undesirable is when you have more VLANs to manage than the number of Spanning Tree Domains. The only other constraint is to limit the number of outgoing BPDUs.

It is required that the number of outgoing BPDUs be less than 500/second on all multiservice switches. This means that if you have 50 VLANs you cannot have more than 20 possible paths back to the root bridge (50 VLANS X 20 LINKS X 0.5 BPDU/sec. = 500), or if you have 100 VLANs you cannot have more that 10 possible paths back to the root bridge. These constraints are not generally exceeded in real networks. If these limits are exceeded, you must use 802.1D Spanning Tree mode.

This method can easily interoperate with legacy IEEE802.1D devices. The legacy devices become part of the VLAN associated with the VLAN port binding with which they are attached. In other words, if three P580 switches are in a network running VLANA and VLANB, and three legacy devices connect up through a port configured in VLAN A, the legacy devices will become part of the VLAN A Spanning Tree. The P580 switches will participate in two Spanning Tree domains, one for VLAN A containing three bridges and one for VLAN B containing 6 bridges. The legacy devices need not be VLAN aware.

Dual-Layer Spanning

Dual-Layer spanning tree mode is a variation of per-VLAN spanning tree mode with many of the same features. However, instead of using normal BPDUs, which are clear (free of VLAN tags) on clear links and tagged on tagged links, as is the case with Per-VLAN Spanning Tree mode, dual-layer uses a proprietary BPDU. This proprietary BPDU is sent to a special multicast address and contains information about which VLAN the BPDU is associated with. This has an advantage over per-VLAN spanning trees in that this method can support multiple VLANs over a non-tagging link, or when connecting to a bridge/router.

In this mode, legacy bridges remain in separate Spanning Tree domains, yet loops between the layer 3 and legacy domains cannot form. For example; if there are three P580 switches in a network running VLAN A and VLAN B, and three legacy devices connect up through a port configured in VLAN A, the legacy devices remain in their own legacy Spanning Tree and do NOT join the Multilayer Spanning Tree.

In the previous example, all three multiservice switches participate in two Spanning Tree domains:

- Domain for VLAN A containing three bridges
- Domain for VLAN B containing three bridges with the legacy domain remaining separate.

The legacy devices need not be aware of the other spanning tree domains. If a loop forms between the Multilayer domains and the legacy domain, the proprietary BPDU is seen returning to the Multilayer domain through the legacy domain and the loop is blocked. The Dual-Layer Spanning Tree method is preferred when inter-operating with large legacy bridge networks. Convergence time is reduced, and management is simplified, by allowing the legacy bridges to remain in their own spanning tree domain

Global Disable

You can globally disable Spanning Tree on all multiservice switches, thus preventing any switch port from participating in Spanning Tree. This is only recommended in a very controlled environment where there is no possibility of a loop being placed in the network.

Spanning Tree Design

Avaya recommends that you consider the following standard IEEE802.1D Guidelines when you design your network, no matter which Spanning Tree mode you decide to use:

- Ensure that core switches can be forced to be root bridges. For per-VLAN Spanning Tree, the function of root bridge can be shared among several core switches
- When using standard timers, ensure that the depth of the tree does not exceed 7 bridges when a default bridge timer is used

Configuring Global Spanning Tree Options

Overview

Global spanning tree options include the configuration and version of spanning tree (common STP or Rapid Spanning Tree).

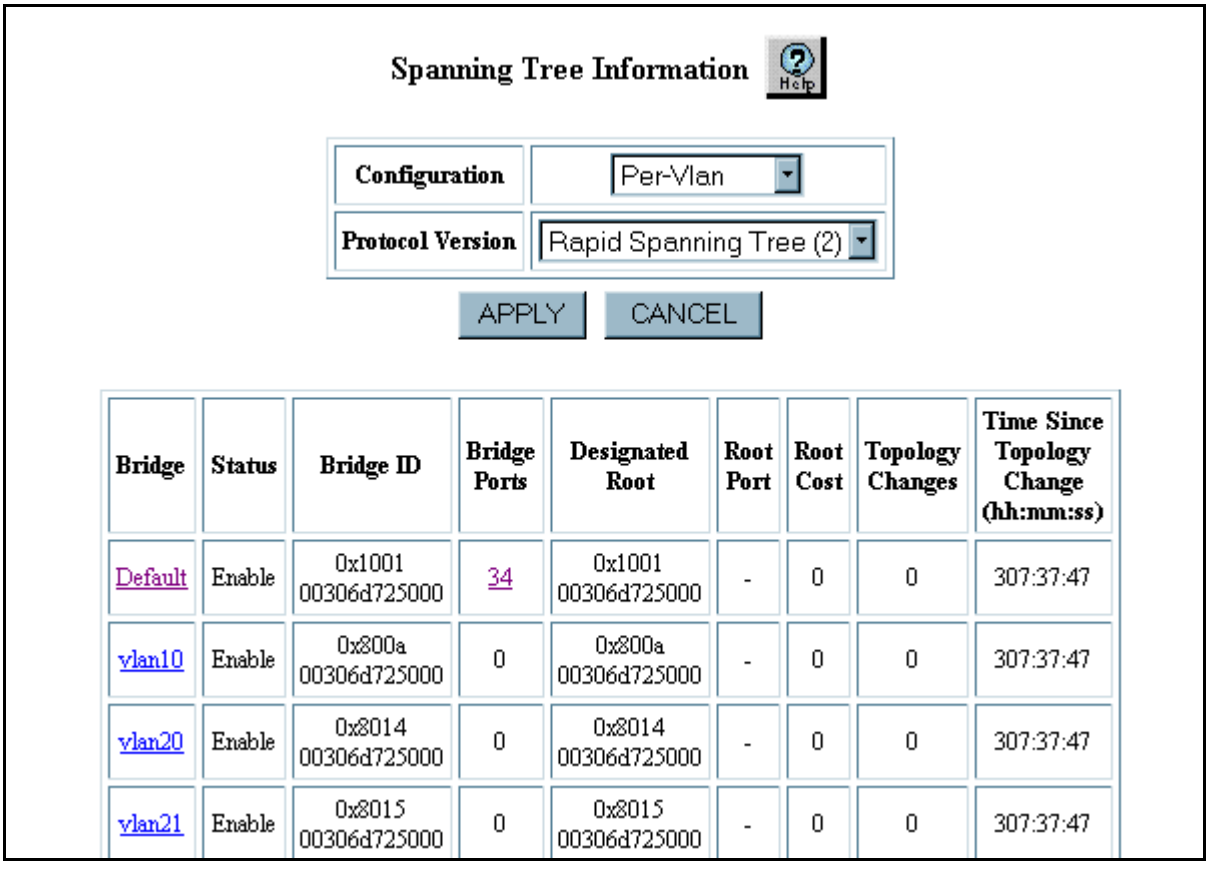
Web Agent Procedure

To configure global spanning tree options on the switch:

1. In the navigation pane, expand the **L2 Switching** folder.
2. Click **Spanning Tree**.

   The Spanning Tree Information Web page is displayed in the content pane. See Figure 7-2.

Figure 7-2. Spanning Tree Information Web page



Bridge	Status	Bridge ID	Bridge Ports	Designated Root	Root Port	Root Cost	Topology Changes	Time Since Topology Change (hh:mm:ss)
Default	Enable	0x1001 00306d725000	34	0x1001 00306d725000	-	0	0	307:37:47
vlan10	Enable	0x800a 00306d725000	0	0x800a 00306d725000	-	0	0	307:37:47
vlan20	Enable	0x8014 00306d725000	0	0x8014 00306d725000	-	0	0	307:37:47
vlan21	Enable	0x8015 00306d725000	0	0x8015 00306d725000	-	0	0	307:37:47

3. In the **Configuration** field, click the type of spanning tree you want the switch to run.

   See Table 7-1 and "Spanning Tree Configurations" earlier in this chapter for more information on the three types of spanning tree.

4. In the **Protocol Version** field, click the version of Spanning Tree that you want the switch to run.

5. Click **APPLY** to save your changes or **CANCEL** to restore the previous settings.

Table 7-1. Spanning Tree Information Configuration

Option	Definition
IEEE 802.1D	Runs a single IEEE 802.1D-compliant spanning tree for the entire bridge. In 802.1D spanning tree mode, there is one root bridge for the whole network, regardless of whether the switches support VLANs or not. Only one active data path is supported for all VLANs from any point in the network to any other point **Note:** When the Spanning Tree mode is set to IEEE 802.1D, bridge protocol data units (BPDUs) are sent out ports in Clear (non-tagged) format even if the port has a tagged format (3Com, IEEE 802.1Q or Cisco ISL) defined.
Per-VLAN	(Default Switch Setting) Runs a separate IEEE 802.1D-compliant spanning tree for each VLAN. In Per-VLAN spanning tree, the switch implements a separate spanning tree domain for each VLAN. Each logical spanning tree has its own Per-VLAN BPDUs, which are tagged with the appropriate VLAN identifier(s).BPDUs are sent clear (untagged) on clear trunks or tagged on tagged trunks. With Per-VLAN spanning tree, a switch can participate in as many Spanning Tree Domains as there are VLANs defined on the switch.This allows for better use of links. Some may be blocked for one VLAN STP and forwarding traffic for another.
	1 of 2

Table 7-1. Spanning Tree Information Configuration

Option	Definition
Dual-Layer	Dual-Layer spanning tree is a variation of per-VLAN spanning tree. However, instead of using normal BPDUs which are clear (free of VLAN tags) on clear trunks and tagged on tagged trunks, as is the case with Per-VLAN Spanning Tree. Dual-layer uses a proprietary BPDUs. These proprietary BPDUs are sent to a special multicast address and contain information about which VLAN the BPDU is associated with.In the Dual-Layer Spanning Tree model, the switch terminates all 802.1D Spanning Tree Domains. The switch does not forward 802.1D BPDUs or participate in any 802.1D Spanning Tree Domains. It only participates in Spanning Tree Domains using a proprietary BPDU, which contains a VLAN identifier (VLAN ID). This proprietary Spanning Tree Protocol will resolve any loops in the switch. Legacy Bridges remain in separate Spanning Tree domains yet Loops between the Dual Layer Domains and Legacy domains cannot form. The Dual-Layer Spanning Tree method is preferred when inter-operating with large legacy bridge networks. Convergence time is reduced, and management is simplified, by allowing the legacy bridges to remain in their own spanning tree domain.
Disable	Globally disables Spanning Tree for the entire switch. **Note:** You can also disable Spanning Tree on an individual port. See "Configuring Spanning Tree Settings for Switch Ports" later in this chapter.
	2 of 2

CLI Command

Use the following CLI commands to configure global Spanning Tree options:

- To set the type of spanning tree that you want the switch to run, `(configure)#` **set spantree config {ieee | per-vlan | dual-layer | disable}**
- To set the version of spanning tree that you want the switch to run, `(configure)#` **set spantree version {common-spanning-tree | rapid-spanning-tree}**

Viewing Global Spanning Tree Information

Web Agent Procedure

To view global spanning tree information:

1. In the navigation pane, expand the **L2 Switching** folder.
2. Click **Spanning Tree**.

   The Spanning Tree Information Web page is displayed in the content pane. See Figure 7-2.

   Table 7-2 provides explanations of the information that is displayed in each field of this Web page.

Table 7-2. Global Spanning Tree Information

Field	Definition
Configuration	Type of Spanning Tree that the switch is running. Options are IEEE 802.1D, Per-VLAN, Dual Layer, and Disable.
Protocol Version	Version of Spanning Tree that the switch is running. Options are **Rapid Spanning Tree** and **common-spanning-tree.** The default setting is **common-spanning-tree**.
Bridge	A bridge instance. In the Figure 7-2, a STP bridge per-VLAN.
Status	Indicates whether a bridge is enabled or disabled.
Bridge ID	ID of the bridge.
Bridge Ports	Number of ports in this bridge.
Designated Root	Root bridge for this spanning tree.
Root Port	Bridge port used to access the root bridge.
Root Cost	Path cost to the root bridge.
Topology Changes	Number of topology changes that have occurred since the last system reset.
Time Since Topology Change (hh:mm:ss)	Number of hours, minutes, and seconds since the last topology change.

CLI Commands

Use the **show spantree config** or **show spantree version** CLI command to view the configuration and version of Spanning Tree that the switch is running.

You can enter these commands from User mode.

Configuring Spanning Tree Bridges

Web Agent Procedure

To configure a spanning tree bridge:

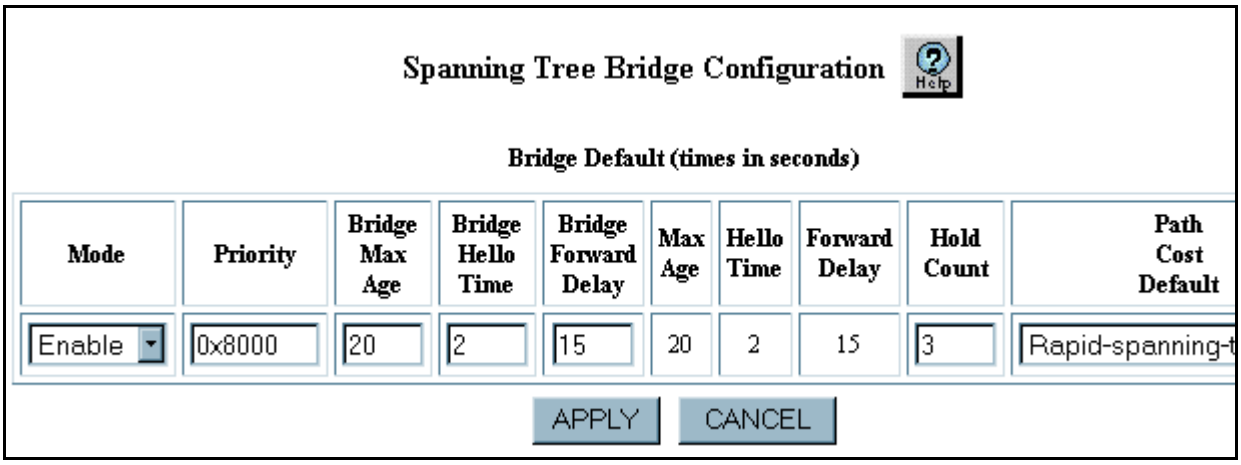
1. In the navigation pane, expand the **L2 Switching** folder.
2. Click **Spanning Tree**.

   The Spanning Tree Information Web page is displayed in the content pane. See Figure 7-2.

3. In the **Bridge** field, click the bridge that you want to configure.

   The Spanning Tree Bridge Configuration Web page is displayed. See Figure 7-3.

Figure 7-3. Spanning Tree Bridge Configuration Web page



4. Configure the bridge as appropriate.

   Table 7-3 provides explanations of each field.

Table 7-3. Spanning Tree Bridge Configuration

Field	Definition
Mode	Enables or disables Spanning Tree for the bridge.The default setting is Enabled
	1 of 3

Table 7-3. Spanning Tree Bridge Configuration

Field	Definition
Priority	Priority of the bridge as hexidecimal value. The valid range for this field is 0x0000 (0) to 0xF000 (61,440) in increments of 0x1000 (4,096). The default setting is 0x8000 (32,768). **Note:** When you upgrade the switch from v5.x application software to v6.x, all bridge priorities are reset to the default setting of 0x8000. Bridge priorities from earlier versions of software are not preserved.
Bridge Max Age	The maximum amount of time that the bridge retains bridging information. When the maximum age expires, the bridge assumes it has lost connection to the network and sends out requests to be readded to the spanning tree. The valid range for this field is 6 to 40 seconds. The default setting is 20 seconds.
Bridge Hello Time	The time between generation of BPDUs by the root bridge. The valid range for this field is 1 to 10 seconds. The default setting is 2 seconds.
Bridge Forward Delay	The time a port takes to change to the forwarding state. The valid range for this field is 4 to 30 seconds. The default setting is 15 seconds.
Max Age	Current maximum age for this spanning tree. The root bridge sets this time.
Hello Time	Current hello time for this spanning tree. The root bridge sets this time.
Forward Delay	Current forwarding delay for this spanning tree. The root bridge sets this time.
Hold Count	The maximum number of BPDUs that are sent out a port in a hello time interval. During any one hello time interval, no more BPDUs than the number that you enter in this field will be sent out a port. The valid range for this field is 1 to 10 seconds. The default setting is 3 seconds.
	2 of 3

Table 7-3. Spanning Tree Bridge Configuration

Field	Definition
Path Cost Default	The type of default path costs that ports in this bridge will use. Options are: • **common-spanning-tree**—uses the 16-bit default path costs from IEEE Std. 802.1D-1998: — For 10 MB ports, 100 — For 100 MB ports, 19 — For 1 GB ports, 4 — For 10 GB ports, 3 • **Rapid-spanning-tree**—uses the 32-bit default path costs from IEEE Std. 802.1t: — 10 Mbps port—2,000,000 — 100 Mbps port—200,000 — 1 Gbps port—20,000 — 10 Gbps port—2,500 **Note:** The switch must be running Rapid Spanning Tree to use the Rapid Spanning Tree default path costs. If the switch is running common Spanning Tree, it uses the common Spanning Tree default path costs regardless of the setting of this field.
	3 of 3

*** Note:** The Avaya Multiservice switches enforce the following relationships, as defined by IEEE 802.1D:

— 2 × (**Bridge Forward Delay** – 1 second) ≥ **Bridge Max Age**

— **Bridge Max Age** ≥ 2 × (**Bridge Hello Time** + 1 second)

5. Click **APPLY** to save your changes or **CANCEL** to restore the previous settings.

CLI Command

Use the following CLI commands to configure a spanning tree bridge:

- To enable or disable Spanning Tree, `(configure)#` **set spantree {enable | disable} {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}**
- To set the priority of the bridge, `(configure)#` **set spantree priority <*priority-value*> {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}**

- To set the maximum age for the bridge, `(configure)#` **set spantree maxage <*maxage-value*> {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}**
- To set the hello time for the bridge, `(configure)#` **set spantree hello <*hellotime-value*> {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}**
- To set the forward delay for the bridge, `(configure)#` **set spantree fwddelay <*fwddelay-value*> {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}**
- To set the hold count for the bridge, `(configure)#` **set spantree hold-count <*hold-count-value*> {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}**
- To set the default path costs for the bridge, `(configure)#` **set spantree default-path-cost {common-spanning-tree | rapid-spanning-tree} {802.1D | vlan {<vlan-id> | name <vlan-name>}}**
- To view the current configuration and status of the bridge, `>` **show spantree {all | 802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}**

Viewing Bridge Port Information

Web Agent Procedure

To view information about spanning tree bridge ports:

1. In the navigation pane, expand the **L2 Switching** folder.
2. Click **Spanning Tree**.

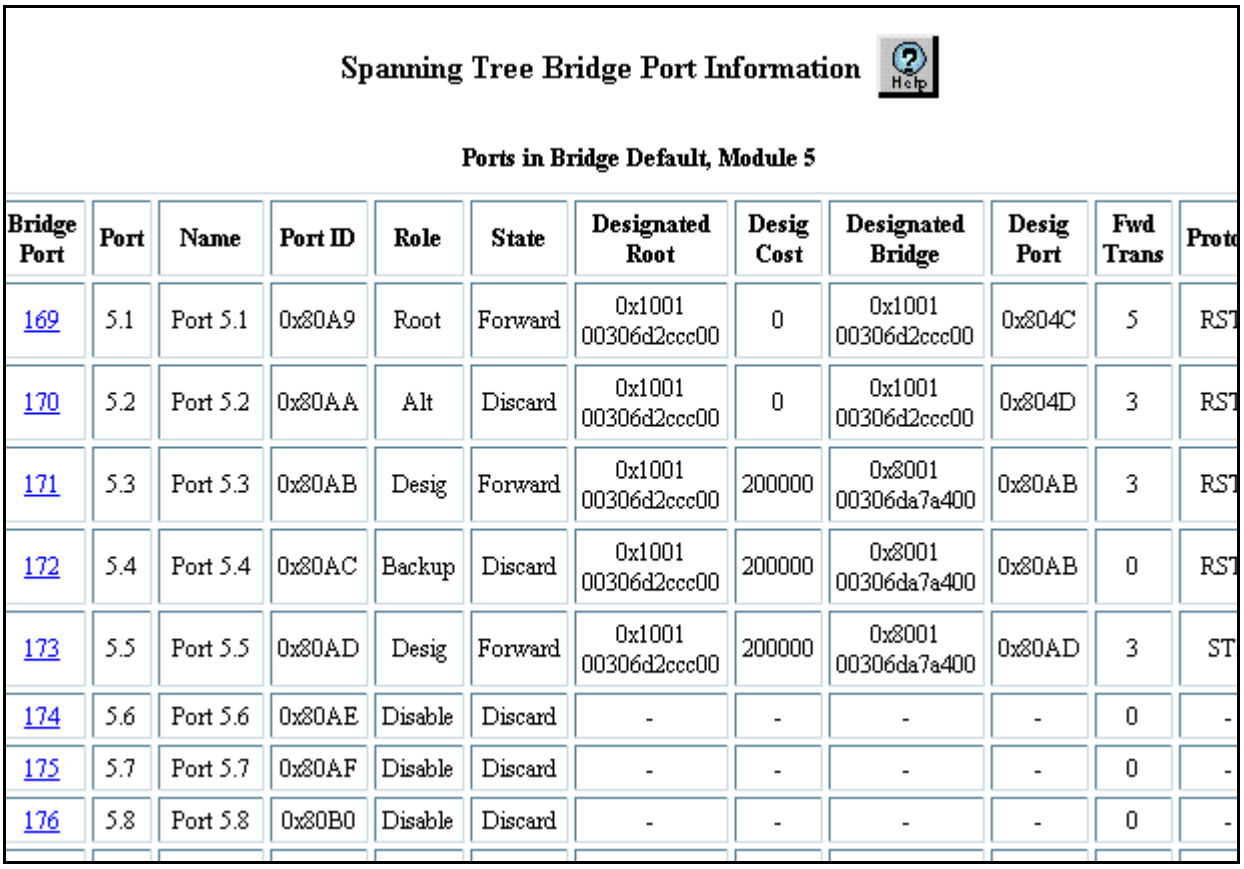
   The Spanning Tree Information Web page is displayed in the content pane. See Figure 7-2.
3. In the **Bridge Ports** field, click the bridge for which you want to view port information.

   The Spanning Tree Per Module Bridge Port Information Web page is displayed in the content pane.
4. In the **Bridge Ports** field, click the ports for which you want to view information.

   The Spanning Tree Bridge Port Information Web page is displayed. See Figure 7-4.

   Table 7-4 provides explanations of the information that is displayed in each field of this Web page.

Figure 7-4. Spanning Tree Bridge Port Information Web Page



Spanning Tree Bridge Port Information

Ports in Bridge Default, Module 5

Bridge Port	Port	Name	Port ID	Role	State	Designated Root	Desig Cost	Designated Bridge	Desig Port	Fwd Trans	Proto
169	5.1	Port 5.1	0x80A9	Root	Forward	0x1001 00306d2ccc00	0	0x1001 00306d2ccc00	0x804C	5	RST
170	5.2	Port 5.2	0x80AA	Alt	Discard	0x1001 00306d2ccc00	0	0x1001 00306d2ccc00	0x804D	3	RST
171	5.3	Port 5.3	0x80AB	Desig	Forward	0x1001 00306d2ccc00	200000	0x8001 00306da7a400	0x80AB	3	RST
172	5.4	Port 5.4	0x80AC	Backup	Discard	0x1001 00306d2ccc00	200000	0x8001 00306da7a400	0x80AB	0	RST
173	5.5	Port 5.5	0x80AD	Desig	Forward	0x1001 00306d2ccc00	200000	0x8001 00306da7a400	0x80AD	3	ST
174	5.6	Port 5.6	0x80AE	Disable	Discard	-	-	-	-	0	-
175	5.7	Port 5.7	0x80AF	Disable	Discard	-	-	-	-	0	-
176	5.8	Port 5.8	0x80B0	Disable	Discard	-	-	-	-	0	-

Table 7-4. Spanning Tree Bridge Port Information

Field	Explanation
Bridge Port	Bridge port number.
Port	Physical port number of the bridge port.
Name	Name that is assigned to the bridge port.
Port ID	The ID that Spanning Tree assigns the port.

1 of 2

Table 7-4. Spanning Tree Bridge Port Information

Field	Explanation
Role	The current role of the port. The options are: • Root • Designated • Alternate • Backup • Disabled For information about each of these roles, see "How RSTP Achieves Rapid Recovery."
State	Current bridging state of the port. The options are: • **Discarding** — The port is either disabled, blocking or listening. • **Learning —** The port is learning new MAC addresses, but not yet forwarding traffic. • **Forwarding —** The port has been selected by Spanning Tree to forward traffic and is forwarding traffic currently.
Designated Root	Root bridge for this spanning tree.
Designated Cost	The path cost to the designated root of the segment that is connected to this port.
Designated Bridge	ID of the designated bridge for this segment.
Designated Port	ID of the designated port on the designated bridge for this segment. For information about each of these roles, see "How RSTP Achieves Rapid Recovery."
Fwd Trans	Number of times that this port has changed from discarding state to forwarding state.
Protocol	Type of BPDUs that the port is transmitting (RSTP or STP). If the switch is running RSTP, and this field displays **STP**, the port has detected the presence of bridges running STP and is operating in common-spanning-tree mode.
	2 of 2

CLI Command

Use the **show spantree port {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}** command to view information about Spanning Tree bridge ports.

Configuring a Bridge Port

Web Agent Procedure

To configure a spanning tree bridge port:

1. In the navigation pane, expand the **L2 Switching** folder.
2. Click **Spanning Tree**.

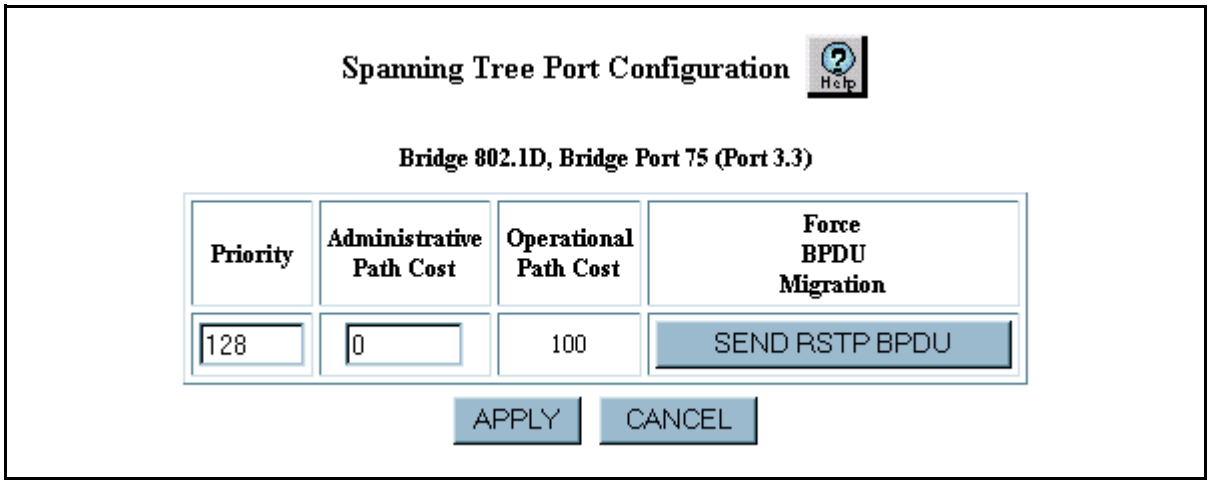
   The Spanning Tree Information Web page is displayed in the content pane. See Figure 7-2.
3. In the **Bridge Ports** field, click the bridge for which you want to configure a port.

   The Spanning Tree Per Module Bridge Port Information Web page is displayed in the content pane.
4. In the **Bridge Ports** field, click the ports that you want to configure.

   The Spanning Tree Bridge Port Information Web page is displayed. See Figure 7-4.
5. In the **Bridge Port** field, click the bridge port that you want to configure.

   The Spanning Tree Port Configuration Web page is displayed. See Figure 7-5.

Figure 7-5. Spanning Tree Port Configuration Web Page



6. Configure the bridge port as appropriate.

   Table 7-5 provides explanations of each field.

7. Click **APPLY** to save your changes or **CANCEL** to restore the previous settings.

Table 7-5. Spanning Tree Port Configuration

Field	Definition
Priority	Priority of the port as a decimal value. A higher priority port (has a lower priority number) is more likely to be chosen as the primary path in the spanning tree when there are two or more paths of equal cost. The valid range for this field is 0 to 240 in increments of 16. The default setting is 128.
Administrative Path Cost	Sets the path cost for this port. The ports that you prefer be used by the spanning tree should have the lowest path cost. If the switch is running common Spanning Tree, the valid range for this field is 0 to 65535. If the switch is running Rapid Spanning Tree, the valid range for this field is 0 to 200,000,000. The default setting is 0. If this field is set to 0, the port uses the default path cost for the bridge.
	1 of 3

Table 7-5. Spanning Tree Port Configuration

Field	Definition
Operational Path Cost	The path cost that Spanning Tree is using for the port: • If **Administrative Path Cost** is set to a nonzero value, this field displays that nonzero value. • If **Administrative Path Cost** is set to **0**, the switch is running *Rapid Spanning Tree*, and **Path Cost Default** for the bridge is set to **Rapid-spanning-tree**, then this field displays the following values (recommended in Table 17-7 of IEEE Std. 802.1w-2001). — 10 Mbps port—2,000,000 — 100 Mbps port—200,000 — 1 Gbps port—20,000 — 10 Gbps port—2,500 For *hunt groups*, this field displays the operational path cost of one link divided by the number of links in the hunt group. For example, if a hunt group comprises four 1-GB ports, and the operational path cost for one port is 20,000, the operational path cost for the hunt group is 5,000 (20,000 ÷ 4). • If **Administrative Path Cost** is set to **0** and the switch is running *common Spanning Tree*, then this field displays the following values (regardless of the **Path Cost Default** setting): — For 10 MB ports, 100 — For 100 MB ports, 19 — For 1 GB ports, 4 — For 10 GB ports, 3 **Note:** If the port is a trunk port (IEEE 802.1q or Multi-Layer tagging is enabled), the operational path cost is one less than the value in the preceding list. For *hunt groups*, this field displays the operational path cost of one link minus 1. For example, if a hunt group comprises four 1-GB ports, and the operational path cost for one port is four, the operational path cost of the hunt group is 3 (4 – 1).
	2 of 3

Table 7-5. Spanning Tree Port Configuration

Field	Definition
Force BPDU Migration	Clicking **SEND RSTP BPDU** in this field forces the bridge port to send out RSTP BPDUs. By forcing a bridge port to send RSTP BPDUs, you can determine whether legacy 802.1D bridges are present on a LAN segment. If you remove a legacy 802.1D bridge from a segment, other RSTP bridges on the segment cannot detect the removal so they continue sending STP BPDUs. However, if you force a bridge port to send RSTP BPDUs, they trigger other RSTP bridges on the segment to generate RSTP BPDUs again. If the switch is running common Spanning Tree, this command has no effect.
	3 of 3

CLI Commands

Use the following commands to configure a Spanning Tree bridge port:

- To set the priority of the bridge port, `(configure)#` **set port spantree priority <*mod-swport-range*> [...,<*mod-swport-range*>] <*bport-priority*> {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}**
- To set the path cost for the bridge port, `(configure)#` **set spantree portcost <*mod-swport-range*> [...,<*mod-swport-range*>] <*port-cost-value*> {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}**
- To force the bridge port to send one RSTP BPDU out the port, `(configure)#` **set port spantree force-protocol-migration <*mod-swport-range*> [...,<*mod-swport-range*>] {802.1D | vlan {<*vlan-id*> | name <*vlan-name*>}}**

Configuring Spanning Tree Settings for Switch Ports

Web Agent Procedure

To configure Spanning Tree settings for a switch port:

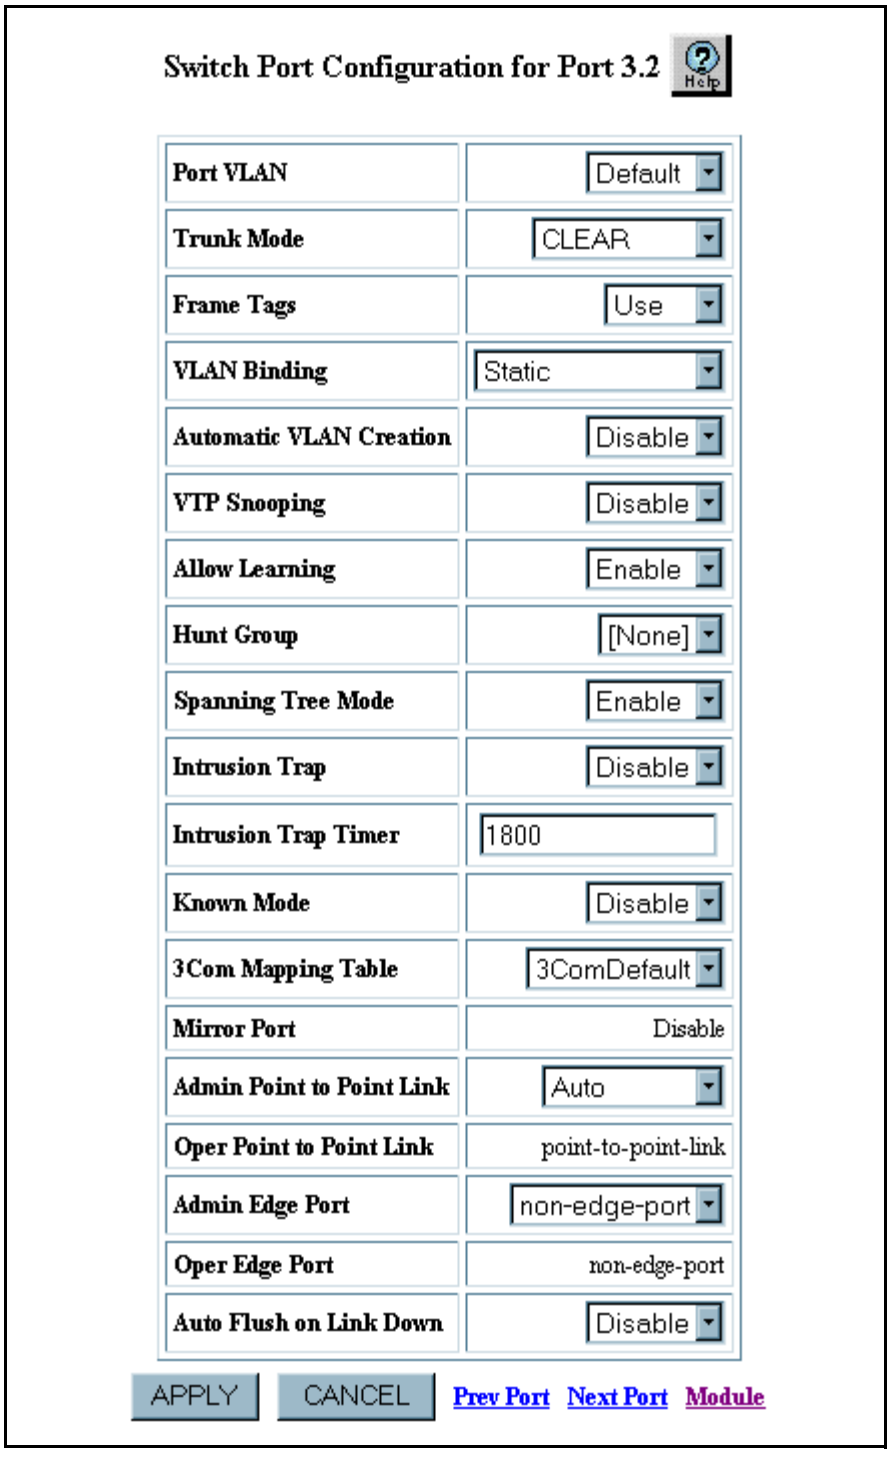
1. In the navigation pane, expand the **Modules & Ports** folder.
2. Click **Configuration**.

   The Module Information Web page is displayed in the content pane.
3. In the **Switch Ports** field, click the switch ports that you want to configure.

   The Switch Ports Web page is displayed in the content pane.
4. In the **Name** field, click the switch port that you want to configure.

   The Switch Port Configuration Web Page is displayed in the content pane. See Figure 7-6.

Figure 7-6. Switch Port Configuration Web Page



5. Enter the appropriate information in the following fields:

   - **Spanning Tree Mode**—Enable or disable Spanning Tree on this port. If you click **Disable**, the port does not participate in Spanning Tree.

   - **Admin Point to Point Link**—Specify whether this port is connected to a shared LAN segment or a point-to-point LAN segment. A point-to-point LAN segment is connected to exactly one other bridge (normally with a direct cable between them). Only point-to-point links and edge ports can rapidly transition to forwarding state.

     If you set this field to **Auto**, the switch automatically detects whether the port is connected to a shared link or a point-to-point link. Ports operating in half duplex are set to **False**, and ports operating in full duplex are set to **True**. You can, however, manually set the type of link. Options are:

     — **ForceTrue**—Defines the port as connected to a point-to-point link.

     — **ForceFalse**—Defines the port as connected to a shared LAN segment.

     — **Auto**—Automatically detects whether the port is connected to a shared link or a point-to-point link. If you select this setting, the **Oper Point to Point Link** field displays the link type that is detected.

   - **Admin Edge Port**—Specify whether this port is an edge port or a nonedge port. An edge port is not connected to any other bridge. Only edge ports and point-to-point links can rapidly transition to forwarding state. Options are:

     — **edge-port**—Defines the port as an edge port.

     — **non-edge-port**—Defines the port as a nonedge port.

     If you set this field to **edge-port**, the **Oper Edge Port** field is also set to **edge-port**. However, if the port receives a BPDU, the **Oper Edge Port** setting changes to **non-edge-port**. (To receive a BPDU, the port must be connected to a bridge and thus is not an edge port.)

     This field replaces the fast start feature in v5.*x* application software. When you upgrade the software from v5.*x* to v6.0 or later, switch ports for which fast start was enabled have **Admin Edge Port** set to **edge-port**.

6. Click **APPLY** to save your changes or **CANCEL** to restore the previous settings.

For information about all other switch port settings, see "Configuring Switch Ports Settings" in Chapter 8, "Configuring Ports."

CLI Commands

Use the following CLI commands to configure Spanning Tree settings for a switch port:

- To enable or disable Spanning Tree on a port, `(configure)#` **set port spanning-tree-mode {<*mod-num*> | <*mod-swport-range*>} [...,{<*mod-num*> | <*mod-swport-range*>}] {disable | enable}**
- To specify whether a port is connected to a shared LAN segment or a point-to-point LAN segment, `(configure)#` **set port point-to-point admin status {<*mod-num*> | <*mod-swport-range*>} [...,{<*mod-num*> | <*mod-swport-range*>}] {force-true | force-false | auto}**
- To specify whether this port is an edge port or a nonedge port, `(configure)#` **set port edge admin state <*mod-swport-range*> [...,<*mod-swport-range*>] {edge-port | non-edge-port}**

8 Configuring Ports

Overview

The following information and procedures are provided in this chapter and are common to both layer 2 and layer 3 module configuration:

- Two Categories of Port Settings
- Configuring Physical Port Settings
- Configuring Switch Ports Settings
- Configuring Port Redundancy
- GBIC Identification
- Network Error Detection and Recovery
- Internal Error Detection and Recovery

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.

Two Categories of Port Settings

The system has two categories of port settings:

- **Physical port settings** — Allows you to set up rules that guide the system's physical layer interaction (for example, enable/disable, speed, auto-negotiation).
- **Switch port settings** — Allows you to specify how the port participates in switching (for example, VLAN mode, trunking).

The following sections explain some of the Avaya P580 and P882 Multiservice switch features and how to configure the ports:

- Physical Port Features
- Switch Port Features
- Relationship Between Different Switch Port Parameters
- Assigning VLANs to a Port and Associated Issues

Physical Port Features

This section describes the following features that are available when configuring the physical parameters on the ports:

- Auto-negotiation
- Remote Fault Detection
- Flow Control

Auto-negotiation

Autonegotiation is an extension of the link test methods used by 10Base-T and 10Base-FL to verify the integrity of the link between devices. Autonegotiation advertises a device's abilities by encoding a 16-bit data packet, called a link code word (LCW), within a burst of 17 to 33 link pulses, called a fast link pulse (FLP) burst. FLP bursts have an approximate duration of 2 µs and are transmitted in 16.8 µs intervals (the same interval as for the normal link pulses used by 10Base-T and 10Base-FL).

However, this does not hold true for the Half Duplex/Full Duplex (HD/FD) selection. If a non-autonegotiating device running FD is connected to a negotiating device, the negotiating device runs at HD, and the link does not operate properly.

*** Note:** A Gigabit Ethernet device negotiates at HD or FD, speed is always one Gigabit.

*** Note:** M5520-TX (P/N M5520-100TX) boards manufactured with a Quality Phy do not auto-negotiate with Xircom brand adapter cards. If you have this problem, disable auto-negotiation on the affected ports, and set the port speed and duplex state manually.

*** Note:** You may experience difficulties with auto-negotiation between some releases of the 10/100Base-TX Module (M5510-100TX, M5520-100TX, M5510R-100TX, M5512R-100TX) and adapter cards using physical interfaces manufactured by National Semiconductor. The symptom is loss of connectivity. If you do. do one of the following:

- Disable auto-negotiation
- Use a patch cable that is longer that 5 meters

***Note:** The factory default for the National Phy Mode is **Enable**.

The LCW contains two fields (the selector field and the technology ability field), which together serve to identify a device's capabilities.

It may seem that because the FLP and the normal link pulse use the same interval at the same frequency, older devices may not be compatible with auto-negotiation. This is, however, not the case. For example, a 10Base-T device that does not have auto-negotiation capabilities sees FLP bursts simply as a link test signal. A 10Base-T device will respond to the FLP burst with its usual normal link pulse signal. At the other end of the link, a 10/100-capable device will recognize normal link pulse and choose 10Mbps mode operation.

Auto-negotiation attempts to find the greatest common denominator for the two devices on the link in the following order of preference:

1. 100Base-TX full-duplex

2. 100Base-T4

3. 100Base-TX

4. 10Base-T full-duplex

5. 10Base-T half-duplex

*** Note:** T4: 100 Mbps with 8B/6T coding scheme

Once the greatest common denominator of settings is determined, each device equipped with auto-negotiation will configure itself automatically. In certain cases where automatic configurations are not desired, auto-negotiation provides a way for these settings to be overridden manually.

*** Note:** Auto-negotiation should be disabled only on 50-series modules that have remote fault detection enabled. Do not disable auto-negotiation on 80-Series gig links.

Table 8-1 lists the gigabit modules that do not support auto negotiation:

Table 8-1. Gigabit Modules not Supporting Autonegotiation

Gigabit Module Model Number	Hardware Revision
M5502-1000SX-F	M or earlier
M5502-1000LX-F	M or earlier
M5502-1000SLX-F	F or earlier
M5504-1000SX-F	H or earlier
M5504-1000LX-F	H or earlier
M5504-1000SLX-F	H or earlier
	1 of 2

Table 8-1. Gigabit Modules not Supporting Autonegotiation

Gigabit Module Model Number	Hardware Revision
M5502R-1000SX-F	J or earlier
M5502R-1000LX-F	J or earlier
M5502R-1000SLX-F	H or earlier
	2 of 2

*** Note:** If a Gigabit module that does not support the autonegotiation is connected to a device that does, disable autonegotiation to ensure proper operation.

Remote Fault Detection

Remote fault detection is a proprietary feature that prevents a loss of traffic if a physical or signaling error occurs on a switch-to-switch fiber link.

A *remote fault* is an error that one switch can detect but the other switch cannot. For example, if a transmit fiber breaks, the remote port continues to receive data and so detects that the link is good. However, the remote port cannot detect that the data it transmits is not received by the local port.

If remote fault detection is enabled and a remote fault occurs, the local switch sends a message to inform the remote switch of the fault. The remote switch then shuts down the remote port.

If the receive signal is restored on the local port, the local port sends a message to the remote switch, which then turns on the remote port again.

***Important:** Remote fault detection must be enabled on both the local port and remote port for the feature to work.

For example, in Figure 8-1:

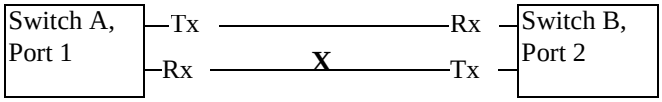
1. Port 1 on Switch A is not receiving traffic from Port 2 on Switch B.

   This problem could occur because:

   - The transmitter on port 2 is not functioning.
   - The receiver on port 1 is not functioning.
   - The fiber is broken or disconnected.

2. Switch A detects the error condition and sends a message to Switch B, which shuts down Port 2. (When the port is shut down, its status changes from **Okay** to **No Link**. The port status is displayed in the **Status** field of the Physical Port Configuration Web page.)

Figure 8-1. Remote fault detection



Enable remote fault detection on both ends of a switch-to-switch connection, in the following two cases:

- When two Gigabit ports that do not support auto-negotiation are connected.
- When a Gigabit port that supports auto-negotiation is connected to a Gigabit port that does not support auto-negotiation.

For a list of Gigabit modules that do not support auto-negotiation, see Table 8-1.

Restrictions:

- Autonegotiation and remote fault detection cannot be enabled at the same time. To enable remote fault detection, autonegotiation must be disabled. When autonegotiation is enabled, remote fault detection is automatically disabled.
- If the switch is operating in *Fabric mode 1*, 80-series Gigabit modules support remote fault detection on only one port. If you enable remote fault detection on more than one port, a loss of traffic occurs.

  If the switch is operating in *Fabric mode 2*, 80-series Gigabit modules support remote fault detection on two ports; however, the ports must be on different fabric ports. If you enable remote fault detection on two ports that are on the same fabric port, a loss of data occurs. For information on the relationship between fabric ports and physical ports, see "How Hunt Groups Load Share" in Chapter 6, "Using VLANs, Hunt Groups, and VTP Snooping."

  50-series modules support remote fault detection on any number of ports.

Flow Control

There are three flow control options on 10/100 Ethernet ports:

- Disable
- Enable
- Enable with Aggressive Backoff

There are four flow control options on Gigabit Ethernet Ports.

- Disable
- Enable
- Enable – Send Only
- Enable – Respond Only

The **Disable** option disables flow control (the default).

The **Enable** option enables IEEE802.3X (XOFF/XON) Flow Control on the line. The Avaya Multiservice Switch is the sender of pauses based upon inbound traffic and the Avaya Multiservice Switch is the receiver of pauses based upon outgoing traffic (and the ability of attached equipment to use the protocol).

With Gigabit links, you can:

- Enable IEEE802.3X for both send and receive (the “Enable” setting)
- Only send pause (XOFF/XON) signals (the “Enable – Send Only” option)
- Only respond to the pause signals (the “Enable – Respond Only” option).

With 10/100 links you can use the **Enable with Aggressive Backoff** option to enable Active Backpressure (creation of a collision) on a link. Active Backpressure only applies to HD links so that when a 10/100 link is in FD, options 2 and 3 are identical.

If you enable flow control, the switch manages the inbound buffers with flow control (IEEE 802.1X XOFF, or Backpressure) applied when a high water mark is reached. Thus, no matter which flow control option is chosen, outgoing pauses or backpressure are only applied to the port on which the parameter is set based upon the inbound traffic for that port in the default queuing mode.

Flow control is not applied across the switch itself. Therefore, if you have a server sending data on one port at 100MB and a client receiving the data at 10MB, the switch will not throttle the data at the Server’s input based on outgoing buffer backup on the client’s port in this default queuing mode.

Switch Port Features

The Avaya P580 and P882 Multiservice switches support up to 1000 VLANs and also support multiple forwarding databases. This means that each VLAN is associated with its own Address Forwarding Table (AFT). Therefore, identical MAC addresses can simultaneously exist on multiple VLANs. The Avaya Multiservice switches provide parameters for configuring VLAN/port associations.

This section provides the following:

- Relationship Between Different Switch Port Parameters
- Assigning VLANs to a Port and Associated Issues
- Setting the Port VLAN attribute in the Switch Port Configuration Web page. This identifies the VLAN to which all untagged frames received on the port are classified. Note that a port has exactly one Port VLAN. Changing this to a new VLAN removes the port from the old VLAN.
- Setting the VLAN Binding attribute in the Switch Port Configuration Web page to **Bind to All**. This causes the port to be bound to all VLANs known to the switch (now and in the future). This is not recommended when the switches on both ends of the link function as a router. This setting should be used on links that connect two Layer 2 switches.
- Setting the VLAN Binding attribute in the Switch Port Configuration Web page to **Bind to Received**. This causes the port to be bound to all VLANs identified by the VLAN tag in IEEE 802.1Q and Multi-Layer tagged frames received on this port. Consequently, ports are bound to those VLANs that actually have members that are reachable through the port.
- Manually creating a VLAN Switch Port via the Web Agent or CLI.

Relationship Between Different Switch Port Parameters

Table 8-2 shows the relationship between Port VLAN / Trunk Mode / VLAN Binding parameters when using the VLAN Operational Rules.

Table 8-2. Relationship Between Switch Parameters

VLAN Binding	Trunk Mode	Description
Static	Clear	**Ingress:** Untagged frames are classified to the VLAN associated with the port on which the frame is received. Tagged frames are classified to the VLAN identified by the VLAN tag in the tag header of the frame. **Forwarding:** Only forward frames to the port for the assigned VLAN. **Egress:** All frames transmitted will be sent with no tagging.
Static	802.1Q Multi-layer	**Ingress:** Untagged frames are classified to the VLAN associated with the port on which the frame is received. Tagged frames are classified to the VLAN identified by the VLAN tag in the tag header of the frame. **Forwarding:** Only forward frames to the port for the assigned VLAN. **Egress:** All frames transmitted out of the port will be tagged using the IEEE 802.1Q/Multi-Layer tag header format. The tag used will be that assigned to the port.
Bind to All	Clear	***NOT RECOMMENDED*** **Ingress:** Untagged frames are classified to the VLAN associated with the port on which the frame is received. Tagged frames are classified to the VLAN identified by the VLAN tag in the tag header of the frame. **Forwarding:** All broadcast frames from all VLANs will be forwarded to the port. **Egress:** All frames transmitted will be sent with no tagging.
Bind to All	802.1Q Multi-layer	**Ingress:** Untagged frames are classified to the VLAN associated with the port on which the frame is received. Tagged frames are classified to the VLAN identified by the VLAN tag in the frame's tag header. **Note:** For 802.1q mode: if a tagged frame is received, but the VLAN for that tagged frame does not exist on the switch, that frame will be placed onto the port VLAN assigned to the port. This may cause unicast and broadcast VLAN traffic from other VLANs to be seen on the port VLAN. To avoid this behavior, you can set the port VLAN into the "discard" VLAN which will drop all untagged frames and tagged frames with unknown VLAN IDs. For Multi-layer: if a tagged frame is received, but the VLAN for that tagged frame does not exists on the switch, that frame will be dropped. **Forwarding:** All broadcast frames from all VLANs will be forwarded to the port. **Egress:** All frames transmitted out of the port to be tagged using the IEEE 802.1Q/Multi-Layer tag header format. The tagged used will be that of the VLAN in which the frame was received.
		1 of 2

Table 8-2. Relationship Between Switch Parameters

VLAN Binding	Trunk Mode	Description
Bind to Received	Clear	***NOT RECOMMENDED*** **Ingress:** Untagged frames are classified to the VLAN associated with the port on which the frame is received. Tagged frames are classified to the VLAN identified by the VLAN tag in the tag header of the frame. **Forwarding:** All broadcast frames from all VLANs learned on the port will be forwarded. **Egress:** All frames transmitted will be sent with no tagging.
Bind to Received	802.1Q Multi-layer	**Ingress:** Untagged frames are classified to the VLAN associated with the port on which the frame is received. Tagged frames are classified to the VLAN identified by the VLAN tag in the frame's tag header. **Note:** For 802.1q mode: if a tagged frame is received, but the VLAN for that tagged frame does not exist on the switch, that frame will be placed onto the port VLAN assigned to the port. This may cause unicast and broadcast VLAN traffic from other VLANs to be seen on the port VLAN. To avoid this behavior, you can set the port VLAN into the "discard" VLAN which will drop all untagged frames and tagged frames with unknown VLAN IDs. For Multi-layer and 3Com mode: if a tagged frame is received, but the VLAN for that tagged frame does not exists on the switch, that frame will be dropped. **Forwarding:** All broadcast frames from all VLANs learned on the port will be forwarded. **Egress:** All frames transmitted out of the port to be tagged using the IEEE 802.1Q/Multi-Layer tag header format. The tagged used will be that of the VLAN in which the frame was received.
		2 of 2

Assigning VLANs to a Port and Associated Issues

There are two ways to assign VLANs to ports on the Avaya Multiservice Switch. You configure the Port(s) to the VLAN desired for the individual port(s), or entire module. Assigning the VLAN this way enables the port(s) to receive information for the assigned VLAN and causes all untagged frames arriving on this port to be assigned to the specified VLAN. The port(s) will still assign incoming tagged packets to the VLAN indicated by the tag.

In the current release of software, a non-trunk port only supports a single VLAN per port. Multiple VLANs per port is NOT recommended and can have adverse effects on network performance.

The second way is to assign the VLAN to a port is by using the following CLI command in **Enable/Configure** mode:

`(configure)#` **set VLAN <*vlan-id or name*> <*options*>**

By selecting the VLAN name or VLAN ID, you can bind the selected VLAN to additional ports. When binding VLANs this way, the port is part of the flooding domain of the selected VLAN. This is an alternative to using the binding types “bind to all” and “bind to receive”, that allows you to add a port(s) to a subset of VLANs in the switch.

*** Note:** This configuration should only be used under special circumstances and with the assistance of customer support as undesirable results may occur (e.g. destination unknown unicast storms).

See Chapter 6, “Using VLANs, Hunt Groups, and VTP Snooping,” for a description of the parameters that will be displayed when viewing VLAN information with the Web Agent and the CLI command syntax.

Although this method of adding VLANs to a port could be thought to support the request for multiple VLANs per port (also referred to as overlapping VLANs), it is NOT recommended on clear (non trunked) ports due to the impact it has on the network. The impact is that destination unknown unicast packets are flooded on the VLAN in which the source host is located. This causes all ports assigned to this VLAN to receive the destination unknown unicast. If enough of these destination unknown unicast packets are being sent it could have a major impact on the network.

The following is an example of a problem with assigning more the one VLAN to a port.

- All PCs and Servers are connected to the Avaya Multiservice Switch 10/100TX switch ports. The following configuring of ports to VLANs are done with the Web Agent.
- PC1 is a member of VLAN1 and is connect to switch Port 1
- PC2 is a member of VLAN2 and is connected to switch Port 2
- SERVER1 is a member of VLAN4 and is connected to switch Port 4

With the **set vlan** CLI command, assign port 1, and port 2 to VLAN4 and also assign port 4 to VLAN1 and VLAN2. Then PC1 and PC2 could communicate with SERVER1 across VLANs without any noticeable problems. However, when PC1 transmits a packet to SERVER1, the following occurs:

- The packet is sent to the Switch
- The switch looks for the MAC address of SERVER1in the Address Forwarding Table for VLAN1 (because this is the actual port VLAN assigned to the port)
- The switch will be unable to find an entry for the MAC address of SERVER1
- The switch will flood this packet to all ports assigned to VLAN1 (regardless of how the VLAN was assigned to the port)

***Note:** The only exception to flooding a port is if 'Known Mode' is set to enable on the port, which stops the flooding of destination unknown unicast packets.

If you had many clients/servers on separate VLANs and they tried to communicate over multiple VLANs using this method, you would flood your network with undesired packets, thus slowing your overall network performance. Currently, the best solution is to use a L3 module to route between the VLANs. This would eliminate the broadcast of destination unknown unicast packets.

Configuring Physical Port Settings

This section contains the following procedures:

- Configuring Physical Ports on 10-Gigabit Ports
- Configuring Physical Ports on Gigabit Ports
- Configuring Physical Ports on Fast Ethernet Ports
- Using the All Ports Configuration Web Page

Configuring Physical Ports on 10-Gigabit Ports

Overview

You can configure the physical port parameters on 10-Gigabit ports by using either the Web Agent or the CLI.

For detailed information about each parameter, see Table 8-3.

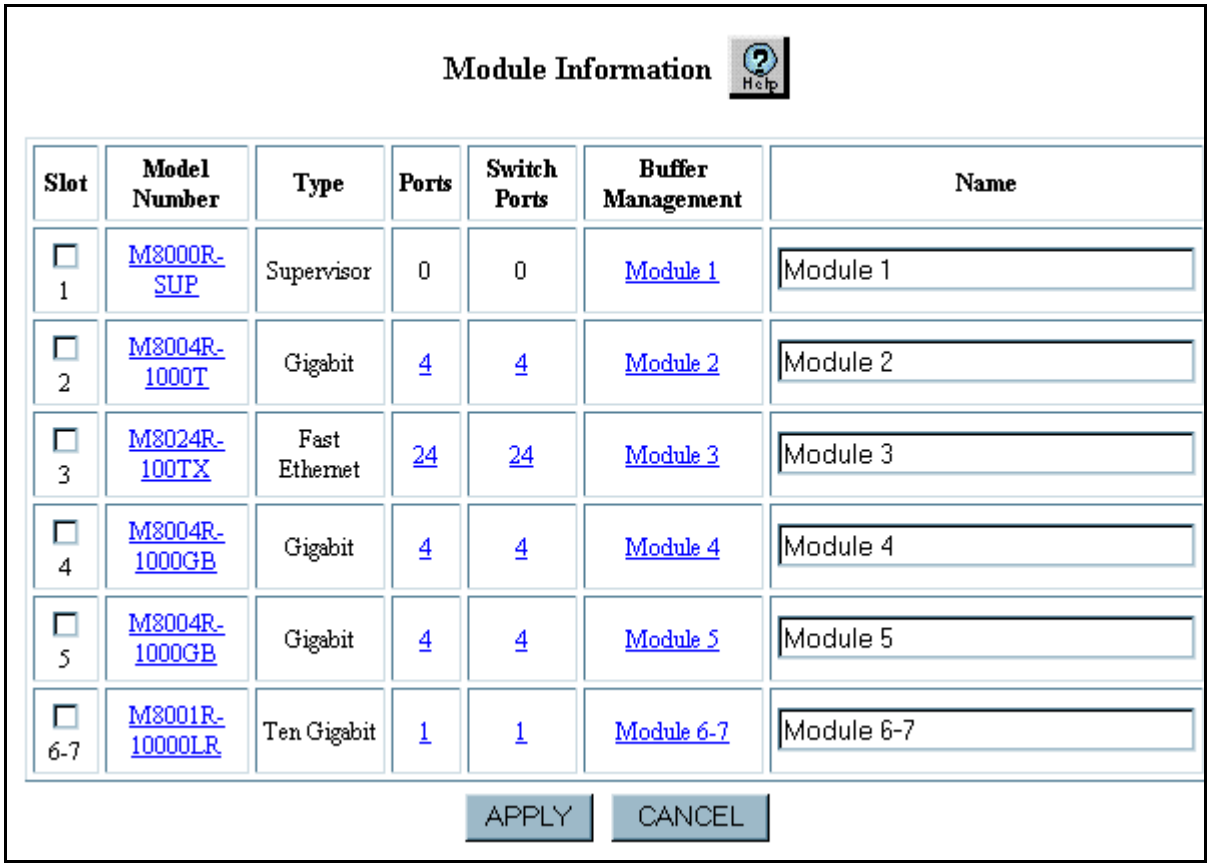
Web Agent Procedure

To configure a port on a 10-Gigabit module by using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder.
2. Select **Configuration**.

   The Module Information Web page is displayed in the content pane (Figure 8-2).

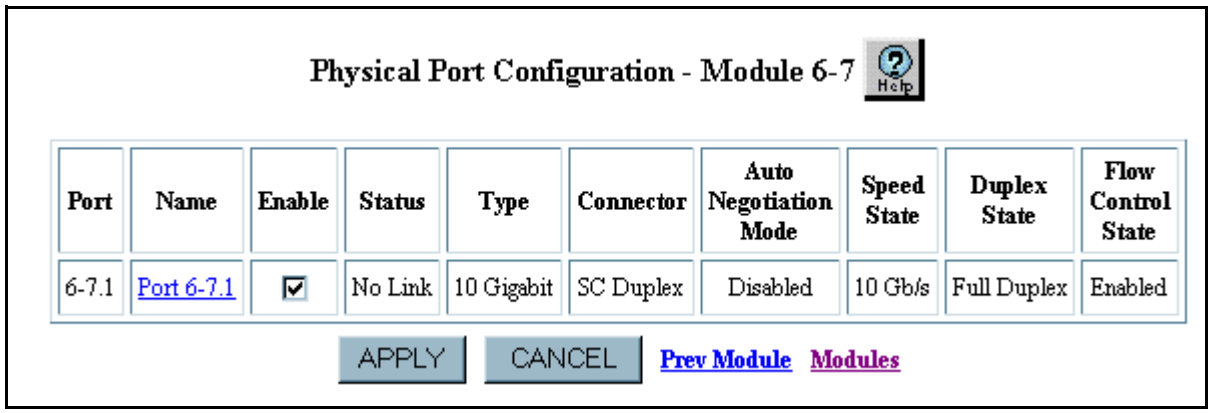
Figure 8-2. Module Information Web Page



3. In the **Ports** column for the 10-Gigabit module that you want to configure, select the port.

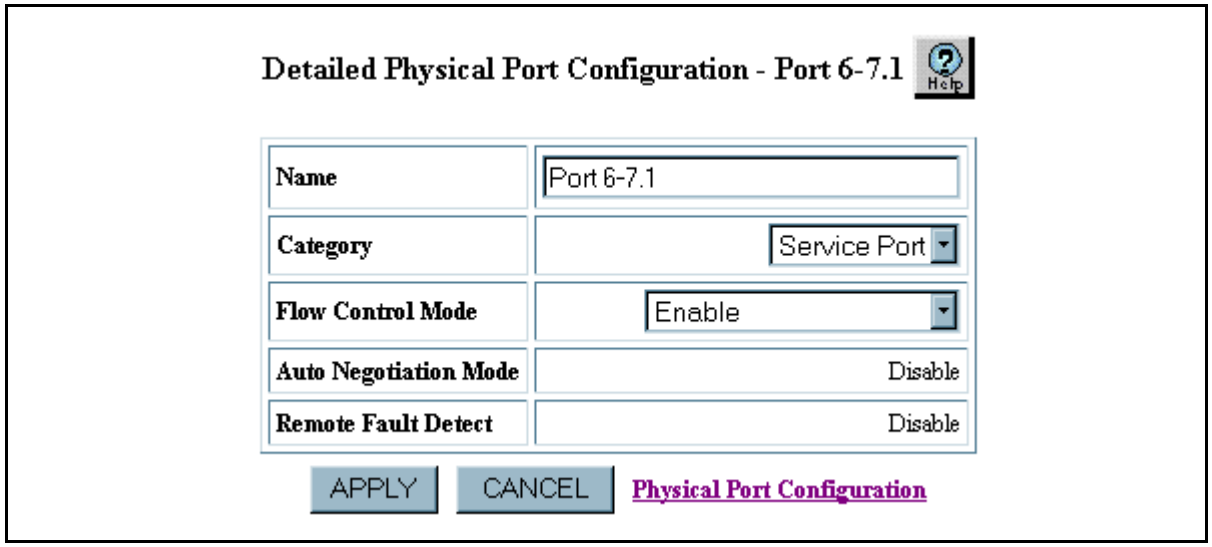
   The Physical Port Configuration Web page for that module is displayed. See Figure 8-3.

Figure 8-3. Physical Port Configuration Web Page for the 10-Gigabit Module



4. In the **Enable** column, select the checkbox to enable the port.
5. Click **APPLY** to save your settings, or **CANCEL** to restore previous settings.
6. In the **Name** column, select the port name.
7. The Detailed Physical Port Configuration Web page is displayed (Figure 8-4).

Figure 8-4. Detailed Physical Port Configuration Web Page for the 10-Gigabit Module



8. In the **Name** field, enter a port name, or you can use the default name.

9. In the **Category** field, select one of the following options:

   — **User Port** if you want the switch to generate only log messages for the port. This setting prevents the switch from generating alarm messages (SNMP traps) for the port.

   — **Service Port** (default setting) if you want the switch to generate both log messages and alarm messages (SNMP traps) for the port.

10. In the **Flow Control Mode** field, select one of the following options:

    Enable – Sets the port to both send and receive pause signals. This setting prevents buffer overflows on both local and remote port.

    Disable – Sets the port to neither send or receive pause signals. Use this setting when flow control is causing congestion in other areas of the network.

    Enable (Send Only) – Sets the port to only send pause signals. This setting prevents buffer overflows on the local port, but not on the remote port.

    Enable (Receive Only) – Sets the port to only receive and respond to pause signals. This setting prevents buffer overflows on the remote port, but not on the local port.

11. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 8-3. Detailed Physical Port Web Page

Parameter	Definition
Name	A name for this port. For example, you could enter a drop connection name or the name of the station or other device that is connected to the port.
Category	• The **User Port** option is intended for connections to end-user nodes. If you select this option, the switch generates only *log messages* for the port. It does not generate *alarm messages* (SNMP traps). This setting prevents the network management station (trap receiver) from being overwhelmed by port status messages that result from users turning workstations on and off. If you want to prevent the switch from generating SNMP trap messages for the 10-Gigabit port, select this option. • The **Service Port** option is intended for connections to servers or other switches. Since the 10-Gigabit module is connected only to other switches, this option is the default setting. If you select this option, the switch generates both *log messages* and *alarm messages* (SNMP traps) for the 10-Gigabit port.
Flow Control Mode	This field determines if the port uses IEEE 802.3z pause control. The pause mechanism allows the port to stop a sending station from sending more packets if the buffers of the receiving port are full. Flow control helps prevent lost or dropped packets. This feature is recommended for use primarily on end station connections. Using this feature on trunk ports can cause unnecessary congestion on the network. Select one of the following options: **Enable** – Sets the port to both send and receive pause signals. This setting prevents buffer overflows on both local and remote port. **Disable** – Sets the port to neither send or receive pause signals. Use this setting when flow control is causing congestion in other areas of the network. **Enable (Send Only)** – Sets the port to only send pause signals. This setting prevents buffer overflows on the local port, but not on the remote port. **Enable (Receive Only)** – Sets the port to only receive and respond to pause signals from the remote port. This setting prevents buffer overflows on the remote port, but not on the local port.
	1 of 2

Table 8-3. Detailed Physical Port Web Page

Parameter	Definition
Auto Negotiation Mode	Auto-Negotiation is not supported on the 10-Gigabit module. The IEEE 802.3ae standard refers to but does not specify an autonegotiation function. 10-Gigabit devices should all operate at 10Gigabits/sec in Full Duplex mode.
Remote Fault Detect	Remote fault detection is not supported on the 10-Gigabit module. The 802.3ae standard provides for automatic detection of link failure. This setting is set to **Disable** and cannot be changed.
	2 of 2

CLI Command

To configure ports on a 10-Gigabit module, use the following CLI command:

```
(configure)# set port
```

When you use the CLI to configure the 10-Gigabit module, enter the the module number in the format of **<*Slot N*>-<*Slot N+1*>**. For example, to enable the 10-Gigabit port if the module were in slots 5 and 6, you would enter **set port enable 5-6/1**.

You must be in Global Configuration mode to enter the **set port** command.

Configuring Physical Ports on Gigabit Ports

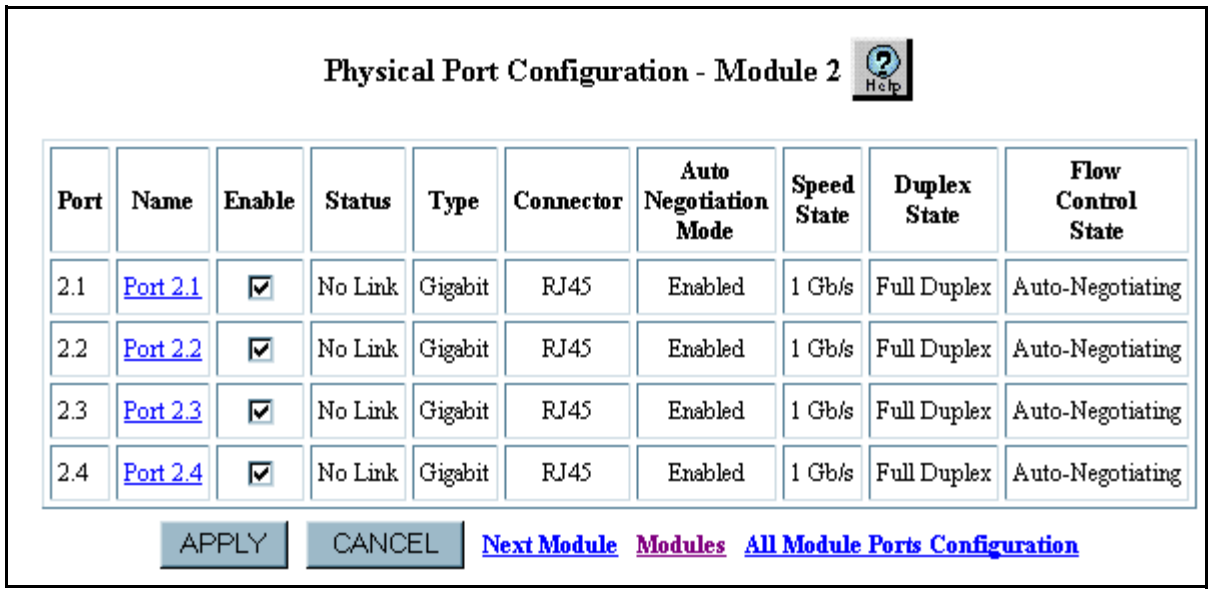
You can configure the physical port parameters on Gigabit ports using either the Web Agent or the CLI.

Web Agent Procedure

To configure ports on a gigabit module using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Configuration**. The Module Information Web page is displayed in the content pane. See Figure 8-2.
2. Select the port from the **Ports** column for the Gigabit module that you want to configure. The Physical Port Configuration Web page for that module is displayed. See Figure 8-5.

Figure 8-5. Physical Port Configuration Web Page for the Gigabit Ethernet Module



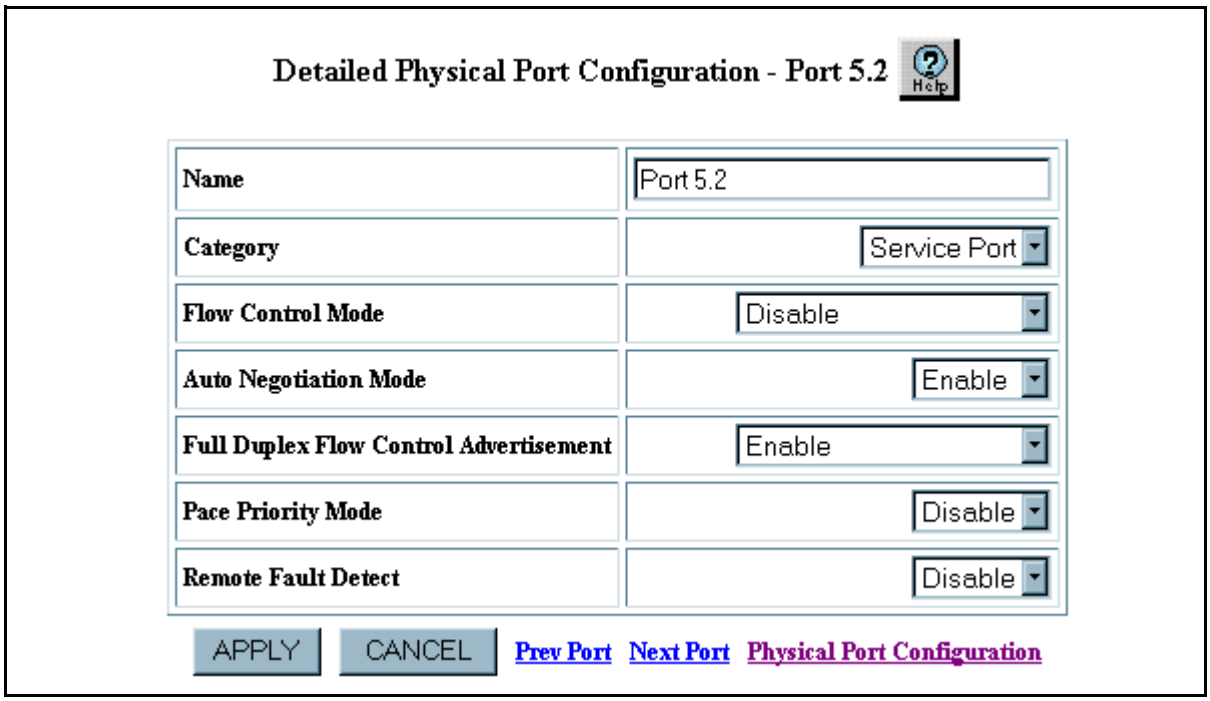
3. Click the checkbox in the **Enable** column to enable a port. A check mark displays.

*** Note:** If the port is already enabled (check mark is visible), and you want to disable the port, click the box to Disable the port (check mark vanishes).

4. Click **APPLY** to save your settings, or **CANCEL** to restore previous settings.

5. Select the port name from the **Name** field. The Detailed Physical Port Configuration Web page is displayed. See Figure 8-6.

Figure 8-6. Detailed Physical Port Configuration Web Page for Port on Gigabit Ethernet Module



6. Enter a port name in the **Name** field, if desired, or use the default name supplied

7. Select one of the following from the **Category** field pull-down menu:

   — **User Port i**f this is an end-station port

   — **Service Port** if this is a trunk port.

8. In the **Flow Control Mode** field, select one of the following options:

   Enable – Sets the port to both send and receive pause signals. This setting prevents buffer overflows on both local and remote port.

   Disable – Sets the port to neither send or receive pause signals. Use this setting when flow control is causing congestion in other areas of the network.

   Enable (Send Only) – Sets the port to only send pause signals. This setting prevents buffer overflows on the local port, but not on the remote port.

   Enable (Receive Only) – Sets the port to only receive and respond to pause signals. This setting prevents buffer overflows on the remote port, but not on the local port.

9. Select **Enable** from the **Auto Negotiation Mode** field pull down menu

10. Select **Enable** from the **Full Duplex Flow Control Advertisement** field pull down menu

11. Select **Enable** from the **Pace Priority Mode** field pull down menu to recognize and use 3Com's PACE priority mechanism.

12. Select **Enable** from the **Remote Fault detect** field pull down menu to detect

13. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 8-4 provides a more detailed description of the Detailed Physical Port Web page parameters.

Table 8-4. Detailed Physical Port Web Page Parameters

Parameter	Definition
Name	A user-definable name for this port (possibly a drop connection name or the name of the station or other device connected to the port).
Category	The **User Port** option is intended for use with switch connections to end user nodes. It is intended for use with switch connections to servers or other switches. The **Service Port** allows the switch to generate both *log messages* and *alarm messages* (traps). The User Port only generates log messages. This prevents your network management station from being overwhelmed by port up/down messages that result from users turning workstations on and off.
	1 of 2

Table 8-4. Detailed Physical Port Web Page Parameters

Parameter	Definition
Flow Control Mode	This field determines if the port uses IEEE 802.3z pause control. The pause mechanism allows the port to stop a sending station from sending more packets if the buffers of the receiving port are full. Flow control helps prevent lost or dropped packets. This feature is recommended for use primarily on end station connections. Using this feature on trunk ports can cause unnecessary congestion on the network. Select one of the following options: **Enable** – Sets the port to both send and receive pause signals. This setting prevents buffer overflows on both local and remote port. **Disable** – Sets the port to neither send or receive pause signals. Use this setting when flow control is causing congestion in other areas of the network. **Enable (Send Only)** – Sets the port to only send pause signals. This setting prevents buffer overflows on the local port, but not on the remote port. **Enable (Receive Only)** – Sets the port to only receive and respond to pause signals from the remote port. This setting prevents buffer overflows on the remote port, but not on the local port
PACE Priority Mode	Determines if the port detects 3Com's copyrighted PACE format as packets pass through the port. PACE allows a packet's priority (higher priority packets move through the switch faster) to be set at the adapter.
Remote Fault Detect	Remote fault detection makes it possible for a Gigabit port at one end of a link to signal status to the other end of the link, even if it does not have an operational receive link. For more information on remote fault detection, see "Physical Port Features" earlier in this chapter. **Note:** Auto-negotiation and remote fault detection can not be enabled at the same time.Auto-negotiation must be disabled to enable remote fault detection. When auto-negotiation is enabled, remote fault detection is automatically disabled.
	2 of 2

CLI Command

To configure ports on a gigabit module using the CLI, enter the following command from Configure mode:

`(configure)#` **set port** <***options***>

*** Note:** When a port on an 80-Series Gigabit Ethernet module is disabled, the port link light blinks continuously. In versions earlier than v5.2.10, the link light extinguished when the port was disabled.

Loopback tests on ports may fail when traffic is present on the link at startup.

Configuring Physical Ports on Fast Ethernet Ports

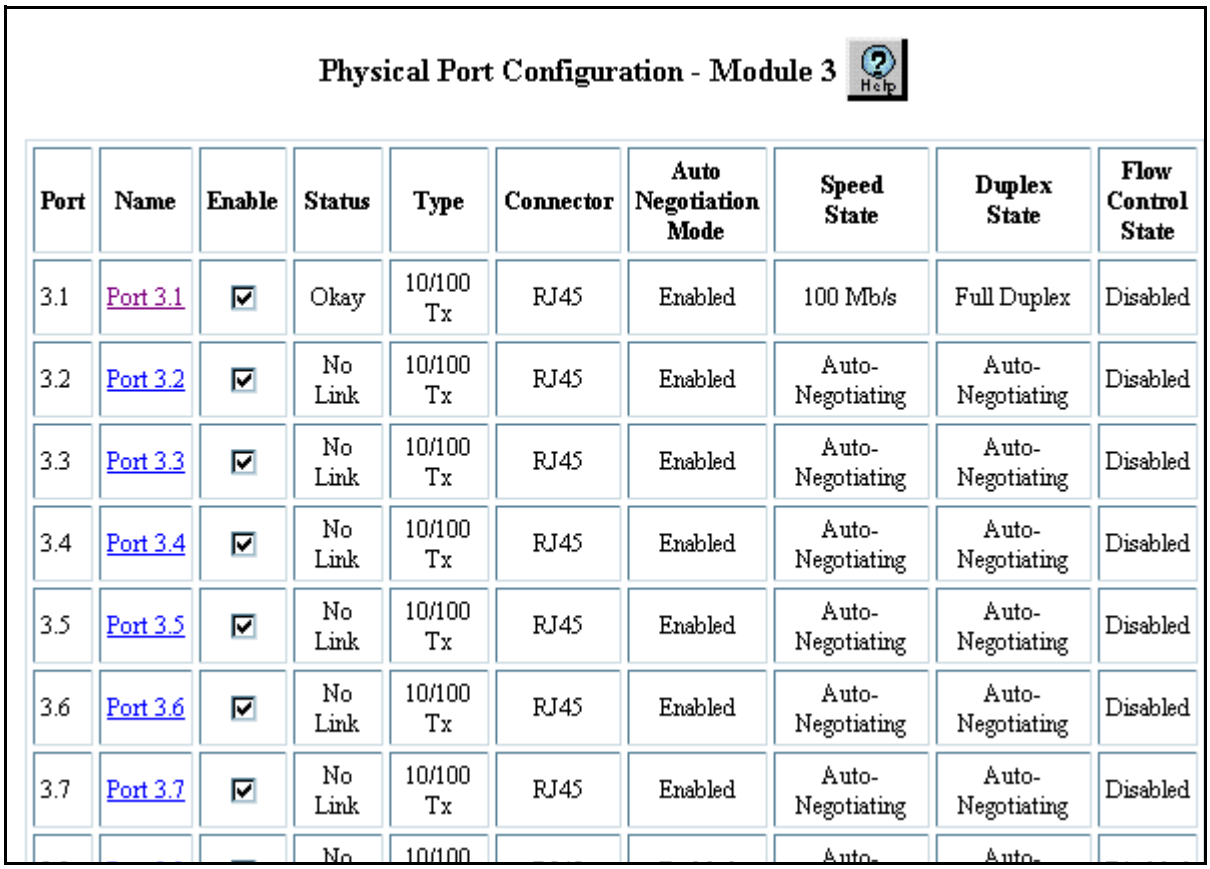
You can configure ports on a Fast Ethernet module using either the Web Agent or CLI commands.

Web Agent Procedure

To configure ports on a Fast Ethernet module using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Configuration**. The Module Information Web page in the content pane. See Figure 8-2.
2. Select the checkbox in the **Slot** column next to the Fast Ethernet module to change the Name field, if desired, change the module name then Click the **APPLY** button.
3. Select the number in the **Ports** column for the Fast Ethernet module that you want to configure. The Physical Port Configuration Web page is displayed. See Figure 8-7.

Figure 8-7. Physical Port Configuration Web Page for the Fast Ethernet Module



Physical Port Configuration - Module 3

Port	Name	Enable	Status	Type	Connector	Auto Negotiation Mode	Speed State	Duplex State	Flow Control State
3.1	Port 3.1	☑	Okay	10/100 Tx	RJ45	Enabled	100 Mb/s	Full Duplex	Disabled
3.2	Port 3.2	☑	No Link	10/100 Tx	RJ45	Enabled	Auto-Negotiating	Auto-Negotiating	Disabled
3.3	Port 3.3	☑	No Link	10/100 Tx	RJ45	Enabled	Auto-Negotiating	Auto-Negotiating	Disabled
3.4	Port 3.4	☑	No Link	10/100 Tx	RJ45	Enabled	Auto-Negotiating	Auto-Negotiating	Disabled
3.5	Port 3.5	☑	No Link	10/100 Tx	RJ45	Enabled	Auto-Negotiating	Auto-Negotiating	Disabled
3.6	Port 3.6	☑	No Link	10/100 Tx	RJ45	Enabled	Auto-Negotiating	Auto-Negotiating	Disabled
3.7	Port 3.7	☑	No Link	10/100 Tx	RJ45	Enabled	Auto-Negotiating	Auto-Negotiating	Disabled

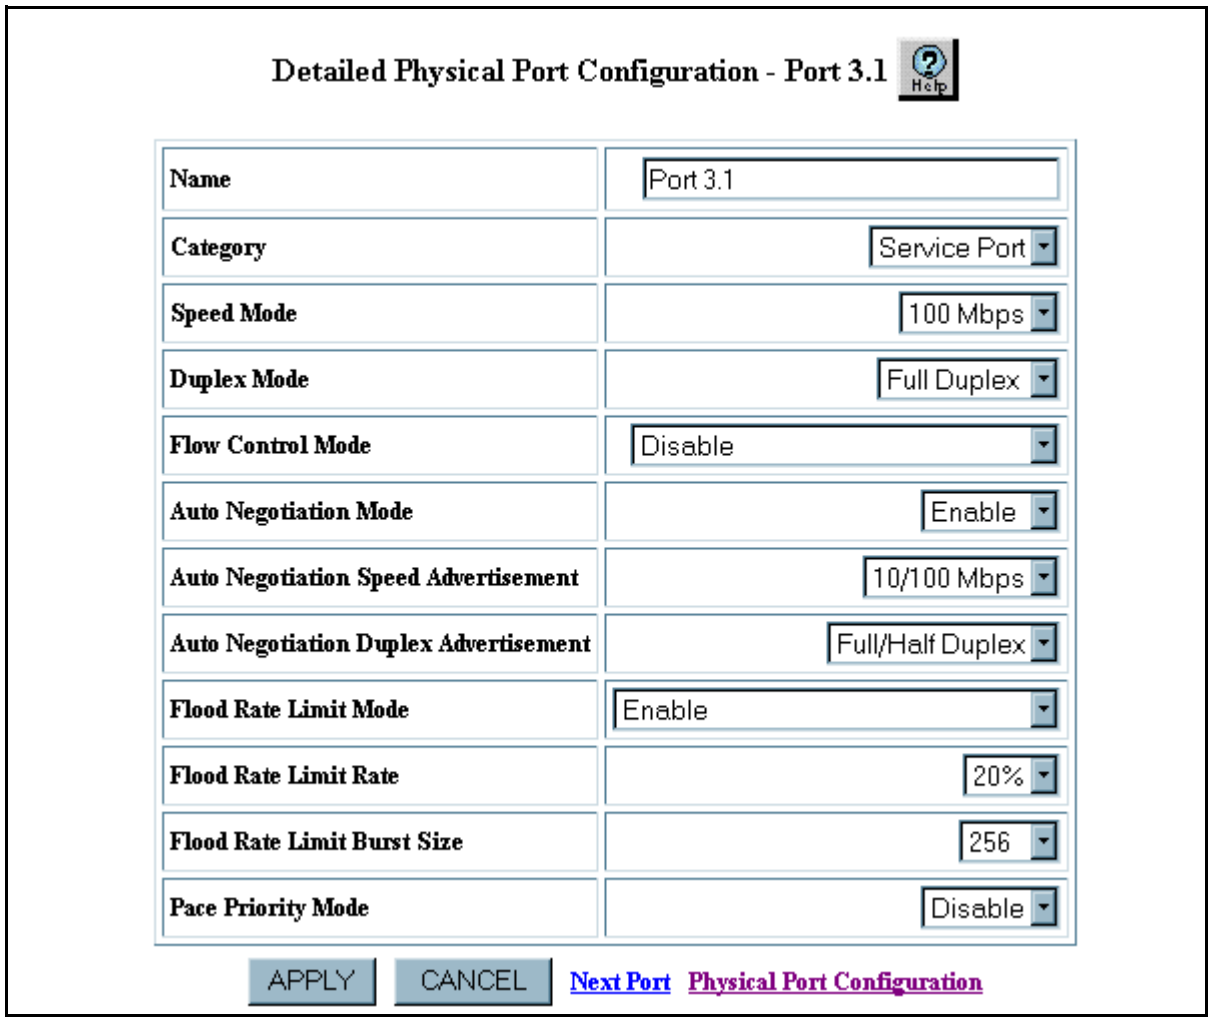
4. By default the checkbox is checked in the **Enable** column. To change the enable remove the check in the checkbox.

5. If you change the enable checkbox **Click APPLY** to save your settings, or **CANCEL** to restore previous settings.

6. Select a port name from the **Name** field (for example, Port 6.1).

   The Detailed Physical Port Configuration Web page for that port is displayed. See Figure 8-8.

Figure 8-8. Detailed Physical Port Configuration Web Page for Port on Fast Ethernet Module



7. Enter a port name in the **Name** field, if desired or use the default name.
8. Select one of the following from the **Category** field pull-down menu:
   - **User Port** if this is an end-station port
   - **Service Port** if this is a trunk port.
9. Select a speed (**10 Mb/s** or **100 Mb/s**) from the **Speed Mode** field pull down menu, If you want to set the port speed manually. If want the port to autonegotiate, you can ignore this field.

***Note:** This feature is only available for 10/100 TX modules.

10. Select a duplex mode (half-duplex or full-duplex) from the **Duplex Mode** field pull-down menu if you want to manually set the port's duplex mode. You do not have to select a duplex mode if you set the port to autonegotiate.

11. In the **Flow Control Mode** field, select one of the following options:

    Enable – Sets the port to both send and receive pause signals. This setting prevents buffer overflows on both local and remote port.

    Disable – Sets the port to neither send or receive pause signals. Use this setting when flow control is causing congestion in other areas of the network.

    Enable (with Aggressive Backoff) – Limits the size of flow control burst on TX and FX ports.

12. Select **Enable** from the **Auto Negotiation Mode** field pull-down menu to enable Auto Negotiation.

    ***Note:** This feature is only available for 10/100 TX modules.

    Auto Negotiation works best when the port or device on the other end of the connection is also set to Auto Negotiation. If you are having problems with Auto Negotiating connections, manually set the modes using the CLI.

13. Select a speed from the **Auto Negotiation Speed Advertisement** pull-down menu. The options are 10Mbps, 100Mbps. or 10/100Mbps.

14. Select a duplex option from the **Autonegotiation Duplex Advertisement** field pull-down menu. The options are Full/Half Duplex and Half Duplex).

    ***Note:** The Speed and Duplex features are only available for 10/100 TX modules.

    The switch sends the Speed and Duplex advertisement values to the device on the other end of the connection at the start of the auto-negotiating process. In general, the defaults are best, but there may be situations when you want to fix one setting, but allow the other setting to auto-negotiate.

15. Select one of the following from the **Flood Rate Limit Mode** pull-down menu:

    Enable – If you want this port to limit the number of unknown unicast and multicast (flooded) packets it tries to forward

    Disable – If you do not want this port to limit the number of unknown unicast and multicast packets it tried to forward

    Enable (include Known Multicasts) – If you want to optionally include known multicast packets in this percentage to further decrease the possibility of the port's output buffer being overwhelmed.

16. Select the percentage of a port's traffic that can be unknown unicast and broadcast packets from the **Flood Rate Limit Rate** field pull-down menu.

    ***Note:** Set this value lower if the port is having overflow problems.

17. Open the **Flood Rate Limit Burst Size** pull down menu.

18. Select a packet limit for the number of packets allowed in a single burst. The values are 1 to 2048.

    ***Note:** Set this value lower than 1024 (the output buffer's capacity) for Fast Ethernet ports. Also set this value lower if the port is experiencing overflow problems.

19. Select **Enable** from the **Pace Priority Mode** pull-down menu if you want this port to recognize and use 3Com's PACE priority mechanism.

20. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 8-5 describes in detail the Fast Ethernet port parameters:

Table 8-5. Fast Ethernet (10/100) Port Parameters

Parameter	Definition
Name	Enter a user configurable name for this port (possibly a drop connection name or the name of the station or other device connected to the port).
Category	Select either User Port or Service Port. • The **User Port** is intended for use with switch connections to end user nodes. • The **Service Port** is intended for use with switch connections to servers or other switches. The Service Port allows the switch to generate both *log messages* and *alarm messages* (traps). The User Port only generates log messages, preventing your network management station from being overwhelmed by port up/down messages that result from users turning workstations on and off.
Speed Mode	Select the speed of the port manually (to either 10 or 100 Mb/s). **If auto-negotiation is enabled, this setting is ignored.**
Duplex Mode	Select the port duplex mode (half- or full-duplex). **If auto-negotiation is enabled, this setting is ignored.**
Flow Control Mode	Determines if flow control is used on this port. For half-duplex links, active backpressure jams the sending Ethernet channel until the port's buffers can receive more packets. This prevents lost or dropped packets. For full-duplex links, IEEE 802.3z pause control allows the port to stop a sending station from sending more packets if the receiving port's buffers are full. For TX and FX ports, there is an additional option for Enable with Aggressive Backoff. This option limits the size of the bursts. Flow Control is recommended for use primarily on end-station connections. Using this flow control on trunk ports can cause unnecessary congestion on the network.
Auto Negotiation Mode	Select the port to auto-negotiate a speed and duplex mode. Auto-negotiate works best when the connection on the other end of the link is set to auto-negotiate as well. If you set a port to auto-negotiate, and the connection is not successful, set the port speed and duplex mode manually.
Auto Negotiation Speed Advertisement	Determines what information the port advertises when it starts auto-negotiating. In most cases, **10/100** are the best settings, but there may be cases when you want to auto-negotiate one parameter, while keeping the other fixed.
	1 of 2

Table 8-5. Fast Ethernet (10/100) Port Parameters

Parameter	Definition
Auto Negotiation Duplex Advertisement	Determines what information the port advertises when it starts auto-negotiating. In most cases, **Half/Full** are the best settings, but there may be cases when you want to auto-negotiate one parameter, while keeping the other fixed.
Flood Rate Limit Mode	Prevents the switch from overwhelming the output buffer on lower-speed ports by placing a threshold on the percentage of port traffic that can be flooded packets (unknown unicasts and multicasts). You can also optionally include known multicast packets in this percentage to further decrease the possibility of the port's output buffer that is being overwhelmed.
Flood Rate Limit Rate	Determines the percentage of a port's forwarded traffic that can be unknown unicast and multicast (flooded). Lower this value if the port has overflow problems.
Flood Rate Limit Burst Size	Enter a value for the limit of packets allowed in a single burst. Accepted values are 1 to 2048. For Fast Ethernet ports, set this value lower than 1024 (output buffer capacity). Lower this value if the port has overflow problems.
Port PACE Priority	Select Enable to allow the port to detect 3Com's proprietary PACE format as packets pass through the port. PACE allows a packet's priority (higher priority packets move through the switch before lower priority packets) to be set at the adapter.
	2 of 2

CLI Command

To configure ports on a Fast Ethernet module using the CLI, enter the following command from Configure mode:

```
(configure)# set port <options>
```

Using the All Ports Configuration Web Page

You can apply the same parameter settings on all of a module's ports using the All Ports Configuration Web page.

To configure all ports on a module:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Configuration**. The Module Information Web page is displayed in the content pane. See Figure 8-2.
2. Select the port number from the **Ports** column for that module. The Physical Port Configuration Web page for that module is displayed. See Figure 8-5 for Gigabit modules and Figure 8-7 for 10/100 modules.
3. **Select All Module Ports Configuration**. The All Ports Configuration Web page is displayed.
4. See "Configuring Physical Ports on 10-Gigabit Ports," "Configuring Physical Ports on Gigabit Ports," and "Configuring Physical Ports on Fast Ethernet Ports" earlier in this chapter to configure the ports.
5. Click **APPLY** to save your changes, or **Restore** to restore previous settings.

Configuring Switch Ports Settings

This section contains the following information and procedures:

- Configuring Switch Ports
- Automatic VLAN Creation
- Configuring MAC Address Lock and Intrusion Detection
- Using the All Module Ports Configuration Web Page
- Viewing Switch Port Settings

Configuring Switch Ports

You can configure how each switch port performs its switching functions (for example, VLAN parameters, hunt group assignments, trunk mode, and frame tag scheme) through the switch port parameters.

You can configure switch port parameters using either the Web Agent or the CLI.

CAUTION:

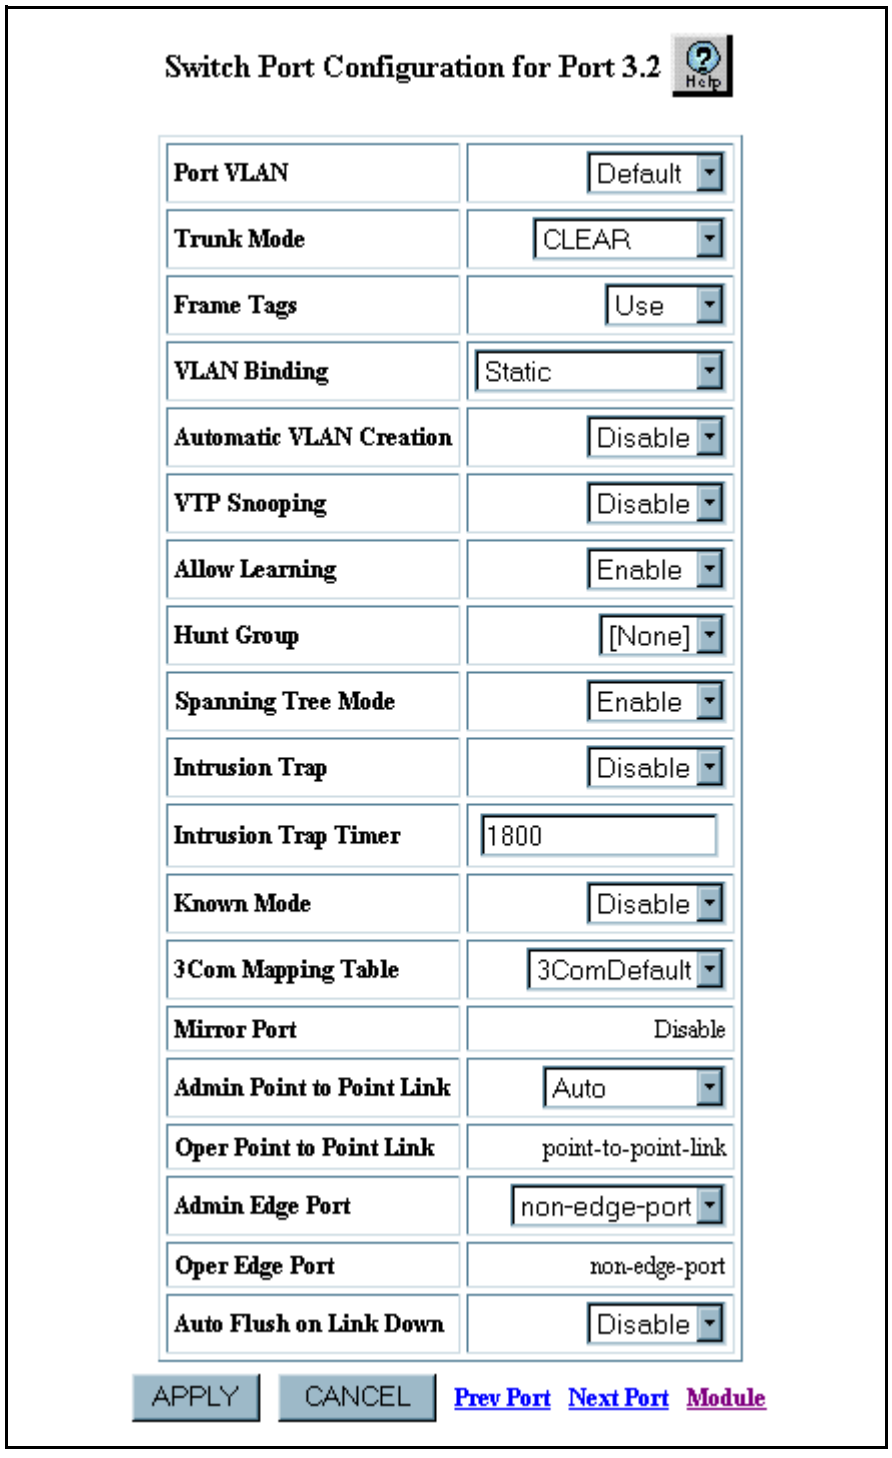
Before changing the VLAN that a port is assigned to, you must remove the port from the hunt group that the port is assigned to. Do *not* attempt to change the VLAN and remove the port from the hunt group simultaneously. If you attempt to change the VLAN that the port is assigned to and remove the port from the hunt group simultaneously, the switch assigns all ports in the hunt group to the new VLAN.

Web Agent Procedure

To configure switch port parameters using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Configuration.** The Module Information Web page is displayed in the content pane. See Figure 8-2.
2. In the **Switch Ports** column, select the switch ports for a module. The Switch Ports Web page is displayed.
3. In the **Name** column, select a port name. The Switch Port Configuration Web page is displayed. See Figure 8-9.

Figure 8-9. Switch Port Configuration Web Page



4. Enter information in the Web page fields as appropriate. See Table 8-6 for an explanation of each field.

5. Click **APPLY** to save your changes, or **CANCEL** to restore default settings.

*** Note:** Do not enable Automatic VLAN Creation and VTP Snooping at the same time.

Table 8-6 describes the Switch Port Configuration Web page fields.

Table 8-6. Switch Port Configuration Web Page Fields

Parameter	Definition
Port VLAN	Specifies the VLAN assignment for this port. All *untagged* frames arriving on this port are assigned to this VLAN. The port still assigns incoming *tagged* packets to the VLAN indicated by the tag.
Trunk Mode	Select the appropriate VLAN trunking format to make the port a trunk, or Clear (default) if you do not want the port to be a trunk. The trunk formats are IEEE 802.1Q, Multi-Layer, and 3Com. **Note:** The 10-Gigabit module supports a **Trunk Mode** setting of **IEEE802.1 Q** or **Clear**.
Frame Tags	Select Ignore if you do not want to use received Frame VLAN tags. Use is the default. **Note:** If you select ignore, the received frames are bound to the port's default VLAN.
VLAN Binding	Select the port's outgoing VLAN binding type. The options are Static (default), Bind-to-All, and Bind-to-Receive. See Table 8-8 for an explanation of these options.
Automatic VLAN Creation	Select Enable to automatically create a VLAN each time the port receives a frame from an unknown VLAN. Disable is the default.
VTP Snooping	Select Enable to allow VTP Snooping on this port. Disable is the default. VTP is a Cisco proprietary Layer 2 trunk port protocol that is used among Cisco switches, over trunk ports, to maintain a VLAN configuration consistency across the network.The switch performs VTP operations when connected to a Cisco switch's trunk that has VTP enabled. The switch queries its Cisco peer for its VLAN configuration. It also learns Cisco VLAN configurations from received Cisco message.
	1 of 4

Table 8-6. Switch Port Configuration Web Page Fields

Parameter	Definition
Allow Learning	Select Disable to prevent the port from learning new MAC addresses and forwarding them to the Supervisor module to be added to the VLANs AFT. Enable is the default. For example, you can set this parameter to Disable, then add a static MAC address entry for this port. If you do not select Disable before you add a static MAC address, unicast flooding can occur. **Note:** If a MAC address is moved to a different VLAN, it will coexist in AFTs, one entry for each VLAN (This is a benefit of having multiple forwarding databases).
Hunt Group	Select a hunt group for which this port will be a member. None is the default. For the 10-Gigabit module, a need for redundancy would probably be the reason that you use a hunt group. **Note:** All ports that participate in the same hunt group must have the same bandwidth. Thus, 10-Gigabit ports can participate only in a hunt group that consists of other 10-Gigabit ports.
Spanning Tree Mode	Select Disable to remove STP on this port. Enable is the default.
Intrusion Trap	Enable or disable intrusion traps. If intrusion traps are enabled, trap messages are sent to the event log when the port receives an unknown source MAC address. One trap is generated per intrusion trap timer setting. For this feature to provide maximum port security, the port must have a static MAC address and learning must be disabled. For more information on intrusion traps, see "Configuring MAC Address Lock and Intrusion Detection" later in this chapter.
Intrusion Trap Timer	Time interval at which intrusion traps are generated. The default setting for the intrusion trap timer is 1800 seconds (30 minutes). The valid range for the timer is 60 to 1800 seconds.
	2 of 4

Table 8-6. Switch Port Configuration Web Page Fields

Parameter	Definition
Known Mode	Enable or disable known mode. The default value is Disable. If known mode is *enabled*, unicast frames that have an unknown destination address are *not* flooded out this port. If known mode is *disabled*, unicast frames that have an unknown destination address *are* flooded out the port. For example, if a known end-station or file server is connected to the port, flooding unknown unicasts out the port is not necessary.
3Com Mapping Table	Select how incoming tagged frames from 3Com equipment are mapped to Avaya VLANs. The default value is 3ComDefault.
Mirror Port	Displays the status of the mirror port (enabled or disabled). This is a Fast Ethernet only option.
Admin Point to Point Link	Specify whether this port is connected to a shared LAN segment or a point-to-point LAN segment. A point-to-point LAN segment is connected to exactly one other bridge (normally with a direct cable between them). Only point-to-point links and edge ports can rapidly transition to forwarding state. If you set this field to **Auto**, the switch automatically detects whether the port is connected to a shared link or a point-to-point link. Ports operating in half duplex are set to **False**, and ports operating in full duplex are set to **True**. You can, however, manually set the type of link. Options are: • ForceTrue—Defines the port as connected to a point-to-point link. • ForceFalse—Defines the port as connected to a shared LAN segment. • Auto—Automatically detects whether the port is connected to a shared link or a point-to-point link. If you select this setting, the **Oper Point to Point Link** field displays the link type that is detected. For more information on Rapid Spanning Tree, see Chapter 7, "Configuring Rapid Spanning Tree."
Oper Point to Point Link	If **Admin Point to Point Link** is set to **Auto**, this field displays the link type that is detected.
	3 of 4

Table 8-6. Switch Port Configuration Web Page Fields

Parameter	Definition
Admin Edge Port	Specify whether this port is an edge port or a nonedge port. An edge port is not connected to any other bridge. Only edge ports and point-to-point links can rapidly transition to forwarding state. Options are: • **edge-port**—Defines the port as an edge port. • **non-edge-port**—Defines the port as a nonedge port. If you set this field to **edge-port**, the **Oper Edge Port** field is also set to **edge-port**. However, if the port receives a BPDU, the **Oper Edge Port** setting changes to **non-edge-port**. (To receive a BPDU, the port must be connected to a bridge and thus is not an edge port.) This field replaces the fast start feature in v5.*x* versions of application software. When you upgrade the software from v5.*x* to v6.0 or later, switch ports for which fast start was enabled have **Admin Edge Port** set to **edge-port**. For more information on Rapid Spanning Tree, see Chapter 7, "Configuring Rapid Spanning Tree."
Oper Edge Port	If **Admin Edge Port** is set to **edge-port**, this field is also set to **edge-port**. However, if the port receives a BPDU, the **Oper Edge Port** setting changes to **non-edge-port**.
Auto Flush on Link down	Select enable Auto Flush for a port to mark all entries learned on that port invalid if the link to that port fails. You should enable Auto Flush on all P580 and P882 ports that are participating in a Port Redundancy configuration with a P330.
	4 of 4

Example

See Table 8-7 as an example of the recommended switch port settings for a Cisco Catalyst 5000™.

Table 8-7. Example 1. Trunk to Cisco Catalyst 5000™

Parameter	Recommended Setting
Port VLAN	**Default** - Causes untagged packets to be assigned to the default VLAN.
Trunk Mode	**Multi-layer** - Causes the port to send frames using the multi-layer format.
VLAN Binding	**Bind to All** - Binds the port to all VLANs known to the switch.
	1 of 2

Table 8-7. Example 1. Trunk to Cisco Catalyst 5000™

Parameter	Recommended Setting
Automatic VLAN Creation	**Enable (Disable if using VTP Snooping)** - Causes the switch to learn new VLAN IDs that arrive at the port, and then bind the port to these VLANs.
VTP Snooping	**Enable** - Causes the switch to update its VLANs as they are created, deleted, or changed on the Catalyst.
	2 of 2

CLI Command

To configure switch ports using the CLI, use the following CLI command:

`(configure)#` **set port <*options*>**

*** Note:** See the examples later in this section for recommendations on how to set particular trunk port connections.

*** Note:** See "VLAN Operation" in Chapter 6 for more information on creating VLANs.

Automatic VLAN Creation

Automatic VLAN creation is done by enabling the Automatic VLAN Creation parameter on an individual port under the 'Module & Ports > Configuration > Switch Ports > name menu. With this parameter enabled, the port will automatically create a VLAN each time it receives a frame from an unknown VLAN ID in received IEEE 802.1Q and Multi-Layer tagged frames.

*** Note:** The automatic VLAN creation feature does not create entries in 3Com Mapping Tables.

When a VLAN is created automatically the VLAN name and VLAN ID are derived from the received tagged frame. The VLAN name will be created as **autoVlan <VLAN ID>*. The VLAN ID will be identical to the VLAN ID of the received tagged frame. As with all VLANs, the VLAN name may be modified to something more descriptive.

Table 8-9 shows the Avaya P580 and P882 Multiservice switch VLAN table. The VLANs *Default* and *Discard* are permanent VLANs assigned to every switch. VLAN *Net90* is a manually created VLAN, while VLAN **autoVlan1001* was created automatically by software.

Table 8-8 describes the **VLAN Binding** field options.

Table 8-8. VLAN Binding Options

Option	Definition
Static	Assigns VLAN membership manually, using the VLAN Switch Ports page described in "Creating and Implementing VLANs" in Chapter 6, "Using VLANs, Hunt Groups, and VTP Snooping."
Bind to All	Binds this port to all VLANs known to the switch. This is an appropriate mode for switch-to-switch connections. **Note:** When a tagged IEEE 802.1Q packet arrives on a port that is set to bind to all and the VLAN does not exist on the switch, the packet is forwarded on to the VLAN assigned to the port default VLAN for that port. To prevent unintended forwarding of unknown VLAN traffic to the port's default VLAN, configure the port default VLAN to Discard. The automatic VLAN creation feature will not work if the port's default VLAN is the discard VLAN, because the switch does not learn for this VLAN.
Bind to Received	Binds this port to any VLAN it receives traffic from. **Note:** If Automatic VLAN Creation is enabled, the port binds to previously unknown VLANs and a VLAN entry is added to the switch VLAN table. If Automatic VLAN Creation is disabled, the port does not bind to any VLAN unknown to the switch.

Table 8-9. Avaya P580 and P882 Multiservice Switch VLAN Table

Name	ID	Group ID	AFT Index
Default	1	2	1
Discard	4097	3	3
Net90	90	90	7
*autoVlan100	100	100	11
* Automatically created by the switch software. All others created manually.			

Automatically Creating VLANs and Frame Tags Parameters

VLANs are created automatically by the switch reading the VLAN tag of all ingress IEEE 802.1Q and Multi-Layer tagged frames. The switch then creates a new VLAN for every new VLAN tag identified. However, if the parameter for **Frame Tags** is set to **Ignore**, the switch will ignore the VLAN tags on ingress frames. The switch assumes that all ingress frames belong to the 'Port VLAN'. Therefore, no new VLANs will ever be created automatically.

Automatic Binding of VLANs to Ports

When a VLAN is either manually or automatically created, the software may automatically bind the VLAN to a port depending on the setting of the 'Trunk Mode' parameter assigned to the port.

- When a VLAN is created manually the software assigns the VLAN to all ports whose **VLAN Binding** is set to **Bind to All**. No other ports will automatically be assigned when a VLAN is created manually.
- When a VLAN is created automatically the software assigns the VLAN to the port it is received on if that port is set to **Bind to All** and **Bind to Received**. Additionally, software will assign the VLAN to all other ports whose **VLAN Binding** is set to **Bind to All.**

following command from **Configure** mode:

`(configure)#` **set port VLAN** <***options***>

Configuring MAC Address Lock and Intrusion Detection

Overview

The MAC Address Lock feature filters frames that do not match the static MAC address that is assigned to a port. If you enable MAC Address Lock, you can also enable Intrusion Detection, which generates trap messages that identify intruding (unknown) source addresses.

The switch stores a history of 64 intruding source MAC addresses in a software buffer. One intruding source MAC address is stored and one trap message is generated per the intrusion trap timer setting. Once a MAC address is stored and the trap is generated, the switch does not send another trap for the address until it is cleared from the buffer.

*** Note:** The default setting for the intrusion trap timer is 1800 seconds (30 minutes). The valid range for the timer is 60 to 1800 seconds.

Web Agent Procedure

To use the Web Agent to enable the MAC Address Lock feature and traps for unknown source addresses:

1. Manually add the "permitted" MAC addresses to the Address Forwarding Table (AFT).

   For information on how to manually add MAC address to the AFT, see "Adding Static Entries to the AFT" in Chapter 9, "Managing the Address Forwarding Table."

2. Open the Switch Port Configuration Web page for the port.

   For information on how to open this Web page, see "Configuring Switch Ports Settings."

3. Set the **Allow Learning** field to **Disable**.

4. Set the **Known Mode** field to **Enable**.

   ***Note:** To prevent the flooding of frames that have unknown destination addresses, set the **Known Mode** field to **Enable** for all ports on the same VLAN as the port for which you are enabling MAC address lock.

5. Set the **Intrusion Trap** field to **Enable**.

6. Adjust the trap timer, if desired.

   The default setting for the intrusion trap timer is 1800 seconds (30 minutes). The valid range for the timer is 60 to 1800 seconds.

7. Click **Apply**.

CLI Commands

To use the CLI to enable the MAC Address Lock feature and traps for unknown source addresses:

1. Enter Global Configuration mode.

   The CLI displays the `(configure)#` prompt.

2. Use the following command to manually add the "permitted" MAC addresses to the Address Forwarding Table (AFT):

   set aft entry <*mac address*> VLAN <*VLAN-id*> port-binding forward <*mod-port-spec*>

3. Use the following command to disable learning on a port or port range:

   set port allow-learning <*mod-swport-range*> disable

4. Use the following command to enable known mode on a port or port range:

   set port known-mode <*mod-swport-range*> enable

   ***Note:** To prevent the flooding of frames that have unknown destination addresses, enter all ports on the same VLAN as the port for which you are enabling MAC address lock.

5. Use the following command to enable intrusion traps on a port or port range:

   set port intrusion-trap <*mod-swport-range*> enable

6. Use the following command to set the intrusion trap timer to other than the default setting:

   set port intrusion-trap-timer <*mod-swport-range*> <*intrusion-trap-timer-value*>

Using the All Module Ports Configuration Web Page

The All Ports Configuration Web page allows you to apply the same parameter settings to all switch ports in a module using a single command.

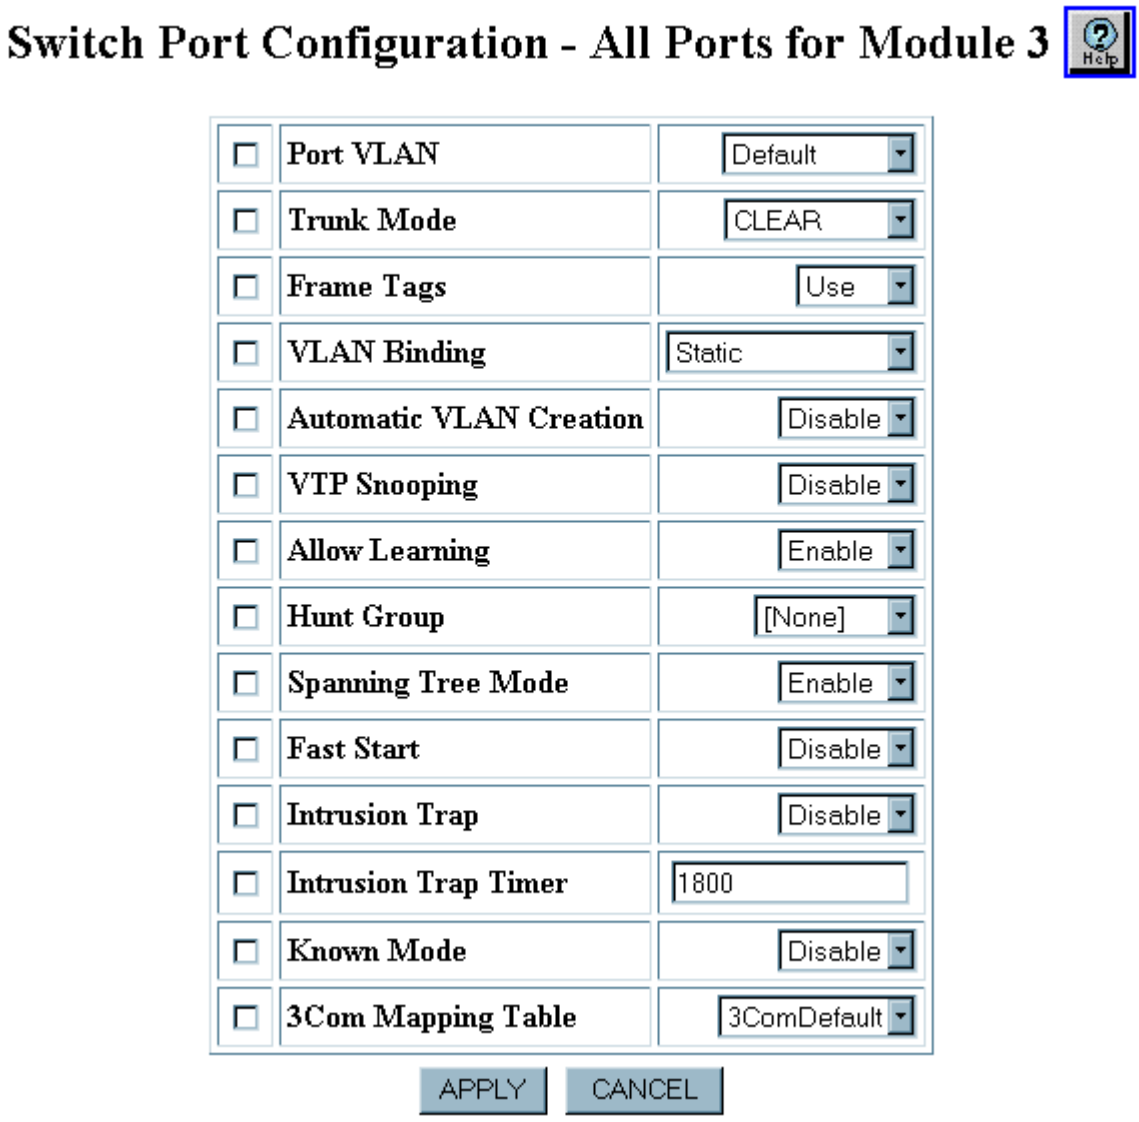
To set all switch ports in a module using a single command:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Configuration**. The Module Information Web page is displayed in the content pane. See Figure 8-2.

2. Select the switch port number listed for that module from the **Switch Ports** column. The Switch Ports Web page is displayed.

3. Select **All Module Switch Ports Configuration**. The Switch Ports Configuration - All Ports Web page is displayed. See Figure 8-10.

Figure 8-10. Switch Ports Configuration - All Ports Web Page



4. Select the check box next to the field you want to configure.
5. Select an option for the specific field from that fields pull-down menu.
6. Select **Restore** to keep the previous settings or **Apply** to accept the new settings.

Viewing Switch Port Settings

You can view switch port parameters through either the Web Agent and the CLI.

Web Agent Procedure

To view switch port parameters using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Configuration**. The Module Information Web page is displayed in the content pane. See Figure 8-2.
2. Select the switch ports for that module from the **Switch Ports** column. The Switch Ports Web page is displayed.
3. See Table 8-10 for more information about the **Switch Port** Web page parameters:

Table 8-10. Switch Port Parameters

Parameter	Definition
Links	Displays associated Web pages.
Port	Displays the port associated with the selected module.
Name	Displays the port name and displays the Switch Port Configuration Web page for the selected module.
Port VLAN	Displays the port VLAN for the selected module.
VLAN Classification	Displays the port VLAN classification for the selected module.
Trunk Mode	Displays the port's trunk mode for the selected module.
Hunt Group	Displays the hunt group of which the port is a member for the selected module.
STAP Mode	Displays whether the spanning tree algorithm protocol is enabled or disabled for the selected module.
MAC Address	Displays the port's MAC address for the selected module.

4. Select one of the following items at the bottom of the Web page for more information about the switch ports:

   — **Next/Previous Module** – Displays the next or previous module's switch port parameters.

   — **Modules** – Returns you to the Module Information Web page.

   — **All Module Switch Ports Configuration** – Displays the Switch Port Configuration All Ports Web page and configures all ports for the selected module.

CLI Command

To view switch port parameters using the CLI, enter the following command from Priv mode:

> show port [{<*mod-num*> | <*mod-swport-range*>} [...,{<*mod-num*> | <*mod-swport-range*>}]]

Configuring Port Redundancy

Overview

You can define two ports as a redundant pair. A redundant pair consists of a primary and secondary port. Only one port is active at a time. In its normal state in a redundant pair, the primary port transmits and receives data and the secondary port is disabled, neither transmitting nor receiving data. The secondary port does not perform load sharing. The primary and secondary ports can be on different types of modules. For example, the primary port can be 10/100 Ethernet and the secondary port can be Gigabit Ethernet. Redundancy for ATM ports is not supported.

If the primary port fails, the secondary port becomes the active port. Failover time is less than 5 seconds. If the secondary port is the active port, there is no change back to the primary port after it recovers. If both ports fail, the first port to recover becomes the active port.

You can globally disable or enable redundancy pairs. When disabled, the active port remains active and the disabled port remains disabled. When port redundancy is enabled, the active port continues in that capacity. There is no change to the primary port if the secondary port is the active port. When you enable or disable redundancy pairs, you enable or disable all active pairs.

A maximum of 20 pairs can be configured and active at any one time. You must create a unique name for your pair. Port redundancy cannot be enabled if Spanning Tree or Rapid Spanning Tree is enabled.

This section contains procedures for the following tasks:

- Creating a Port Redundancy Pair
- Deleting a Port Redundancy Pair
- Disabling or Enabling Port Redundancy

Creating a Port Redundancy Pair

Overview

You can use the Web Agent or CLI to create a port redundancy pair. The switch supports a maximum of 20 pairs.

You must create a unique name for the pair. If the name is not unique to the pair, it is assumed that the existing pair is being modified

After creating a redundancy pair, enable port redundancy globally for all configured pairs. See "Disabling or Enabling Port Redundancy" for more information.

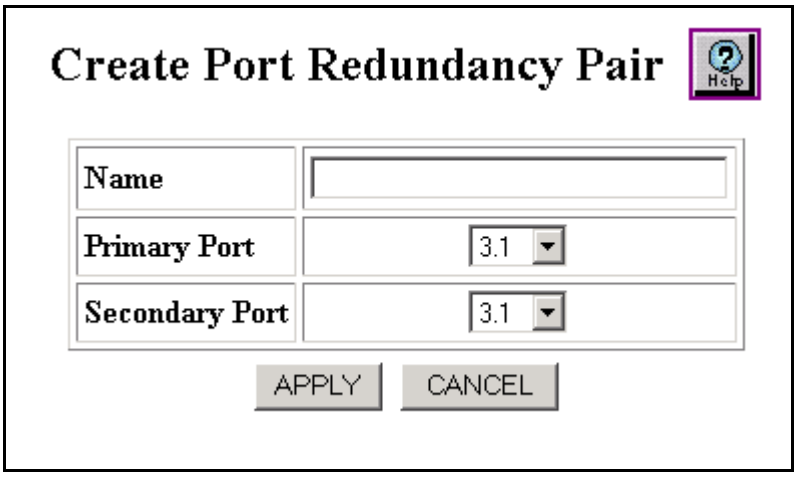
*** Note:** You must globally disable Spanning Tree and Rapid Spanning Tree before you can create a port redundancy pair.

Web Agent Procedure

To create a port redundancy pair

1. In the navigation pane, expand the **L2 Switching** folder, and then click **Port Redundancy**. The Port Redundancy Configuration Web page is displayed in the content pane.
2. Click **Create**. The Create Port Redundancy Pair Web page is displayed.

Figure 8-11. Create Port Redundancy Pair Web Page



3. Complete the following:
   - **Name**—Enter the pair's name.
   - **Primary Port**—Select the pair's primary port from the pull-down menu.
   - **Secondary Port**—Select the pair's secondary port from the pull-down menu.
4. Click **Apply**.

CLI Command

To create a port redundancy pair, use the following CLI command:

`(configure)#` **set port-redundancy name <*redundant-name*> <*primary-port*> <*secondary-port*>**

Deleting a Port Redundancy Pair

Web Agent Procedure

You can delete a port redundancy pair from the Port Redundancy Configuration Web page.

To delete a port redundancy pair:

1. In the navigation pane, expand the **L2 Switching** folder, and then click **Port Redundancy**. The Port Redundancy Configuration Web page is displayed in the content pane.
2. Select the redundancy pair(s) you want to delete.
3. Click **Delete**.

The redundancy pair is deleted.

CLI Command

To delete a port redundancy pair with the CLI, use the following command in Configure mode:

`(configure)#` **no port-redundancy <*pair-name*>**

Disabling or Enabling Port Redundancy

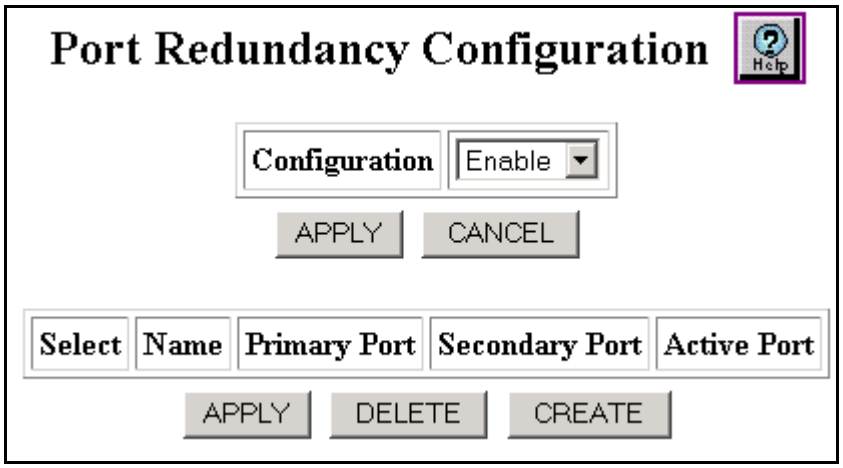
Web Agent Procedure

You can globally enable or disable port redundancy in the Port Redundancy Configuration Web page. Disabling redundancy does not delete the pairs.

To enable or disable a port redundancy

1. In the navigation pane, expand the **L2 Switching** folder, and then click **Port Redundancy**. The Port Redundancy Configuration Web page is displayed in the content pane.

Figure 8-12. Port Redundancy Configuration Web Page



2. In the **Configuration** field, select one of the following from the pull-down menu:
   - Enable—Enables all existing redundancy pairs.
   - Disable—Disables all existing redundancy pairs.

All existing redundancy pairs are enabled or disabled.

CLI Command

You can globally enable or disable port redundancy with the CLI. Disabling redundancy pairs does not delete the pairs.

To enable or disable port redundancy, use the following CLI command:

`(configure)#` **set port-redundancy {enable | disable}**

GBIC Identification

The GBIC Identification feature identifies the type of GBICs that are installed in 80-series gigabit modules with GBIC interfaces. The feature identifies the following types of connectors:

- GBIC SX (short wavelength)
- GBIC LX (long wavelength)
- GBIC LX/LH (long wavelength/long haul)

If any nonstandard connectors are present, the feature identifies them as GBIC LX/LH. You can view this information by using the Web Agent, CLI, or SNMP.

Web Agent Procedure

To use the Web Agent to view the type of GBICs that are connected to the modules in your switch, open the Physical Port Configuration Web page. The type of GBIC that is connected to each port is displayed in the **Connector** field on this Web page.

To open the Physical Port Configuration Web page:

1. In the navigation pane, expand the **Modules & Ports** folders, and then click **Configuration**.

   The Module Information Web page is displayed in the content pane.

2. In the **Ports** column, click the number for the module that you want to view GBIC types for.

   The Physical Port Configuration Web page is displayed. See Figure 8-5 for Gigabit modules and Figure 8-7 for 10/100 modules.

CLI Command

To view the type of GBICs that are connected to the modules in your switch, use the following CLI command:

> show port physical [{<*mod-num*> | <*mod-swport-range*>} [...,{<*mod-num*> | <*mod-swport-range*>}]]

The type of GBIC that is connected to each port is displayed in the `Connector` field.

Network Error Detection and Recovery

Overview

Functionality

The Network Error Detection and Recovery (NEDR) feature monitors Ethernet CRC errors on Ethernet data ports and compares the rate of errors that occurs to the threshold values that you set. This feature helps you find the source of the errors and take corrective action if necessary.

If the rate of CRC errors equals or exceeds the threshold value, the NEDR feature can:

- Disable the port
- Log the event in the event log

*** Note:** The Link LED on the 10-Gigabit module does not indicate a disabled status. The LED should flash green when you or the Network Error Detection and Recovery (NEDR) feature disable the port. However, the LED remains solid green.

If NEDR shuts down a port, the switch will forward traffic through a redundant port if you:

- Enable a protocol that supports redundancy, such as OSPF, VRRP, or STP, on the port
- Provide an alternate path

If you do not enable a protocol that supports redundancy on the port and provide an alternate path, you may lose traffic if the port shuts down.

If you replace a module that has a port that NEDR has disabled, the switch will load the startup configuration for the module and reenable the port. However, if you save the running configuration to the startup configuration while a port is shut down, you will have to manually reenable the port after you replace the module or reset the switch.

To reenable a turned off port, use the **set port enable** CLI command.

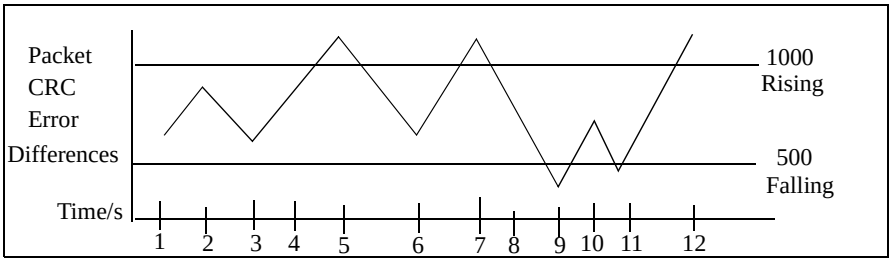
Hardware Support

You can enable NEDR only for ports on gigabit and 10-gigabit Ethernet modules and only by using CLI commands. You cannot enable the feature on any other types of ports, for example, Ethernet 10/100, ATM Uplink, or supervisor, and cannot enable the feature by using the Web Agent.

*** Note:** If the switch has an 80-series supervisor module installed, monitoring begins as soon as you turn on the switch. If the switch has a layer 3 supervisor module installed, monitoring may not begin until approximately 35 seconds after you turn on the switch.

Example

Figure 8-13. Example of How NEDR Works



In Figure 8-13, the rising threshold is set to 1000, and the falling threshold is set to 500. The first time that the rate of CRC errors exceeds the rising threshold, NEDR logs the event in the event log or disables the port. If you set the feature to **notify**, NEDR does not log another event until the rate of CRC errors drops below the falling threshold and then exceeds the rising threshold again. If you set NEDR to **disable-port**, the feature disables the port the first time that the rate of CRC errors exceeds the rising threshold. The port remains disabled until you reenable it, or until you remove and replace/reinsert a module of the exact type.

Enabling NEDR

You can enable NEDR only for ports on gigabit and 10-gigabit Ethernet modules and only by using CLI commands.

To enable NEDR:

1. Enter Global Configuration mode.

   The CLI displays the `(configure)#` prompt.

2. Enter the following command:

   set port network-error-detection {<*mod-port-range*> | all} [action {notify | disable-port | off}] [rising-threshold <*rising-threshold-value*>] [falling-threshold <*falling-threshold-value*>] [interval <*interval-seconds*>]

See Table 8-11 for explanations of the keywords, arguments, and options in this command.

The CLI displays a confirmation message if you successfully enable NEDR for a port or range of ports. If no message displays, the port information was invalid.

*** Note:** Occasionally, if you enable NEDR on a 10-Gigabit port, the following error message is displayed:

`Set port network-error-detection failed:` *<10-Gigabit port number>*`.`

NEDR is successfully enabled regardless of the error message. Ignore the error message.

3. To verify that NEDR is enabled on a port, enter the following CLI command:

show port network-error-detection [<*mod-port-range*>]

If the port is not listed as enabled for NEDR, reenable NEDR on the same port or ports again.

Table 8-11. Keywords, Arguments, and Options

Keywords, Arguments and Options	**Explanation**
<mod-port-range>	The slot number of the module, and, either port number, or range of port numbers. Enter the port ranges in the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**.
all	Disables NEDR on all ports on all modules in the switch. **all** can be used only with **off**. **Note:** The **set port network-error-detection all action off** CLI command does not turn off NEDR on 10-Gigabit ports. To turn off network error detection for specific 10-Gigabit ports, use the **set port network-error-detection <*mod-port-range*> action off** command.
	1 of 2

Table 8-11. Keywords, Arguments, and Options

Keywords, Arguments and Options	Explanation
action {notify | disable-port | off}	Action that NEDR performs when the rate of errors exceeds the threshold. The options are: • **notify**—Logs the event in the event log • **disable-port—**Disables the port and logs the event in the event log. A port is disabled if the rate of errors equals or exceeds the threshold. Make sure a redundant protocol is configured. • **off—**Disables NEDR on the port or ports that you specify. The default setting is **notify**.
<*rising-threshold value*>	The rising threshold. The number of CRC errors that triggers NEDR to log an event in the event log or disable the port. The default setting is 100 (minimum is 1; maximum is 65535).
<*falling-threshold value*>	The falling threshold. After exceeding the rising threshold, NEDR does not log another event in the event log until the rate of CRC errors falls below the falling threshold and then exceeds the rising threshold again. The default setting is half the rising threshold value (minimum is 0; maximum is 65535).
<*interval-seconds*>	How often NEDR checks the number of errors occurring against the thresholds. Enter a number of seconds. The default setting is 2 seconds (minimum is 1; maximum is 65535).
	2 of 2

*** Note:** If you set the rising threshold value and the falling threshold value close together, events may be logged more often if the **Notify** option is selected.

*** Note:** No configuration event is generated in the event log when you enable or disable NEDR. However, you can find out which ports have NEDR enabled by using **show port network-error-detection** CLI command.

Internal Error Detection and Recovery

Overview

The Internal Error Detection and Recovery (IEDR) feature shuts down a port if its rate of internal errors exceeds the threshold setting. Currently you can enable this feature only by using the CLI and only for ports on 50-series and 80-series Gigabit Ethernet modules. By default IEDR is disabled.

*** Note:** The 10-Gigabit module does not support IEDR.

This feature is designed to serve as mechanism for detecting and isolating hardware failures.

If IEDR shuts down a port, the switch will forward traffic through a redundant port if you:

- Enable a protocol that supports redundancy, such as OSPF, VRRP, or STP, on the port
- Provide an alternate path

If you do not enable a protocol that supports redundancy on the port and provide an alternate path, you may lose traffic if the port shuts down.

If you replace a module that has a port which has been disabled by IEDR, the switch will load the startup configuration for the module and reenable the port. However, if you save the running configuration to the startup configuration while a port is shut down, you will have to manually reenable the port after you replace the module or reset the switch.

This section contains the following information and procedures:

- Enabling IEDR
- Enabling IEDR for Hunt Groups
- Setting the IEDR Threshold for Internal Errors
- Viewing IEDR Settings
- Viewing IEDR Settings for Hunt Groups
- Error Messages
- System Messages

Enabling IEDR

*** Note:** To enable IEDR for ports in a hunt group, you must use a different command that enables the feature globally for all ports in all hunt groups on the switch. For information on this command, see "Enabling IEDR for Hunt Groups."

*** Note:** Do not use ISL Tagging (Trunk Mode: Multilayer) on ports that have IEDR enabled.

IEDR is disabled by default on gigabit ports. To enable the feature:

1. Enter Global Configuration mode.

   The CLI displays the `(configure)#` prompt.

2. Enter the following command:

   set port internal-error-shutdown {<*slot*>/<*port*> | <*slot*> | all-ports} enable

To disable IEDR, enter:

set port internal-error-shutdown {<*slot*>/<*port*> | <*slot*> | all-ports} disable

Enabling IEDR for Hunt Groups

This procedure globally enables IEDR for all ports in all hunt groups on the switch.

*** Note:** All ports that have IEDR enabled, whether they are administratively disabled or enabled, assume the hunt group IEDR setting if they are placed in a hunt group. If a port has IEDR enabled *before* you place it in a hunt group, the **show port internal-error-config** command displays the port as IEDR-enabled regardless of the hunt group IEDR setting. However, the port in fact assumes the hunt group IEDR setting (whatever that setting is).

To globally enable IEDR for all ports in hunt groups:

1. Enter Global Configuration mode.

   The CLI displays the `(configure)#` prompt.

2. Enter the following command:

   set huntgroup internal-error-shutdown enable

To globally disable IEDR for all ports in hunt groups, enter:

set huntgroup internal-error-shutdown disable

Setting the IEDR Threshold for Internal Errors

You can set the threshold at which the switch will shutdown a port. By default this threshold is set to 10 internal errors in a 5-second time period. You can set the threshold to any number between 5 and 500 internal errors in a 5-second time period. This setting is global for all ports that have been enabled for IEDR including ports configured for hunt groups.

To set the IEDR threshold:

1. Enter Global Configuration mode.

   The CLI displays the `(configure)#` prompt.

2. Enter the following command:

   set internal-error-threshold <*threshold*>

Viewing IEDR Settings

To view a list of the ports that have IEDR enabled:

1. Enter Global Configuration mode.

   The CLI displays the `(configure)#` prompt.

2. Enter the following command:

   show port internal-error-config

Viewing IEDR Settings for Hunt Groups

To view the global IEDR setting for hunt groups:

1. Enter Global Configuration mode.

   The CLI displays the `(configure)#` prompt.

2. Enter the following command:

   show huntgroup internal-error-config

Error Messages

*** Note:** If IEDR detects errors on a port in the hunt group, it tests each port within that hunt group. When it finds the first bad port, it removes it from the hunt group and disables the port.

The IEDR feature may display the following error messages:

- The following message indicates the IEDR feature is enabled on the specified ports
  - `internal-error-config enabled on port <x>/<y>`
- The following message indicates that internal CRC errors are occurring:
  - `Description: CRC err: Rx'ed 5 bad pkts since last poll - last from port3.2`
- The following message indicates that IEDR shut down a port because its rate of internal errors exceeded the threshold:
  - `Description: Shutdown Port 3.2 due to CRC error exceeding threshold of 10`

System Messages

The **set port internal-error-shutdown all-ports enable** command produces system messages similar to the following:

- Upon success:

```
Internal-error-config enabled on port 3/1
Internal-error-config enabled on port 3/2
Internal-error-config enabled on port 4/1
Internal-error-config enabled on port 4/2
Internal-error-config enabled on port 4/3
Internal-error-config enabled on port 4/4

Slot 6 has 10/100 ethernet ports that are
unsupported and will not be set.
```

*** Note:** IEDR is supported only for 50-series and 80-series Gigabit modules.

```
Internal-error-config enabled on port 7/1
Internal-error-config enabled on port 7/2
Internal-error-config enabled on port 7/3
Internal-error-config enabled on port 7/4

Feature enabled for port(s) specified

If CRC Error Threshold is exceeded then port
will automatically be shut down.

TRAFFIC LOSS WILL RESULT if no redundant port
is configured and port shutdown occurs.

To UNDO use: "set port internal-error-
shutdown {slot/port} disable"
```

- Upon Failure:

The **set port internal-error-shutdown all-ports disable** command produces system messages similar to the following:

- Upon Success:

```
Internal-error-config disabled on port 3/1
Internal-error-config disabled on port 3/2
Internal-error-config disabled on port 4/1
Internal-error-config disabled on port 4/2
Internal-error-config disabled on port 4/3
Internal-error-config disabled on port 4/4

Slot 6 has 10/100 ethernet ports that are
unsupported and will not be set.

Internal-error-config disabled on port 7/1
Internal-error-config disabled on port 7/2
Internal-error-config disabled on port 7/3
Internal-error-config disabled on port 7/4

Feature disabled for port(s) specified

Port(s) will not be shut down if CRC Threshold
Exceeded.
```

- Upon Failure:

```
Don’t Shutdown all ports failed
```

The **set internal-error-threshold** <***5-500***> command produces system messages similar to the following:

- Upon Success:

  ```
  Threshold value set to 10
  ```

- Upon Failure (The only failure that can occur is that the number entered was not in the range 5 – 500):

  ```
  Threshold value must be between 5 and 500
  ```

9 Managing the Address Forwarding Table

Overview

The Address Forwarding Table (AFT) is a mapping table of MAC addresses with their associated port locations.The AFT is used by the Avaya Multiservice switch to correctly forward (bridge) frames destined for a particular MAC address to the correct physical port. The AFT performs several major functions, learning new MAC addresses, aging out old MAC addresses, and providing a management interface to display, add, modify, and remove AFT entries (MAC addresses). AFT's are maintained throughout the system, on media modules and the supervisor module.

There is one AFT created for each VLAN. The entire Avaya Multiservice Switch can store up to 24,000 MAC addresses and 1000 VLANs.

*** Note:** Although the Avaya Multiservice Switch can support 1000 VLANs, (in Fabric mode 2 with all 80-series modules), restrictions apply for the size of the Hash tables, the size of the Address Forwarding Tables and the number of VLANs. This information is detailed in the following pages in this chapter.

- New address Learning - When a MAC address of a packet is unknown, it must be learned by the AFT
- Address Aging - On a periodic basis (which you can set) the addresses in the AFT are aged out, MAC addresses not received for a period of time (e.g. 5 minutes) are removed from the AFT along with their port associations.
- Static Address Management - Static MAC address entries are saved in memory upon module reset. These static address always exists in the AFT, and are not aged out or overwritten.Static entries are those manually configured. Each VLAN and its AFT has 18 AFT entries automatically entered by the switch that are reserved for internal use.
- Address Management - The user has the ability to display entries, add entries to the AFT, remove entries from the AFT, and change entries within the AFT. This is accomplished through the Web Agent interface or CLI commands. A user can also make an entry or entries static, which saves the entries upon module reset.

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Chapter contents

The following information and procedures are in this chapter.

- Relationship between VLANs, AFT and Hash Table Sizes
- Hash Table guidelines for creating VLANs
- Examples of Configuring VLANs and Hash Table Size
- AFT Default Settings
- Address Forwarding Table, Auto-Sizing, Auto Increment and Threshold
- Total Entries, Address Memory, Age and Super-Age Timers
- Instance Table Information
- Searching the Switch AFT
- Adding Static Entries to the AFT
- Option 82 for DHCP

Relationship between VLANs, AFT and Hash Table Sizes

Hash Tables and Buckets

For every VLAN created, a hash table is allocated. The hash table keeps track of the memory locations where the learned AFT entries (MAC addresses) are stored. These memory locations are referred to as *buckets*.

An individual *bucket* can range in size from 1 to 128 AFT entries (MAC addresses) in powers of 2 (1, 2, 4, 8, 16, and so on). The *bucket capacity* for a VLAN is the sum of all the capacities of the *buckets* assigned to that VLAN. If you add the number of AFT entries and divide by the *Bucket Capacity*, you will get the percentage of *Bucket Utilization* (Figure 9-2).

To view the *Bucket* parameters from the Web Agent, expand the **Layer 2 Switching** > **Address Forwarding Table** > **Configuration** folders. And view the **Bucket Capacity** and **Bucket Utilization** fields.

Hash Table and Bucket memory usage guidelines

The Avaya Multiservice switch provides 60K of memory that is used to store hash tables and AFT buckets. In general, no more then 20K should be used for Hash Tables, leaving 40K available for AFT entries (MAC addresses). It is important to take into consideration the number of VLANs that the switch will have when determining the hash table size for each VLAN, as the total amount of memory used for the Hash Tables should not exceed 20K. As more VLANs are created, the smaller each VLANs Hash Table should be.

It is not advised to allow the total hash table memory usage to exceed 20K. More than 20K will limit the potential for growth (i.e., address learning). If your Avaya Multiservice switch is approaching this threshold, adjust the Initial Hash Table Size accordingly.

For example, if you anticipate creating more than 20-30 VLANs, the Initial Hash Table Size default setting must be reduced. The rule of thumb is the following formula:

Number of VLANs *times* Initial Hash Table Size <20K

Example 1: 30 x 512 ~ 15K

Example 2: 1000 x 16 ~ 16K

*** Note:** When you create a VLAN, there are 18 internal MAC entries automatically created for that VLAN. If you were to configure Example 2, the switch would indicate that 36K of memory is currently in use. This is because you have 16K of memory for the VLAN Hash Tables and another 20K of memory used for the 18,000 MAC entries.

Hash Table guidelines for creating VLANs

Maximum Number of VLANs

In order to support the maximum number of VLANs, VLAN ID numbers should be chosen from the range of 1 to 1000. The numbering of VLANs has no impact on memory usage within the switch. The numbering of VLANs only effects the total number of supported VLANs. This is not a memory usage issue, however, it does effect the way the switch uses or Hashes VLAN ID’s. If VLAN ID’s are used outside of the range of 1 to 1000, there is a possibility of unavailable VLAN ID’s. This again does not effect switch memory usage. If VLANs are already created with VLAN ID’s outside of the 1-1000 range, there is no need to reassign new VLAN ID’s. The only potential issue could be particular VLAN ID’s might not be available.

VLAN ID number outside the recommended range

If you use an ID number outside the range of recommended numbers, you may find certain VLAN ID numbers cause “collisions” and are unavailable when you try to create a new VLAN.
The reason for this has to do with the way that the switch hashes VLAN numbers into internal tables. The ‘hash’ algorithm used by the switch takes VLAN numbers (which have a range from 1 to 4094) and ‘hashes’ them to a range of 1 to 1024 by picking 10 bits from the VLAN number. This smaller range is used as an index into tables within the switch, and is referred to as the ‘index’. Because the algorithm maps 4096 VLAN numbers to 1024 indexes, it is easy to see that more then one VLAN number may map to the same index. This is known as a ‘hash collision’.

When a collision occurs, the switch attempts to resolve the hash collision by using 10 other bits from the VLAN ID. When the switch cannot find 10 bits that uniquely hash the VLAN number, an error is returned to the user that the VLAN number cannot be added. In order to avoid this situation, it is suggested that the user pick VLAN numbers only from the range of 1 to 1000, as when this is the case, it is guaranteed that the switch can hash this set of VLAN numbers without having any collisions that cannot be resolved. Specifically, the 10 least significant bits are picked, which hash each VLAN number (in the range of 1 to 1000) to an index that is the same as the VLAN number. In this situation there will never be any collisions. See the Maximum Number of VLANs Supported section for valid ranges.

Reassign VLAN IDs?

It is not recommended to reassign VLAN numbers in a switch that is currently using VLAN numbers outside the range of 1 to 1000. The only potential issue is an attempt to add more VLANs, it is possible that particular VLAN numbers will not be available to use.

Optimal Bucket Utilization and Hash Table size

To achieve optimal Bucket Utilization, Hash Table size should be ¼ times the number of expected AFT entries (MAC entries) on the VLAN.

*** Note:** Each VLAN has 18 AFT entries automatically entered by the switch that are reserved for internal use. (See Chapter 1, “Introduction,” for a detailed listing of Self Addresses).The total amount of AFT memory used for Hash Tables should not exceed 20K.

Maximum Number of VLANs

Table 9-1 shows the maximum number of VLANs that are supported depending on the version of code that is running and the Fabric Mode that the switch is operating in.

*** Note:** Application software earlier than v5.0 does not support Fabric modes. Fabric mode 2 supports only 80-series hardware.

If the number of VLANs supported in the switch is 500, then the VLAN ID numbers used in the switch should be chosen from the range of 1 to 500 rather then 1 to 1000.

Table 9-1. Number of Supported VLANs

Version of Code	Fabric Mode 1	Fabric Mode 2
Pre-5.0	1000	N/A
5.0-Pre 5.2.10	500	500
5.2.10 and above	500	1000

Table 9-2 is an example of the size of the hash table created per the number of VLANs created using the guidelines.

Table 9-2. Number of VLANs and Hash Table Size:

Number of VLANs	Hash Table Size
2	8192
4	4096
10	2048
20	1024
40	512
80	256
150	128
300	64
600	32
1000	16

Examples of Configuring VLANs and Hash Table Size

Example 1

All VLANs have the *same* Hash Table Size

Configure the Avaya Multiservice switch for 45 VLANs. This requires the default Hash Table size be reduced according to the hash table guidelines.

45 VLANs *times* 1024 = 46,080 or ~ 46K

This exceeds the 20K guideline for the amount of AFT memory used for Hash Tables. Even if the Hash Table size is reduced by half, it still exceed the guideline. In order to comply with the guideline, the Initial Hash Table Size for each VLAN should be 256 or less. 45 VLANs *times* 256 = 11,520 or ~ 12K This leaves 48K for AFT entries (60K - 12K = 48K).

Example 2

VLANs have the *different* Hash Table Sizes

Configuring the Avaya Multiservice switch for 6 “large” VLANs (4000 users on each), 30 “medium” VLANs (1000 users on each), and 50 “small” VLANs (200 users on each). Using the “Hash Table Guidelines”, each “large” VLAN should have a hash table size around (1/4)*4000=1000, each “medium” VLAN around (1/ 4)*1000=250, and each “small” VLAN should have a hash table size around (1/4)*200=50. We also want to make sure that the total amount of memory used for the Hash Tables is less then 20k, which it will be (~17K).

6 VLANs *times* 1024 = 6,144

30 VLANs *times* 256 = 7,680

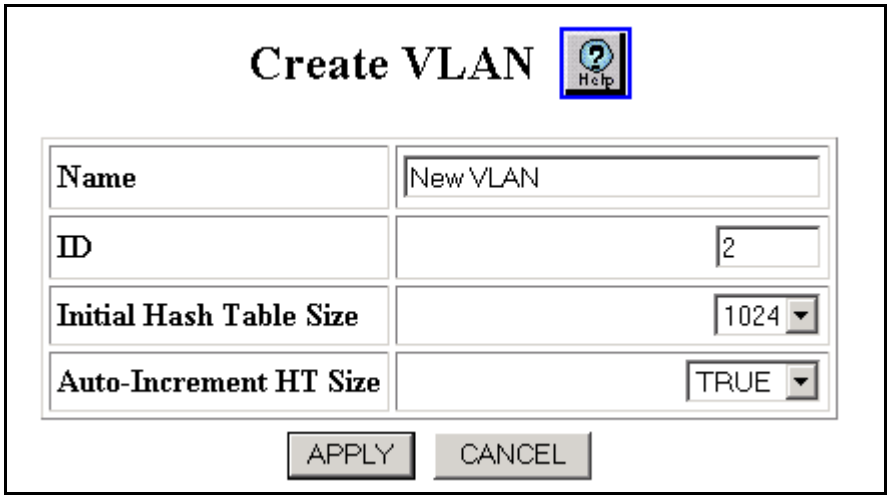
50 VLANs *times* 64 = 3,200

Total AFT Memory used or Hash Tables = 17,024 or ~ 17K

This meets the guideline.

AFT Default Settings

Figure 9-1. Create VLANs



Name—Name for the new VLAN

ID—VLAN ID number, valid number between 1-4094 (1 is reserved for the Default VLAN)

Instance ID—Instance table representation of a VLAN. The Instance ID does not directly correlate to the VLAN ID

Initial Hash Table Size—1024, the default setting for the Initial Hash table size.

Auto- Increment HT Size—Default setting is **True**. Enable to automatically correct for undersized hash table size.

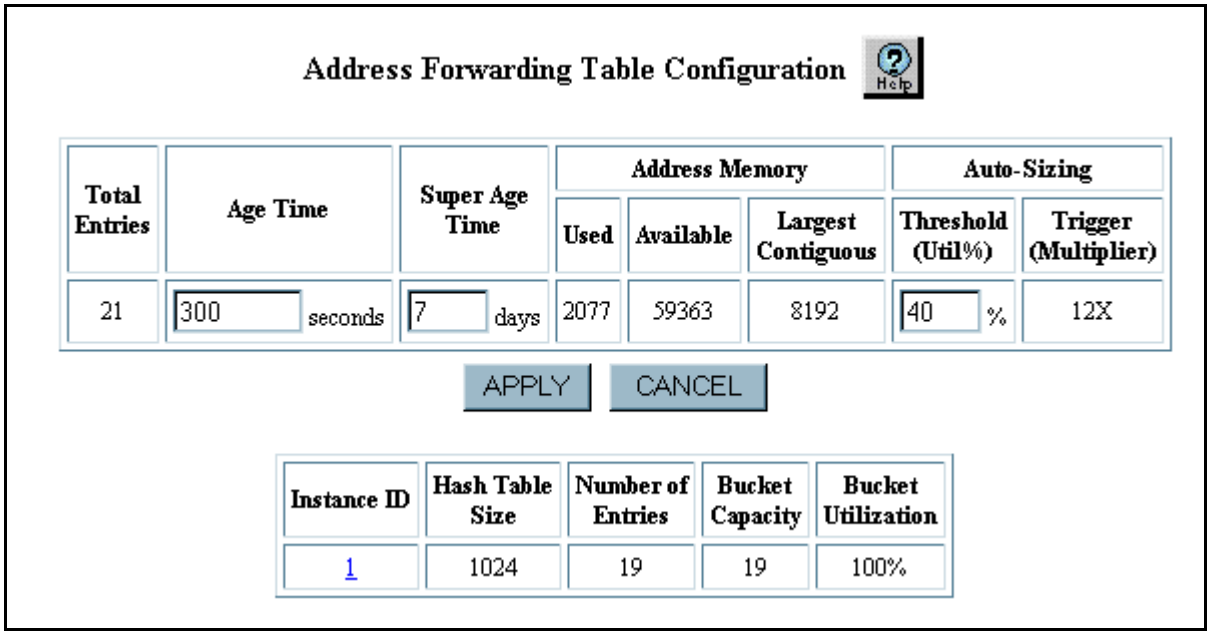
Address Forwarding Table, Auto-Sizing, Auto Increment and Threshold

Default Settings

In the navigation pane, expand the **L2 Switching** > **Address Forwarding Table** folders, and then click **Configuration**. The Address Forwarding Table Configuration Web page is displayed in the content pane. See Figure 9-2.

Auto-Sizing is a feature that allows the Avaya Multiservice Switch to automatically correct for an undersized hash table, which can result in poor AFT memory utilization. Auto-Sizing consists of three parameters: **Auto-Increment**, **Trigger (Multiplier)**, and **Threshold (Util%)** in the Address Forwarding Table Configuration

Figure 9-2. Address Forwarding Table Configuration



Auto-Increment

Auto-Increment is the Enable/Disable setting for the Auto-Sizing feature. If Auto-Increment is enabled, the switch will automatically increase the hash table size provided two conditions are met.See Figure 9-1 for enabling or disabling Auto-Increment HT size.

Trigger (Multiplier)

The first condition is the *Trigger (Multiplier)*, which sets a minimum threshold for the number of learned entries a VLAN must have before the switch will consider increasing the Hash Table size. This prevents the Avaya Multiservice Switch from resizing Hash Tables for VLANs who have only a few AFT entries. This threshold is 12 times the Hash Table size.

For Example: If a VLAN is created with an Initial Hash Table Size of 64, the trigger would be 12 *times* 64 = 768 learned entries. If there are more than 768 learned entries, the second condition is evaluated.

Threshold (Util%)

The second condition is met if Bucket Utilization drops below the *Threshold (Util%)*, which by default is 40%. The switch will compensate for the poor utilization by doubling the Hash Table. 40% is recommended for most applications.

Resizing results in the switch utilizing AFT memory more efficiently. When the switch resizes the Hash Table, all entries are flushed and must be relearned.

It is highly recommended the default of 40% not be changed. The affects of varying the Threshold (Util%) are as follows:

- *Increasing* Threshold: Gives better memory utilization at the expense of more frequent flushing, flooding, and relearning.
- *Decreasing* Threshold: Gives less efficient memory utilization but does not flush, flood and relearn as often.

Total Entries, Address Memory, Age and Super-Age Timers

Overview

You can see the Total Entries and Address Memory. You can configure address age, auto-sizing threshold, and Super Age timers. See Figure 9-2.

Total Entries

Display how many address entries are contained in the table.

Address Memory

Displays how much address memory is currently being used, the amount of memory still available, and the largest contiguous block of memory that can exist.

Address Age and Super Age timers

- **Address Age Time -** defines the length of time addresses remain active in the address forwarding table.
- **Super Age Time** - defines how long inactive addresses are stored in the address forwarding table before being deleted.

Configuring the AGE and Super Age Timers

Web Agent Procedure

You can configure the Age Timer and Super Age Timer by using either the Web Agent or the CLI.

Enter a value in the **Age Time** field. Aged out addresses become invalid until the switch sees another packet with the aged out entry's source address.The standards-recommended default is 300 seconds.

*** Note:** If you have ATM Uplink Modules in the chassis and you set the **Age Timer** value to less than 30 seconds, you must change the **Flush MAC Timer** on each ATM Uplink module to be less than or equal to the new **Age Timer** value.

Enter a new value in the **Super Age Time** field. The valid entries are 1-30 days. The Super Age Timer marks all invalid table entries, then checks to see if they remain invalid for the specified super age interval. This clears the table of entries that are no longer used. The default value is 7 days.

Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Commands

To change the aging values for all instances of the Address Forwarding Table from the CLI, enter one of the following commands from **Configure** mode:

- `(configure)#` **set aft agetime** <***age-time-value***>
- `(configure)#` **set aft super-agetime** <***age-time-value***>

Instance Table Information

The Instance Table Information, is a summary of a AFT instance or Hash Table settings and utilization per VLAN. See Figure 9-2.

Instance ID, Hash Table Size, Number of Entries, Bucket Capacity, and Bucket Utilization

Instance ID - Individual AFT Hash Table settings and utilization, each ID would correlate to a VLAN.

Hash Table Size - Memory space available for the Hash Table per Instance or VLAN

Number of Entries = The Total Number of Entries (MAC Addresses) for that VLAN

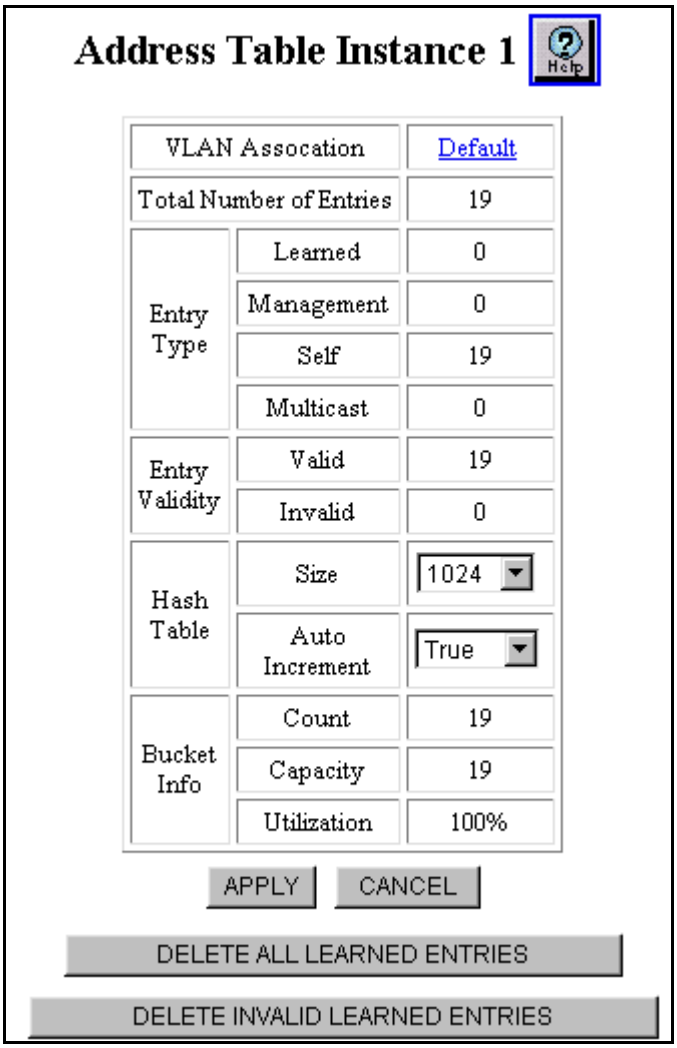
Bucket Capacity = The bucket capacity for a VLAN is the sum of all the capacities of the buckets assigned to that VLAN.

Bucket Utilization = The number of entries divided by the Total Bucket Capacity times 100 =% of utilization.

Instance Table

Click the **ID** number in the **Instance IDs** column, to achieve finer control of a particular VLAN table. The Address Table Instance Web page displays for the instance ID selected.(See Figure 9-3)

Figure 9-3. Address Table Instance Web page



See Table 9-3 for a definition of the Address Table Instance Web page parameters.

Table 9-3. Address Table Instance Parameters

Parameter	Definition
VLAN Association	Name of the VLAN this Hash table is associated with.
Total Number of Entries	Displays the total number of entries (MAC addresses) for this VLAN.

1 of 2

Table 9-3. Address Table Instance Parameters

Parameter	Definition
Entry Type	Displays the entry type for this VLAN. Options include: • **Learned** - Entry is dynamically learned. • **Management** - Entry is configured by the user statically. • **Self** - Entries for internal addresses contained within the switch. • **Multicast** - Entries belong to Multicast Groups.
Entry Validity	Displays the entry validity for this VLAN. Options include: • **Valid** - Entry exists. • **Invalid** - Entry has aged out but still exists even though the port binding is invalid.
Hash Table	Select one of the following hash table parameters: Size - Select the size to alter the space available for this address table. The selection range is 16 to 8192 incrementing in the powers of 2. **Note:** The number of addresses for a given hash table is 4:1 (for example, if you have a hash table of 16 bytes, the VLAN can hold 64 addresses in its table instance.To achieve optimal Bucket Utilization, Hash Table size should be ¼ times the number of expected AFT entries (MAC entries) on the VLAN. **Auto Increment** - Select an option to allow the hash table to grow dynamically larger if more addresses are discovered. The options are: **True** - Enables auto-increment of the hash table. **False** - Disables auto-increment to prevent the hash table from growing dynamically at the risk of extra flooding.
Bucket Info	Displays parameters to monitor the efficiency of the hash table allocations. Options include: • **Count** - The Total Number of Entries (MAC Addresses) for that VLAN • **Capacity** - Indicates bucket capacity. • **Utilization** - The number of entries divided by the total bucket capacity.
	2 of 2

3. Click:

   - **APPLY** to save your changes
   - **CANCEL** to restore previous settings
   - **Delete All Learned Entries** if you want to relearn the entire table
   - **Delete Invalid Learned Entries** if you want to delete all entries that are currently aged out.

*** Note:** If you change the Hash Table Size, the switch relearns all addresses in that table, causing the switch to flood packets for a few seconds.

Controlling Re-configuration of AFT Sizes Using the Web Agent

To manually control reconfiguration of AFT sizes using the Web Agent:

1. In the navigation pane, expand the **L2 Switching** > **Address Forwarding Table** folders, and then click **Configuration**. The Address Forwarding Table Configuration Web page is displayed in the content pane. See Figure 9-2.

2. Enter a new **Auto-Sizing Utilization Threshold** percentage in the Threshold (Util%) field, if you just want to cause the table reconfiguration to occur at a different level of usage efficiency. The default value of 40% is recommended for most applications.

   ***Note:** Raising the value in the **Threshold (Util%)** field might cause the table to be relearned more frequently, and makes address space usage more efficient.

Controlling Reconfiguration of AFT Sizes Using the CLI

To manually control reconfiguration of AFT sizes using the CLI, enter the following command in **Enable/Configure** mode:

```
(configure)# set aft auto-sizing-threshold
```

Searching the Switch AFT

The AFT can hold 24,000 entries on each switch. The switch Web Agent provides a utility that allows you to filter which addresses it displays, making the list more manageable. Multiple criteria can be selected to produce a sophisticated filter. The parameters are treated as “ands,” meaning that displayed addresses must meet all selected criteria.

You can search the AFT using the Web Agent or the CLI.

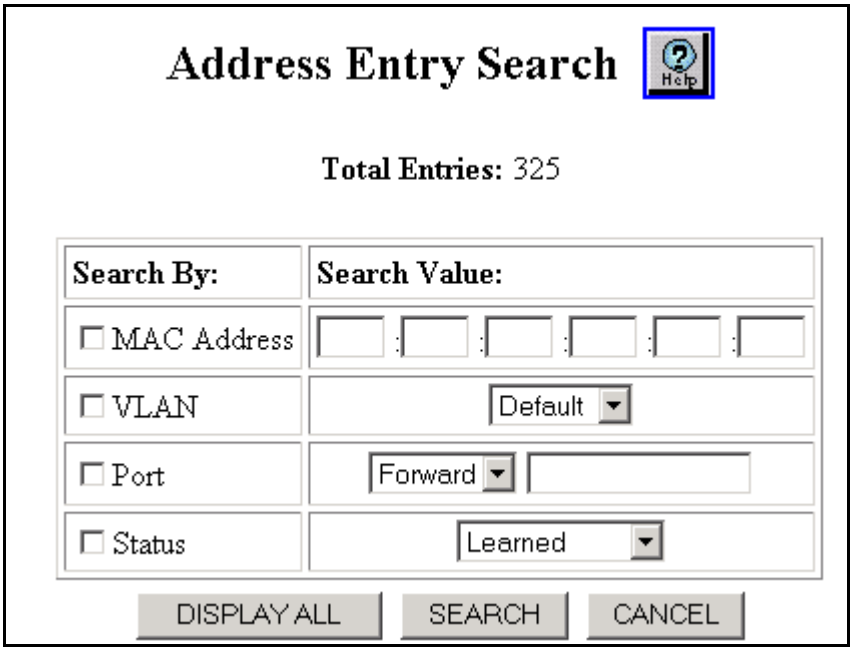
Web Agent Procedure

To filter the AFT using the Web Agent:

1. In the navigation pane, expand the **L2 Switching** > **Address Forwarding Table** folders, and then click **Address Search**.

   The Address Entry Search Web page is displayed in the content pane. See Figure 9-4.

Figure 9-4. Address Entry Search Web Page



***Note: DISPLAY ALL** ignores any parameter not checked in the left column. To view all addresses in the table, select **DISPLAY ALL** without selecting any filters.

2. To search using a MAC Address:

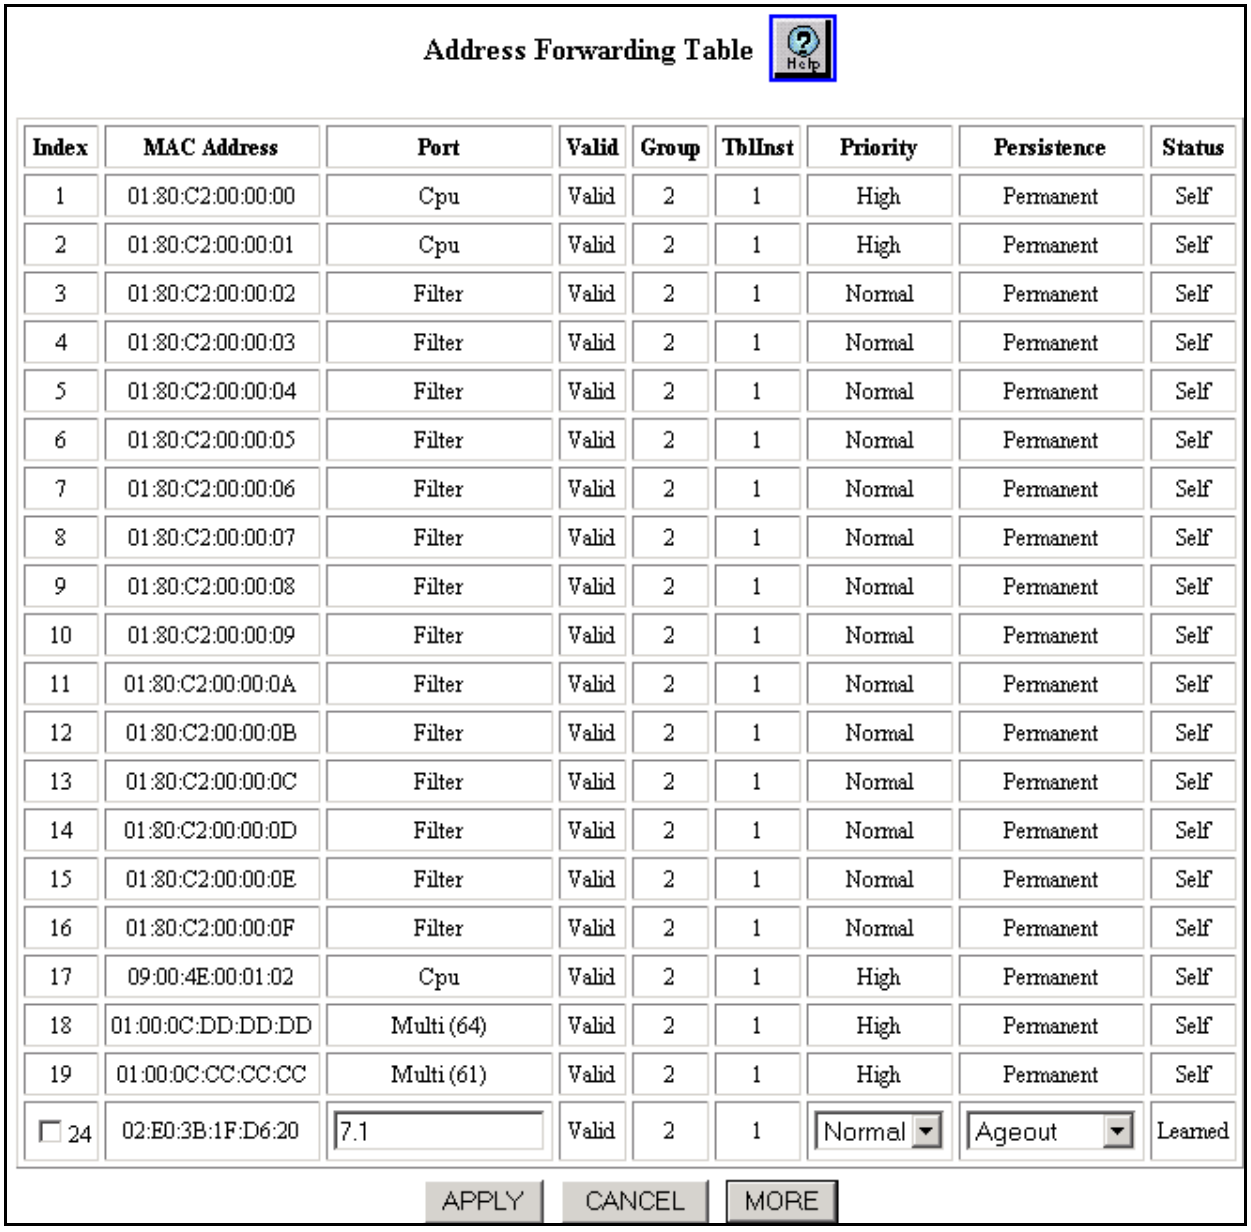
   — Select the MAC Address check box in the **Search By:** column

   — **Enter a MAC address in the** Search Value column.

   ***Note:** You must enter the complete MAC address.

3. To search using a VLAN:
   - Select the VLAN check box in the **Search By:** column
   - Select a VLAN name from the **VLAN** field pull-down menu in the **Search Value:** column.
4. To search using a Port:
   - Select the Port check box in the **Search By:** column
   - Select a search variable from the Port field pull-down menu. The options are: Forward, Filter, and CPU
   - Enter a port number to display only the entries associated with that switch port.
5. To search using a bridging status:
   - Select the **Status** check box in the **Search By:** column
   - Select a status from the pull-down menu to show only ports of a specific bridging status. The options are: Learned, Management, Self, and Multicast.
6. Select **DISPLAY ALL** to open the Address Forwarding Table Web page. See Figure 9-5.

Figure 9-5. Address Forwarding Table Web Page



Address Forwarding Table

Index	MAC Address	Port	Valid	Group	TblInst	Priority	Persistence	Status
1	01:80:C2:00:00:00	Cpu	Valid	2	1	High	Permanent	Self
2	01:80:C2:00:00:01	Cpu	Valid	2	1	High	Permanent	Self
3	01:80:C2:00:00:02	Filter	Valid	2	1	Normal	Permanent	Self
4	01:80:C2:00:00:03	Filter	Valid	2	1	Normal	Permanent	Self
5	01:80:C2:00:00:04	Filter	Valid	2	1	Normal	Permanent	Self
6	01:80:C2:00:00:05	Filter	Valid	2	1	Normal	Permanent	Self
7	01:80:C2:00:00:06	Filter	Valid	2	1	Normal	Permanent	Self
8	01:80:C2:00:00:07	Filter	Valid	2	1	Normal	Permanent	Self
9	01:80:C2:00:00:08	Filter	Valid	2	1	Normal	Permanent	Self
10	01:80:C2:00:00:09	Filter	Valid	2	1	Normal	Permanent	Self
11	01:80:C2:00:00:0A	Filter	Valid	2	1	Normal	Permanent	Self
12	01:80:C2:00:00:0B	Filter	Valid	2	1	Normal	Permanent	Self
13	01:80:C2:00:00:0C	Filter	Valid	2	1	Normal	Permanent	Self
14	01:80:C2:00:00:0D	Filter	Valid	2	1	Normal	Permanent	Self
15	01:80:C2:00:00:0E	Filter	Valid	2	1	Normal	Permanent	Self
16	01:80:C2:00:00:0F	Filter	Valid	2	1	Normal	Permanent	Self
17	09:00:4E:00:01:02	Cpu	Valid	2	1	High	Permanent	Self
18	01:00:0C:DD:DD:DD	Multi (64)	Valid	2	1	High	Permanent	Self
19	01:00:0C:CC:CC:CC	Multi (61)	Valid	2	1	High	Permanent	Self
24	02:E0:3B:1F:D6:20	7.1	Valid	2	1	Normal	Ageout	Learned

APPLY CANCEL MORE

*** Note:** Some types of entries in the Address Forwarding Table can be modified. The port, priority, and persistence can be changed for learned and management entries.

7. To change an entry in the Address Forwarding Table, select the check box in the Index column for MAC address entry.

8. Make the change(s) to the port, priority, or Persistence for the selected entry

9. After making the changes, click the **APPLY** button to save the current settings, or **CANCEL** to restore previous settings. or **MORE** to view more statistics.

10. See Table 9-4 to review your search criteria options:

Table 9-4. Address Forwarding Table Parameters

Parameter	Defines the...
Index	Index number of this address entry in the switch address forwarding table.
MAC Address	MAC address associated with this entry. This address is learned by the switch as an address to forward to the associated port.
Port	Port associated with this MAC address table entry.
Valid	An entry is valid until it ages out, at which time it becomes invalid. Aging out occurs when a frame with the entry's MAC address is not received during the address age time interval. An invalid entry can become a valid entry again if the entry's MAC address is learned (frame is received) during the super age time interval. If the super age time interval expires before the MAC address is learned again, the entry is deleted."
Group	Group number associated with this MAC address.
TblInst	Address table instance number associated with this MAC address.
Priority	Priority level associated with traffic forwarded to this MAC address. The options are: • **Normal** • **High** You can set this parameter on all learned entries.
	1 of 2

Table 9-4. Address Forwarding Table Parameters

Parameter	Defines the...
Persistence	The persistence of the entry in the table can be set on all learned entries: • **Permanent -** The address is not aged out of the table. • **Invalid -** This entry is cleared from the table each time the switch resets. • **Ageout -** Address is cleared from the address forwarding table when the timeout interval expires. This is the state of all entries dynamically learned by the switch. It ensures that MAC addresses that are not active on the network do not remain in the switch address forwarding table indefinitely. **Note:** If you create a static, filtered MAC address and set **Persistence** to **Ageout**, the switch does not filter the MAC address. The switch continues to forward frames that have the MAC address for a source address or destination address. To correct this, set Persistence to Permanent.
Status	The status of the address entry. Options include: • Learned • Management • Self AFT Self Entries: • 01:80:C2:00:00:00 - 01:80:c2:00:00:0F - 802.1D reserved addresses, the first entry is the Spanning Tree BPDU destination address. • 09:00:4E:00:01:02 - 3Com Vlan Trunk (VLT) control frame address. Used for Spanning Tree blocked ports that are in 3Com tagging mode and connected to a port on the other side of a connection that is blocked by Spanning Tree. • 01:00:0C:CC:CC:CC - Cisco control frame. Used to synchronize VLAN membership across switches (VTPs) or router information between routers (CDPs). The switch sends them to the CPU, the CPU will set the BPDU bit when transmitting out a port trunking ISL. • 01:00:0C:DD:DD:DD - If Cisco Group Management Protocol (CGMP) is enabled, the switch may listen to gather multicast group information
	2 of 2

CLI Command

To search the switch AFT using the CLI, enter the following command in **Configure** mode:

```
(configure)# show aft entry
```

Adding Static Entries to the AFT

You can add static entries to the AFT manually using the Web Agent or the CLI.

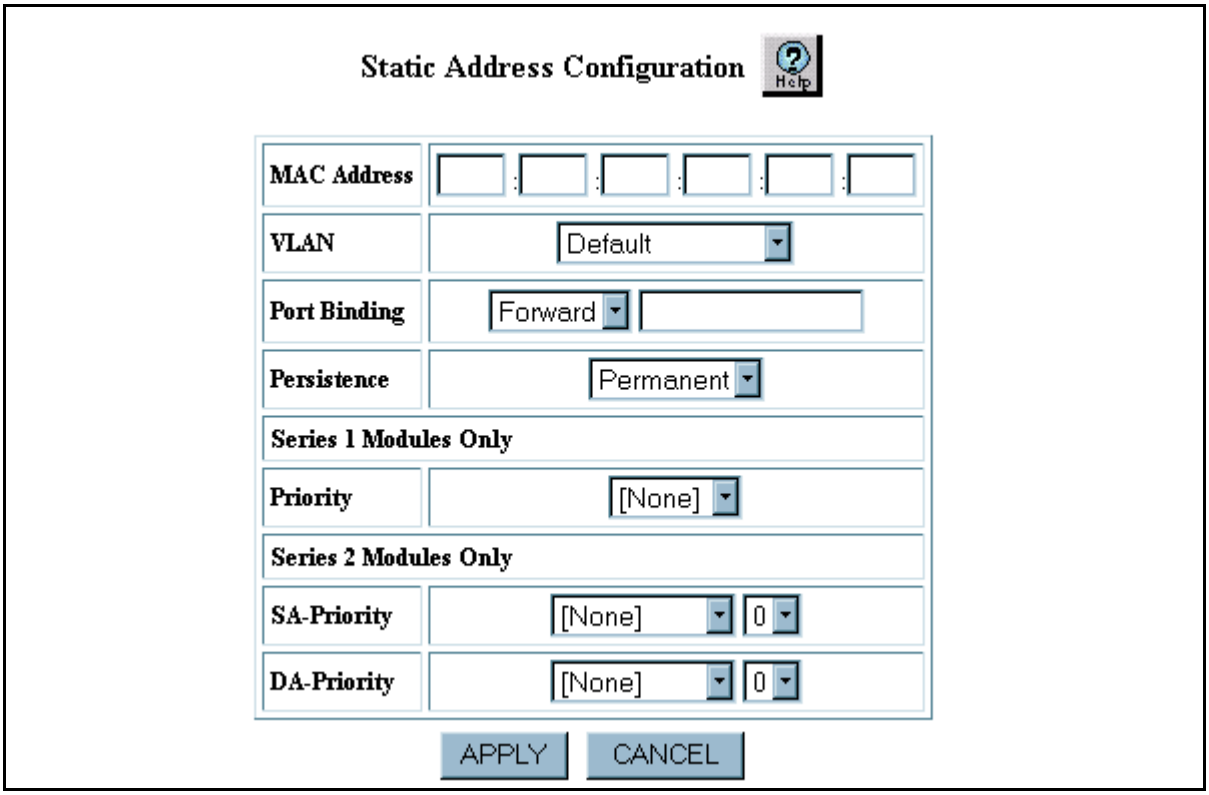
Web Agent Procedure

To add an AFT address manually using the Web Agent:

1. In the navigation pane, expand the **L2 Switching** > **Address Forwarding Table** folders, and then click **Address Configuration**.

   The Static Address Configuration Web page is displayed in the content pane. See Figure 9-6.

Figure 9-6. Static Address Configuration Web Page



2. In the **MAC Address** field, enter the MAC address that you want to add to the table.

3. In the **VLAN** field, select a VLAN for this entry.

4. In the **Port Binding** field, select **Forward** or **Filter**:

   - **Forward**—Frames that have this source or destination MAC address are forwarded. If you select **Forward**, you must enter the port that you want associated with this MAC address.

   - **Filter**—Frames that have this source or destination MAC address are dropped.

5. In the **Persistence** field, select **Permanent** or **Ageout**:

   - **Permanent**—The address is saved in nonvolatile memory and is not aged out of the AFT.

   - **Ageout**—Address is cleared from the AFT when the age time interval expires. All dynamically learned entries are set to **Ageout**. This setting ensures that MAC addresses that are not active on the network do not remain in the AFT indefinitely.

6. In the **Priority** field, select a priority level for packets that are forwarded to this MAC address. Options are **None**, **Normal**, and **High**.

   High priority addresses move to the front of the switch packet buffers automatically.

*** Note:** This field applies only to 50-series modules.

7. In the **SA-Priority** field, select the priority that you want associated with frames that have this source MAC address:

   - **None**

   - **Aft**—Uses the priority that is assigned to the MAC address in the AFT. If you select this option, you must select a numerical priority for the MAC address. Priorities range from 0 to 7.

   - **Max-port-aft**—Determines the priority of a frame by using the higher of the:

     — Physical port priority or tag priority

     — Source MAC address priority

     If you select this option, you must select a numerical priority for the MAC address. Priorities range from 0 to 7.

   - **Port**—Uses the priority of the physical port, Cisco ISL tag, or 802.1p tag.

*** Note:** This field applies only to 80-series modules.

8. In the **DA-Priority** field, select the priority that you want associated with frames that have this destination MAC address:

   - **None**
   - **Aft**—Uses the priority that is assigned to the MAC address in the AFT. If you select this option, you must select a numerical priority for the MAC address. Priorities range from 0 to 7.
   - **Max-port-aft**—Determines the priority of a frame by using the higher of the:
     - Physical port priority or tag priority
     - Destination MAC address priority

     If you select this option, you must select a numerical priority for the MAC address. Priorities range from 0 to 7.
   - **Port**—Uses the priority of the physical port, Cisco ISL tag, or 802.1p tag.

*** Note:** This field applies only to 80-series modules.

9. Click **APPLY** to save your changes, or **CANCEL** to reset fields to their default values.

CLI Command

To add entries to the AFT manually using the CLI, enter the following command in Configure mode:

`(configure)#` **set aft entry <*mac address*>**

Option 82 for DHCP

Overview

Option 82 for DHCP includes two suboptions:

- Suboption 1 identifies the slot and physical port number from which the DHCP request was received.
- Suboption 2 identifies the IP address and, if available, the system name of the switch.

By default, both of these suboptions are enabled. You can change the status of either or both suboptions by using the Web Agent, Command Line Interface (CLI), or SNMP.

Web Agent Procedure

To use the Web Agent to change the status of option 82:

1. Open the IP Global Configuration Web page.

   For information on how to open this Web page, see "Enabling IP Routing Global Parameters" in Chapter 12, "Configuring IP Routing."

2. Enable or disable the suboptions as appropriate in the **DHCP Option - Circuit Info** and **DHCP Option 82 - Agent Info** fields.

CLI Commands

To use the CLI to change the status of option 82, enter Global Configuration mode and use the following commands:

- To enable suboption 1, use the following command:

  `(configure)#` **ip dhcp circuit-info**

- To enable suboption 2, use the following command:

  `(configure)#` **ip dhcp agent-info**

- To disable suboption 1, use the following command:

  `(configure)#` **no ip dhcp circuit-info**

- To disable suboption 2, use the following command:

  `(configure)#` **no ip dhcp agent-info**

10 Configuring Redundancy Options

Overview

The Avaya Multiservice Switch can be configured with redundancy to provide fault tolerance.

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Supervisor Functions

The Avaya Multiservice Switch supervisor module manages the resources of the switch, provides access to these resources and supports a number of network protocols. These resources include configuration information, spanning tree topology, address forwarding tables, routing tables, and network statistics. Access is provided to these resources via console CLI (command line interface), Web interface, SNMP, and telnet.

Redundant Supervisors

The redundant supervisor module is an auxiliary “standby” supervisor module that acts as a fault-tolerant supervisor in the event that the “Active” supervisor fails. Once the redundant supervisor is installed, loaded with the same software version as the Active supervisor, and synchronized, it is ready to act as a redundant or “standby” supervisor. If the Active supervisor does fail, the redundant supervisor assumes control of network operation. No user intervention is required for the Supervisor module failover. The management view is accessible without changing IP or MAC addresses. The event is logged and indicated via an SNMP trap. Notification of this change is provided by a console message and an event log message.

Multiple Software Versions

Up to two software versions can be saved on the Avaya Multiservice Switch supervisor module.

Redundant Switch Controller and Element

The Avaya Multiservice Switch can be configured with a redundant switch element module and a redundant switch controller module, for backup of both switch fabric components. In the event of a failure in a switch element or controller module, the redundant module will take over the function of the failed component.

VRRP

The Avaya Multiservice Switch supports Virtual Router Redundancy Protocol (VRRP), an IETF protocol designed to support redundant LAN routers, as well as load balancing of traffic. VRRP is transparent to host end stations. All configurations are done at the Avaya Multiservice Switch, no additional end station configurations are required. See "Configuring VRRP" in Chapter 12, "Configuring IP Routing," for more details.

Hunt Groups

Ports within a hunt group by default are redundant to one another since hunt groups perform load balancing among the ports. The hunt group load-shares the traffic between two switches allowing the bandwidth to be multiplied. The use of hunt groups also increases reliability since the links behave as hot standby links to one another. If the traffic is shared over multiple links, and one of the links is "lost", the traffic will be automatically redistributed over the other links and the communications will continue without interruption. See "Using Hunt Groups to Aggregate Bandwidth" in Chapter 6, "Using VLANs, Hunt Groups, and VTP Snooping," for more details.

Redundant Power Supplies

Only two power supplies are required to support a fully-loaded Avaya Multiservice Switch. A third power supply can be installed to provide backup should one of the other two fail. The power subsystem provides N+1 power supply redundancy.

Supervisor Failover conditions

The Active supervisor module fails over to the standby supervisor if one of the following events occur:

- Active Supervisor module removal
- Active Supervisor Software reset
- Active Supervisor Module Reset Push button
- Active Supervisor loses power
- Active Supervisor system crash either hardware or software

Chapter Contents

The following information and procedures are provided in this chapter:

- Configuring Redundant Hardware
- Installing and Enabling Redundant Hardware
- Replacing the Primary Controller
- Replacing an Element
- Configuring Supervisor Module Redundant Ethernet Console IP Addresses
- Synchronizing the Active and Standby CPUs

Configuring Redundant Hardware

The Avaya Multiservice switches provide a redundant backplane to ensure that should a Supervisor module, Switch Controller, or fabric element fail, it will continue to forward data properly. The redundant components are available as separate options.

Redundant Switch Controllers and Elements

Controllers and Elements are located at the rear of the switches. When the redundant controller and element are installed, the backplane consists of the following parts (Figure 10-1or Figure 10-2):

- One Primary Controller
- One Redundant Controller
- Six Active elements
- One Redundant element

Figure 10-1. Layout of the Redundant Backplane for P580 Chassis

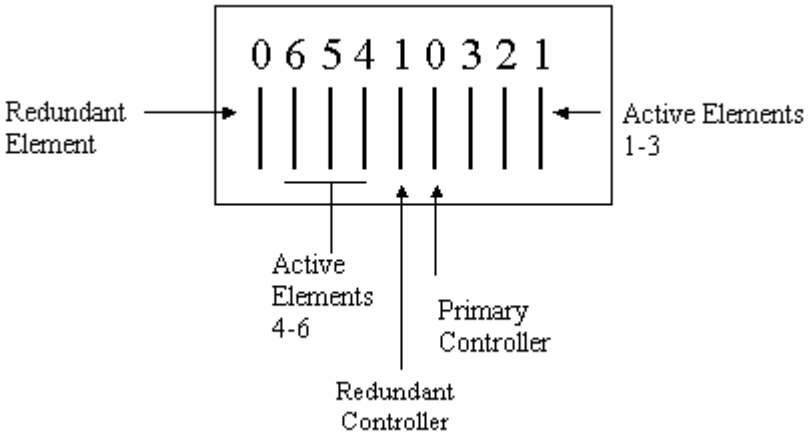
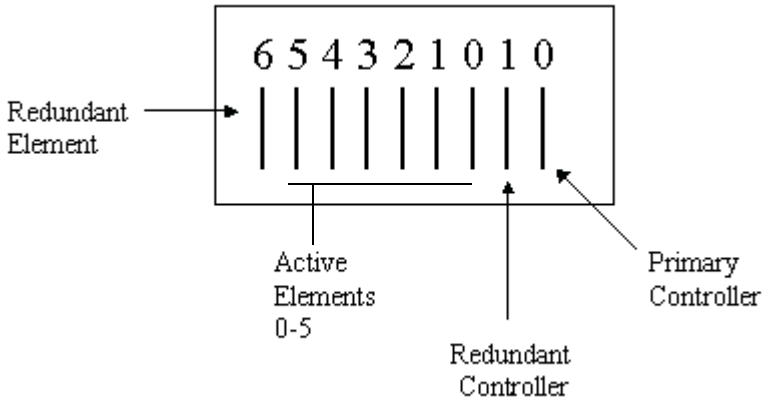


Figure 10-2. Layout of the Redundant Backplane for P882 Chassis



The Avaya Multiservice switches are equipped with diagnostics to:

- Monitor the status of the Primary switch controller.
- Monitor the status of a Active element.
- Monitoring of Health Check status messages sent between the Active and Standby supervisors.

If the Active supervisor fails, the Standby supervisor assumes supervisor operation. Information about a failed Supervisor (CPU) displays in the CPU Redundancy Status Web page.

If the Primary controller fails, the redundant controller takes over switch controller operation. When an element fails, diagnostics run automatically to test the hardware. Information about a failed controller or element displays in the Switch Fabric Status Web page. The redundant element takes over the function of the failed element until it is replaced.

*** Note:** In the event that the Primary switch controller fails, the switch will reset and record an entry of the event in the shutdown log.

If the Active element fails, the Active and Standby Supervisor modules can no longer communicate, or be synchronized. No health messages will be passed. However, the Standby Supervisor module will become Active if the current Active Supervisor module fails.

Installing and Enabling Redundant Hardware

The standard switch configuration does not include a redundant controller or element. The correct process is to:

- Install the redundant hardware. See "Installing Redundant Hardware."
- Enable the redundant hardware by using either the Web Agent or the CLI. See "Enabling Redundant Hardware."

After the redundant hardware is enabled, the applicable redundant component takes over if either a controller or any element fails.

Installing Redundant Hardware

To install redundant hardware:

1. Turn off all of the switch power supplies.
2. Replace the Primary switch controller with the new switch controller. This ensures that the switch checks the status of the new switch controller.
3. Insert the previously Primary switch controller into the Redundant Controller Slot 1 (see Figure 10-1 for a P580 or Figure 10-2 for a P882).

*** Note:** Controllers and elements are keyed to ensure that a controller is not inserted into an element slot.

4. Install the redundant element into the redundant element slot 0 for the P580, (see Figure 10-1), or slot 6 for the P882, (see Figure 10-2).
5. Turn on the switch power supplies and log in.
6. Enable the Redundant Hardware as described in the next section.

Enabling Redundant Hardware

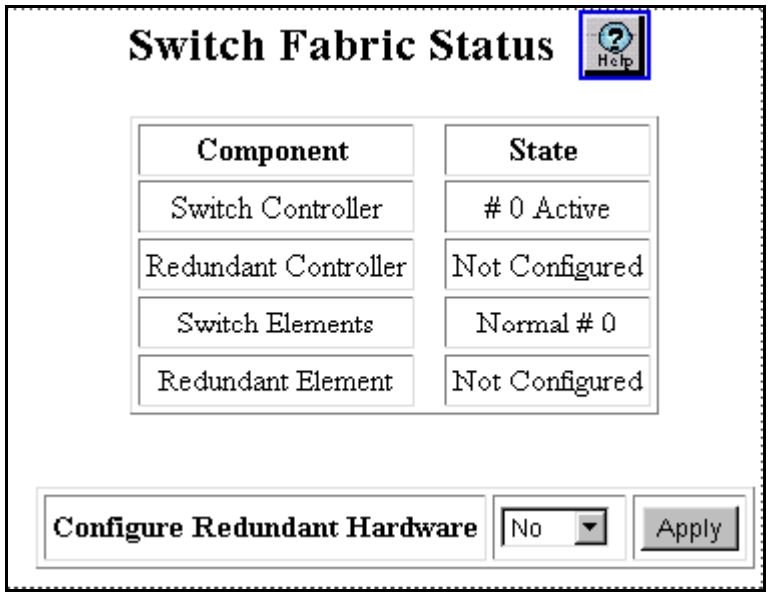
Web Agent Procedure

To enable redundant hardware using the Web Agent:

1. In the navigation pane, expand the **System > Configuration** folders, and then click **Switch Fabric**.

   The Switch Fabric Status Web page is displayed in the content pane. See Figure 10-3.

Figure 10-3. Switch Fabric Status Web Page



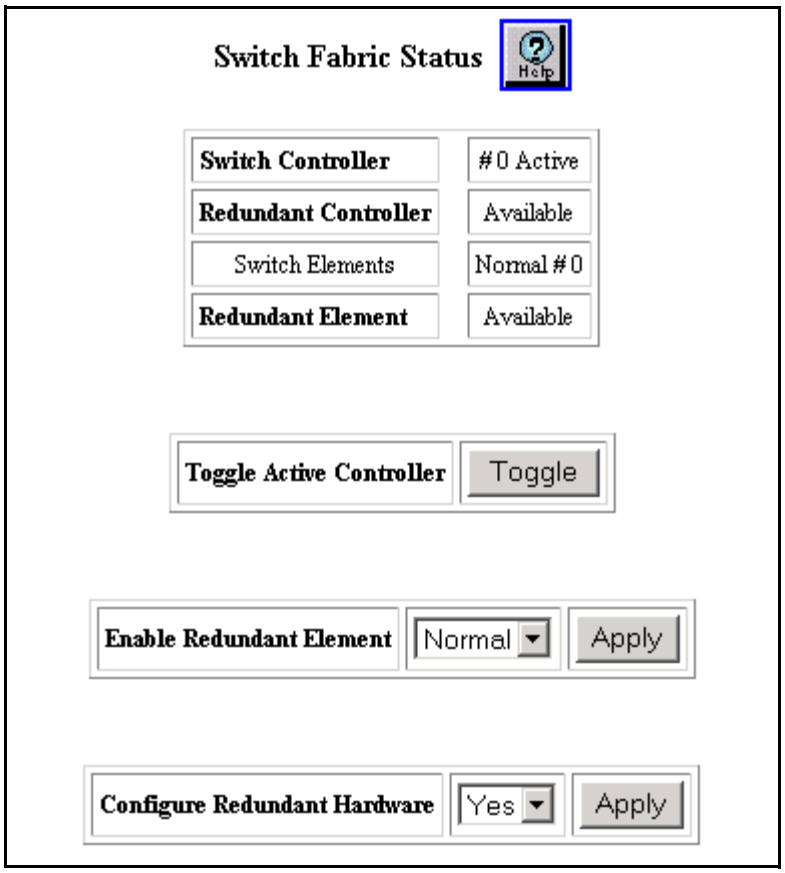
2. Select **Yes** from the **Configure redundant Hardware field** pull-down menu.

3. Click **APPLY**. The Switch Fabric Status should now show the redundant controller and element **Available**. (Figure 10-4)

4. Save the configuration.

*** Note:** If you disable and then reenable redundant hardware after the Primary controller fails, the state of the failed controller changes from **Failed** to **Available**.

Figure 10-4. Switch Fabric Status Web Page



CLI Command

To enable the redundant hardware by using the CLI, enter the following command from **Enable/Configure** mode:

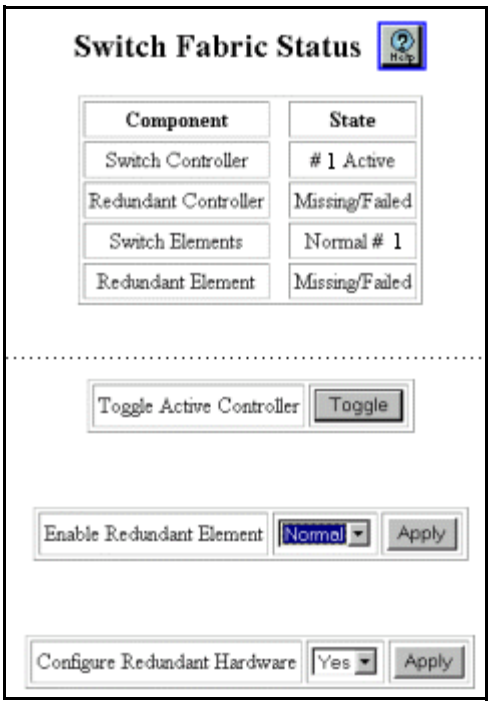
`(configure)#` **set fabric configure-redundant-hardware enable**

Replacing the Primary Controller

If a Primary controller fails, notification is provided by:

- An event log message
- An error message that displays in the **Switch controller** field.

Figure 10-5. Switch Fabric Status Web Page



If a redundant controller is installed in the switch and enabled, the switch resets and the redundant controller automatically takes over the operation of the failed Primary controller.

To replace the failed Primary controller:

1. Disable redundant hardware.
2. Copy the running config to the startup config.
3. Synchronize if you have redundant supervisor modules.
4. Power off the switch power supplies.

WARNING: To avoid bodily harm and equipment damage, you must power off the switch before you remove switch controllers or elements.

5. Replace the failed controller.
6. Power on the switch power supplies and login.
7. Enable the redundant hardware (see "Enabling Redundant Hardware" earlier in this chapter).
8. In the navigation pane, click **Switch Fabric** again (**System > Configuration** folders). The Switch Fabric Status Web page redisplays with the **Switch Controller** field displaying:

   ```
   # 1 Active
   ```

Replacing an Element

If an element fails, you are notified by:

- An event log message
- An error message that displays in the **Switch Elements** field

To replace a failed element:

1. Disable redundant hardware.
2. Copy the running config to the startup config.
3. Synchronize if you have redundant supervisor modules.
4. Power off the switch power supplies.

WARNING: To avoid bodily harm and equipment damage, you must power off the switch before you remove switch controllers or elements.

5. Power on the switch power supplies and login.
6. Enable the redundant hardware (see "Enabling Redundant Hardware" earlier in this chapter).
7. Replace the failed element.
8. Power on the switch's power supplies and login.

Web Agent Procedure

To enable the replacement element using the Web Agent:

1. In the navigation pane, expand the **System > Configuration** folders, and then click **Switch Fabric**. The Switch Fabric Status Web page is displayed in the content pane. See Figure 10-5.
2. Select **Yes** from the **Configure Redundant Hardware** field pull-down menu.
3. Click **APPLY** to enable the redundant element.

   The Switch Elements field displays:

   ```
   Normal # 0
   ```

   This indicates that the redundant element is now enabled.
4. Save the configuration changes

CLI Command

To enable an element using the CLI, enter the following command from Configure mode:

> **set fabric configure-redundant-hardware enable**

Configuring Supervisor Module Redundant Ethernet Console IP Addresses

Web Agent Procedure

If you need to communicate with both the Active and Standby Supervisors, configure an alternate Ethernet Console IP address for both. Configure an alternate IP address for both slot 1 and 2 CPU's, by using the Web Agent or the CLI.

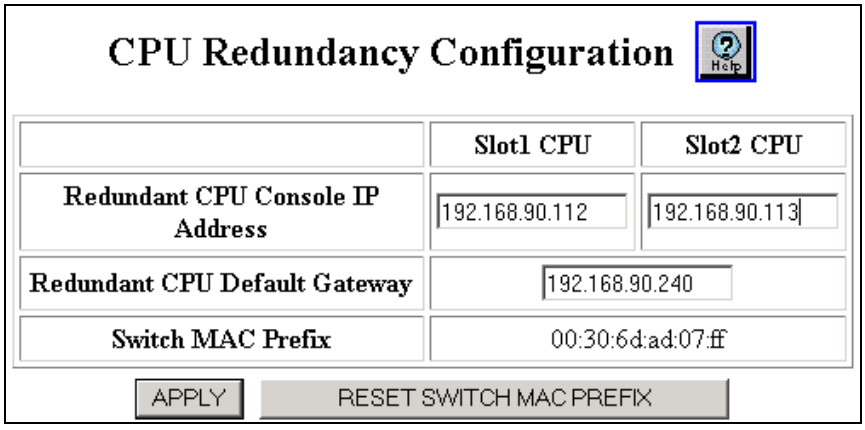
*** Note:** Do not Login to a Standby Supervisor module to manage the switch. The Standby Supervisor does not have accurate switch status information until it becomes the Active Supervisor.

To identify the Active or Standby supervisor, look at the LED display (marquee), or open the **System > Configuration > CPU Redundancy > Status/Statistics** folders from the Web Agent, or use the **show CPU status** CLI command.

To configure a redundant console IP address by using the Web Agent:

1. In the navigation pane, expand the **System > Configuration > CPU Redundancy** folders, and then click **Configuration**. The CPU Redundancy Configuration Web page is displayed in the content pane. See Figure 10-6.

Figure 10-6. CPU Redundancy Configuration Web Page



2. Enter the internal IP addresses of the **slot 1 and slot 2 CPUs** in the **Redundant CPU Console IP Address** field. The addresses cannot be the same as the console IP or each other. The Gateway must be on the same subnet as the IP addresses. This sets the IP addresses for the Ethernet Consoles in slot 1 and slot 2 CPU's. The default values are 0.0.0.0.

3. Enter the redundant CPU default gateway IP address in the **Redundant CPU Default Gateway** field. This is the default gateway that the standby Ethernet console IP interface uses. It does not have to match the Active CPU's default gateway. The default value is 0.0.0.0.

4. The Switch **MAC Prefix** field displays the configured MAC prefix used to build the local MAC address for each routing interface. The first 36 bits of the switch MAC prefix make up the first 36 bits of this local MAC address. If the startup.txt does not set the switch MAC address, the supervisor's MAC address is used. Resetting this will cause it to be restored to the supervisor's MAC address. Only the Active supervisor uses this value. The default value is the manufacturer's supervisor MAC address.

5. Click **APPLY** to save your changes,

6. **RESET SWITCH MAC PREFIX** to reset the MAC prefix on your switch.

CLI Command

To configure redundant CPU IP addresses using the CLI, use the following commands from **Enable/Configure** mode:

- `(configure)#` **cpu_redundancy console {slot1 | slot2}** <***ip address***>
- `(configure)#` **cpu_redundancy default-gateway** <***ip address***>

*** Note:** When the redundant CPU console IP address and default gateway are assigned and addresses are removed by using the commands:

> no cpu_redundancy slot1
> no cpu_redundancy slot2
> no cpu_redundancy gateway

The router is no longer reachable and cannot be pinged or used by the Web Agent from redundant addresses. The initial console IP address is still usable.

Synchronizing the Active and Standby CPUs

Overview

After installing the active and standby supervisor modules, you must synchronize them. You can use the Web Agent or CLI to perform the synchronization.

*** Note:** If the Active Supervisor module and the Standby (redundant) Supervisor module are synchronized while the active supervisor module is being heavily used, TFTP transfers from the active supervisor module to the Standby Supervisor module may time out or fail.

*** Note:** After you synchronize supervisor modules, the standby supervisor can be accessed only with the default passwords. The synchronization does not synchronize passwords.

*** Note:** P580 and P882 supervisor modules must be in the same Fabric mode (1 or 2) to communicate and to synchronize.

During synchronization, the boot code, software images (app 1 and app 2), and switch configuration are synchronized. If the boot code or software images on the standby supervisor differ from those on the active supervisor, the active supervisor overwrites the boot code or software images on the standby supervisor. (The boot code and software images on the active supervisor overwrite those on the standby regardless of which supervisor has earlier or later software.) To synchronize the two supervisor modules, the active supervisor performs TFTP uploads to the standby supervisor. This process takes place over an internal subnet and VLAN.

Once the software images and boot code on the standby and active supervisors are synchronized, the active supervisor then performs a TFTP upload of the startup.txt and the.int files.

Once this step is done, and the standby supervisor is reset, and the synchronization is complete.

The section contains the following information and procedures:

- Supervisor Module Redundancy Statistics
- Active and Standby Roles
- Synchronization Procedure

Supervisor Module Redundancy Statistics

The active supervisor periodically sends a health report message to the standby (redundant) supervisor. If the standby supervisor is enabled, it responds with a health report reply.

Health report statistics are displayed on the CPU Redundancy Status Web page. To view this Web page, expand the **System** > **Configuration** > **CPU Redundancy** folders, and then click **Status/Statistics**. The following statistics are displayed:

- **Health Reports Sent** - Displays the number of health reports sent.
- **Health Reports Received** - Displays the number of health reports received.
- **Health Reports Timeouts** - Displays the number of health reports timeouts.
- **Health Reports Missed** - Displays the number of health reports missed.

Click **REFRESH** to reset all CPU redundancy statistics.

Active and Standby Roles

Table 10-1 shows possible supervisor configurations and, for each configuration, which supervisor is active and which is standby.

Table 10-1. Supervisor Module Configurations

Configuration	Active and Standby Role
One supervisor in slot 1	The slot 1 supervisor module is active. If a supervisor module is inserted into slot 2, the slot 1 supervisor module remains active and sends health reports to the standby supervisor in slot 2.
	1 of 2

Table 10-1. Supervisor Module Configurations

Configuration	Active and Standby Role
One Supervisor module in slot 2	The slot 2 supervisor module is active. If a supervisor module is inserted into slot 1, the slot 2 supervisor module remains active and sends health reports to the standby supervisor in slot 1. If the slot 2 supervisor is removed, reset, or fails, or if the switch is reset, the slot 1 supervisor becomes the active supervisor, and the slot 2 supervisor becomes the standby. The slot 1 supervisor then sends health reports to the slot 2 standby supervisor.
Supervisor modules in slot 1 and slot 2	The slot 1 supervisor module is active and sends health reports to the standby supervisor in slot 2. If the slot 1 supervisor module is removed, reset, or fails, the slot 2 supervisor becomes the active supervisor. If you then replace the slot 1 supervisor, it remains the standby supervisor and receives health reports from the active supervisor in slot 2. If the slot 2 supervisor is removed, reset, or fails, or if the switch is reset, the slot 1 supervisor becomes the active supervisor, and the slot 2 supervisor becomes the standby. The slot 1 supervisor then sends health reports to the slot 2 standby supervisor. **Note:** If the switch is reset, slot 1 supervisor is the active supervisor, unless it has failed.
	2 of 2

*** Note:** If the standby supervisor becomes active, it assumes the IP and MAC addresses of the previously active supervisor.

*** Note:** When a standby supervisor is installed in the switch there is a 45-second wait period before the standby supervisor is fully initialized.



CAUTION: Do not remove the standby supervisor module during initialization.

Synchronization Procedure

Web Agent Procedure

*** Note:** You must synchronize the active and standby supervisor modules for the standby (redundant) supervisor to use the same configuration as the active supervisor. Save the running configuration to the startup configuration so that if the active supervisor fails, the current configuration information is not lost. Synchronizing the two supervisors copies the startup configuration file to the standby supervisor.

*** Note:** The P580 and P882 Supervisor modules must be in the same Fabric mode to synchronize.

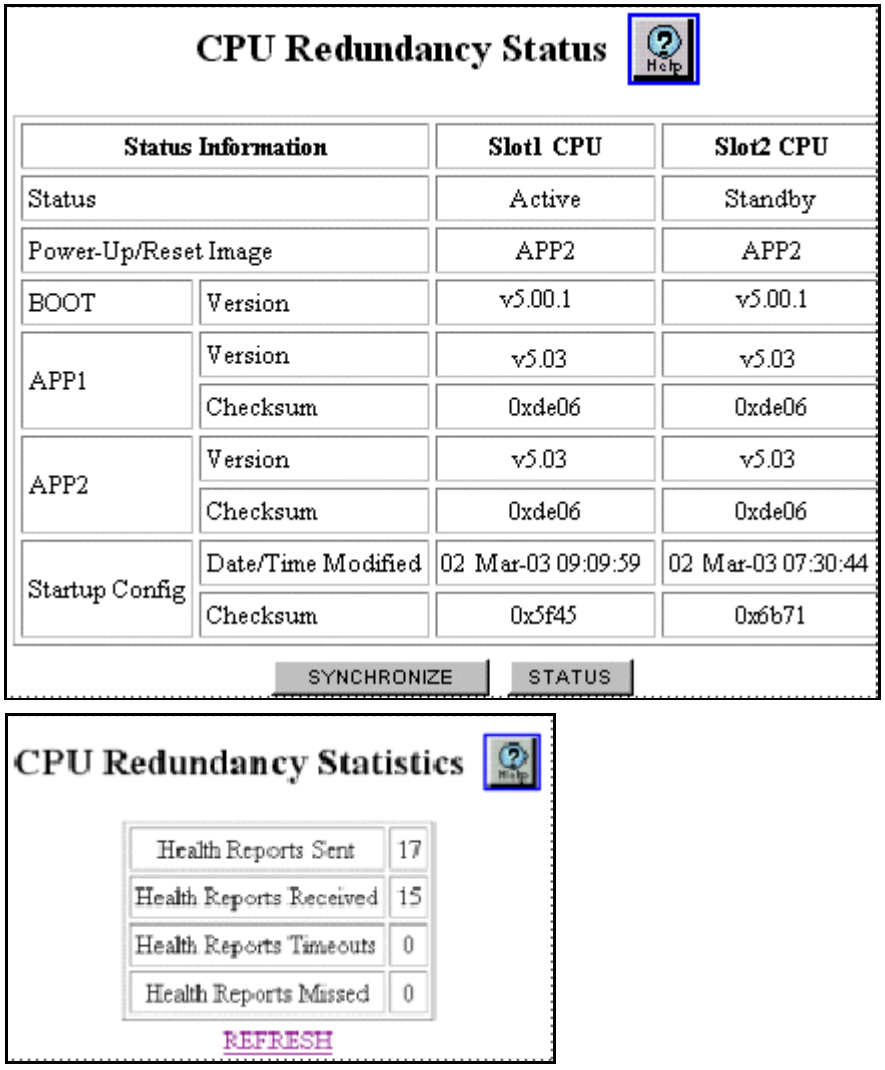
To synchronize the active and standby supervisor modules:

1. Copy the running configuration to the startup configuration. For information on how to copy the running configuration to the startup configuration, see "Copying Configuration Files" in Chapter 2, "Configuring Redundancy Options."

2. In the navigation pane, expand the **System** > **Configuration** > **CPU Redundancy** folders, and then click **Status/Statistics.**

   The CPU Redundancy Status Web page is displayed in the content pane. See Figure 10-7.

Figure 10-7. CPU Redundancy Status Web Page



3. Monitor the CPU redundancy statistics to ensure that the health messages between the supervisor modules are being sent and received. Click **REFRESH** to ensure that the counters increment.

See Table 10-2 for an explanation of the fields on the CPU Redundancy Status Web page.

Table 10-2. CPU Redundancy Status Web Page Fields

Parameter	Definition
Status	Displays the functional status of the CPU modules.
Power-Up/Reset Image	Displays the image the CPU will use upon startup or reset.
	1 of 2

Table 10-2. CPU Redundancy Status Web Page Fields

Parameter	Definition
APP1	• **Version** - Displays the version of the image in application 1 (APP1) for each supervisor module. • **Checksum** - Displays the checksum for the image running in APP1.
APP2	• **Version** - Displays the version of the image in application 2 (APP2) for each supervisor module. • **Checksum** - Displays the checksum for the image in APP2.
Startup Config	• **Date/Time Modified** - Displays the date and time that the configuration (startup.txt) was modified. • **Checksum** - Displays the checksum for the image in the Startup Config.
	2 of 2

4. Click **SYNCHRONIZE** to start the synchronization process.

If both supervisors are running v6.0 or later application software, the switch performs the following checks after you synchronize the active and standby supervisors:

- Version of application software.
- Startup image (app 1, app 2, cardapp 1, or cardapp 2)

If any of these settings changed during the synchronization, the standby supervisor module automatically resets so that it is failover-ready. In earlier versions of software, you had to manually reset the standby supervisor, if these settings changed.

***Important:** If you are updating the boot code, you must manually reset the standby supervisor after synchronizing it with the active supervisor.

*** Note:** For this automatic reset feature to work, the active and standby supervisors must *both* be running v6.0 or later application software. The first time that you synchronize the supervisor modules after upgrading the active supervisor from v5.*x* application software to v6.*x*, you must manually reset the standby supervisor.

CLI Command

To synchronize the active and standby supervisor modules, use the following CLI command:

```
(configure)# cpu_redundancy synchronize
```

11 Configuring DNS Client

Overview

This chapter provides an overview of the Domain Naming System (DNS) Client support and procedures for configuring DNS client support on the Avaya Multiservice switch:

- DNS Client Support
- DNS Client on the Avaya Multiservice Switch
- Configuring the DNS Client

DNS Client Support

Overview

DNS is a distributed database of name servers which supply name-to-address translations for DNS clients. The DNS servers collectively direct DNS Clients to the DNS Server responsible for a Name to Address mapping; and that Server provides the specific mapping being requested.

There are three types of DNS servers:

- The root name server — Located at the top of the DNS database tree. It contains pointers to the master name servers for each of the top-level domains. The root name servers handle the domains such as .COM, .EDU, .GOV, etc.
- The master name server — Located in the middle of the DNS database tree. It contains pointers to the individual name servers for each of the subdomains within its top-level domain.
- The individual name server — Located at the bottom of the DNS database tree. It contains detailed address information for the local hosts in the subdomain.

Recursive vs. Iterative

The goal of DNS is to resolve a fully qualified domain name (FQDN) to an IP address. This work can either be done by the DNS server or the DNS client. These approaches are referred to as Recursive or Iterative, respectively.

In the Recursive approach, a client sends a query to the server. Assuming recursion is enabled, the server then looks for resolution first locally in its own database, then in its local cache, finally by going through the DNS tree until it finds a server that can give an authoritative answer to the query. In this model, the client is referred to as a Stub Resolver. Typically, Stub Resolvers are implemented on devices with limited resources such as embedded systems or Personal Computers.

The Avaya Multiservice switch is a Stub Resolver, i.e. it only supports a Recursive lookup not Iterative.

There are four most common answers a DNS server can provide are:

- **Authoritative -** a positive answer returned to the client with the Authoritative Answer (AA) bit set in the response.
- **Positive** - an answer that contains the resource record (RR) or list of RRs that match the query.
- **Referral** - an answer that contains a list of alternate servers the client can use to resolve the name. This type of answer is given if Recursion is not supported.
- **Negative** - this answer indicates that an Authoritative server reported that the name (or record type) does not exist in the DNS name space. The server can be the preferred server or, if recursion is enabled, that another server responded back to the preferred server with the authority bit set.

In the Iterative approach, the client sends a query to the server. If recursion is disabled, and the server cannot answer the query, the server will responds with a Referral answer. The client will then use that information to query another DNS server. This process will continue until a server responds with an Authoritative response.

Sample Recursive DNS Query

Typically, a network will have a local DNS server which provides translations for devices in the local network. That server will also "know" the IP address of the Root servers. When a client attempts to communicate with a device by its name, the IP portion of the protocol stack will recognize that a name (not an address) has been specified. It will then contact the local DNS server and request the name be translated into an IP address. For example: if a user types **`ping www.avaya.com`**, that name will be sent to the local DNS server. If the DNS server has that information, it will respond with the appropriate IP address.

If the DNS server does not have that address translation, it will contact the root server using the IP address for that root server. In this example, it would use the IP address for the **.COM** root server.

Root servers do not have specific translations for names, but they do have the database of master name servers. The root server would the avaya.com master name server using the avaya.com IP address.

The root server would then send a request to the master name server requesting the translation. In small subdomains, a single DNS server may be adequate for servicing the subdomain and be able to translate all DNS requests. So the master name server may act as the individual name server by providing the translations.

In large subdomains, a single DNS server may be inadequate to handle all of the address translations. In these cases, the master name server may point to other individual name servers. In this case, we are pinging the Web server **www.avaya.com**. Assume for this example that Avaya has a single DNS server for the entire subdomain of avaya.com. The avaya.com DNS server would use the IP address of the Avaya Web server.

The client would then insert the IP address into the IP packets going to the destination. In this example, it would put the IP address in the HTTP packet going to the Avaya Web site.

The benefit of using DNS is that you need only know the name of the server instead of the IP address for which you are trying to communicate. Also, if the IP address of the server changes, you need only update the DNS database.

Many vendors provide DNS servers. Consult the DNS Server vendor's documentation for information on configuring the DNS server. Virtually every IP protocol stack includes DNS client capabilities. The Avaya multiservice switch is a DNS client only.

DNS Client on the Avaya Multiservice Switch

The Avaya Multiservice switch has DNS client capabilities that work in conjunction with the Command Line Interface (CLI). This feature is useful for testing connectivity (ping and traceroute) as well as copying files from TFTP servers.

It is not meant for use with the normal operation of the Avaya switch. Therefore, you should not use a DNS name to specify the location of the Avaya Multiservice switch's RADIUS server, SNTP Time server, etc. These parameters must be specified as IP addresses. However, DNS does work for the Avaya Multiservice Switch's Help File HTTP server.

In order for the DNS client functionality to work, it must be enabled on the switch and at least one name server IP address must be specified.

The following DNS Client features are supported:

- The DNS Client on the Avaya Multiservice switch is Stub Resolver (i.e. it does not support Iteration). This means that referral answers will be discarded. A Recursive DNS server must be used if you intend on using FQDN for which the configured DNS servers do not have resource records.
- Up to six DNS name servers can be configured.
- Up to six default DNS Domain Suffixes can be configured.
- When DNS is enabled, DNS names can be used in place of IP addresses for `ping`, `telnet`, `connect, traceroute and copy tftp` CLI commands.
- The DNS client name can consist of up to 255 characters.
- A single label within a DNS client name can consist of 63 characters.

Once enabled, you have the option of specifying a DNS suffix to add to unqualified host names. If so, you need only type the host name, in place of the IP address, without the suffix when entering commands. The suffix will be appended to host names that have no suffix.

For example: You can specify the *avaya.com* suffix. So when you want to ping **hostA** in the avaya.com domain, you need only type the host name without the suffix.

Avaya(configure)# **ping hostA**

instead of:

Avaya(configure)# **ping hostA.avaya.com**

Order of Operations for DNS on the Switch

You can configure up to six DNS servers and up to six DNS suffixes. When you use a name instead of an IP address in a command, the Avaya Multiservice switch will first check to see if the name is a fully qualified domain name (ex: hostA.avaya.com). If it is, this will be sent to the first DNS server in the list of servers.

If it is not a fully qualified domain name (ex: just hostA) and a suffix(es) has been specified, the first suffix will be appended to the name; and then sent to the DNS server. If no suffix(es) is configured, it will be sent as is for resolution.

If the server cannot resolve the name, the Avaya Multiservice switch will use the same method with the second configured suffix. Once it has exhausted the configured DNS suffixes, it will attempt the query with the second DNS server.

This process will continue until the name is either resolved, there are no more DNS servers in the list, or the DNS server returns an error. Table 11-1 lists the error messages and their meanings that you may encounter when using the DNS client.

Table 11-1. DNS Error Codes

Error Message	Explanation
Bad ARGS	indicates that DNS is Disabled but the user entered a host name.
Name Too Long	the name sent is too long. RFC 1034 limits DNS names to 255 characters.
Bad Name	indicates that the name was in some way invalid
Label Too Long	indicates that the label of a DNS name was too long. RFC 1034 limits labels to 63 octets.
Time-out	indicates that the DNS query has expired. This implies that the query could not be answered at the present time.
Server Failure	indicates that the DNS server is unable to answer due to a failure on the DNS server itself.
Non-Existent Name	indicates that the DNS server authoritatively claims that the DNS name does not exist.
NIY	Indicates the DNS server does not support the requested service. You may have specified the correct DNS server but the server has not yet been configured to respond to DNS queries.
	1 of 2

Table 11-1. DNS Error Codes

Error Message	Explanation
Refused	indicates that the DNS server refuses to answer the query for administrative reasons. Possibly due to security implementations on the DNS server.
No RRs	indicates that the DNS server authoritatively claims that there are no RRs that match the specified name, type and/or class.
No Recursion	indicates that one of the configured DNS servers does not support Recursion.
Irrelevant	indicates that a response message was received that does not match the query sent.
	2 of 2

Configuring the DNS Client

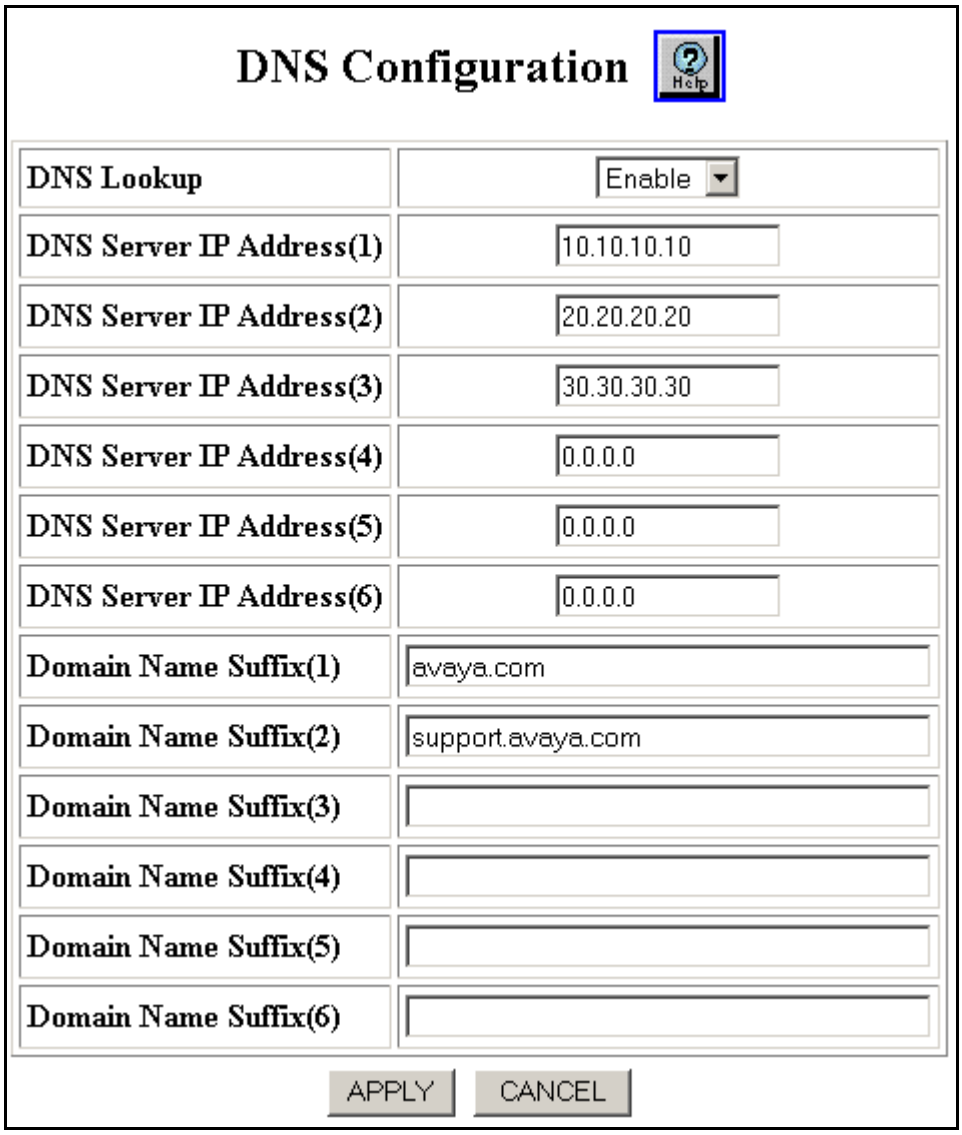
You can configure DNS client from the Web Agent or the CLI.

Web Agent Procedure

To configure a DNS client by using the Web Agent:

1. In the navigation pane, expand the **System > Configuration** folders, and then click **DNS**. The DNS Configuration Web page is displayed in the content pane. See Figure 11-1.

Figure 11-1. DNS Configuration Web Page



2. Select **Enable** from the **DNS Lookup** field pull-down menu to enable DNS on the switch.

   ***Note:** When DNS is enabled and configured, you can enter DNS names in place of IP addresses for the following CLI commands. Again, DNS is for use with the CLI only for testing connectivity and copying files from TFTP servers, not the normal operation of the Avaya Multiservice switch.

   - ping
   - telnet
   - connect
   - trace (traceroute)
   - copy tftp

3. Enter the DNS server IP address for each DNS server in the **DNS Server IP Address** fields (fields 1-6).
4. Enter the domain name suffix for each DNS server in the **Domain Name Suffix** fields (fields 1-6).
5. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Commands

The following commands configure DNS client:

- To enable DNS client, `(configure)#` **ip domain-lookup**
- To add a DNS server address, `(configure)#` **ip name-server** <***a.b.c.d***>
- To add a DNS suffix, `(configure)#` **ip domain-list** <***name***> and **ip domain-name** <***name***>
- To display the DNS Client configuration, `(configure)#` **show hosts**

CLI Example

The following CLI example enables DNS, configures three DNS servers, two suffixes, and then displays the DNS configuration:

```
Avaya(configure)# ip domain-lookup
DNS domain name lookup is ENABLED.

Avaya(configure)# ip name-server 10.10.10.10
Added name server '10.10.10.10' to the name server
list.
```

```
Avaya(configure)# ip name-server 20.20.20.20
Added name server '20.20.20.20' to the name server
list.

Avaya(configure)# ip name-server 30.30.30.30
Added name server '30.30.30.30' to the name server
list.

Avaya(configure)# ip domain-list avaya.com
Added domain name 'avaya.com' to the domain name
list.

Avaya(configure)# ip domain-list support.com
Added domain name 'support.avaya.com' to the domain
name list.

Avaya(configure)# show hosts
DNS domain name lookup is ENABLED.

Configured DNS name servers:
        10.10.10.10
        20.20.20.20
        30.30.30.30
Configured DNS default domain names:
        avaya.com
       support.avaya.com
```

12 Configuring IP Routing

Overview

The Avaya Multiservice Switch supports the full suite of IP protocols for unicast and multicast packet routing and control. It is compatible with the installed base of IP routers and does not require changes to software in hosts or other routers. For Unicast, the Switch supports OSPF, RIP, and RIP2 interior gateway protocols. For Multicast, it implements IGMP, and DVMRP.

The following information and procedures provided in this chapter pertain to layer 3 module configurations only:

- Routing Function
- Requirements for IP Routing
- Routing Configuration Quickstart
- Displaying Existing IP Interfaces
- Creating and Assigning IP Interfaces to the VLAN
- Enabling IP Routing Global Parameters
- IP Multinetting
- Configuring Short-Lived IP Protocol Filters
- Creating IP Static Routes
- Creating a Static Route to a Null Interface
- Creating IP Static ARP Entries
- Creating a BOOTP/DHCP Server Entry
- Configuring Route Redistribution
- IP Multicast
- Monitoring Switch Performance Using IP Statistics
- IP Multicast Statistics

- Configuring VRRP
- Configuring IRDP
- Configuring LDAP
- Configuring a Static Route for the PPP Console
- Configuring the IP Interface for the PPP Console

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.

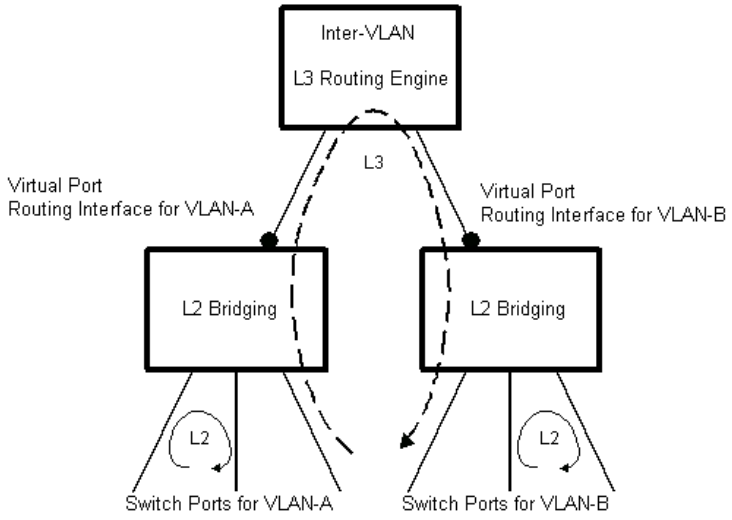
Routing Function

The Routing function is logically independent from and sits “on top” of the Layer 2 VLANs. The function of the Inter-VLAN router is to route traffic between VLANs (subnets). The router has an “interface”, or “virtual port”, for each VLAN that it is routing traffic.

This interface is configured similarly to a physical interface, but is not associated with any particular physical interface.(Figure 12-1)

For more detailed information on the switch routing operations, see “Routing with Layer 2 and Layer 3 Modules”, in Chapter 1, “Introduction.”

Figure 12-1. Intra and Inter VLAN Bridging and Routing



Requirements for IP Routing

Hardware Requirements

To configure your switch as an IP router, you must configure your switch with the following hardware:

- Layer 3 supervisor module

*** Note:** Licensed layer 3 80-series media modules are recommended for best routing performance.

Software Requirements

To configure your switch as an IP router, with 80-series modules (Supervisor and Media modules) you must configure your switch with a minimum of version 5.0 and above.

Minimum Configuration Requirements

Your switch must be configured as follows to successfully create IP routing:

- VLANs (subnets) have been created to address current or future network configurations, or if you plan to expand your current network.
- Assign an IP address, and subnet mask to each configured IP interface.
- IP unicast forwarding must be globally enabled.
- At least one routing protocol (RIPv1,RIP v2, OSPFv2) must be enabled, if communication between routers is necessary.

Routing Configuration Quickstart

This section provides an overview of the how to configure the Avaya Multiservice switch as a IP router.

To configure the switch as a IP router:

1. Ensure that physical port parameters and cabling are correctly configured before creating VLANs. See Chapter 8, "Configuring Ports".
2. Create a VLAN (each VLAN is an IP subnet). See "Creating and Implementing VLANs," in Chapter 6, "Using VLANs, Hunt Groups, and VTP Snooping."
3. Assign switch ports to the VLAN. See "Configuring Switch Ports Settings", in Chapter 8,"Configuring Ports."
4. Configure an IP interface, IP address and subnet mask, for each VLAN and associated subnet.
5. Globally enable IP Forwarding (Routing).
6. Enable the routing protocol on an interface.
7. Specify how the router will communicate with other routers, by configuring RIPv1, RIPv2 or OSPF v2 parameters.

*** Note:** The IP interfaces should be set to the ADMIN state **Down**.

8. Set the IP interface Admin State to **Up**.

Displaying Existing IP Interfaces

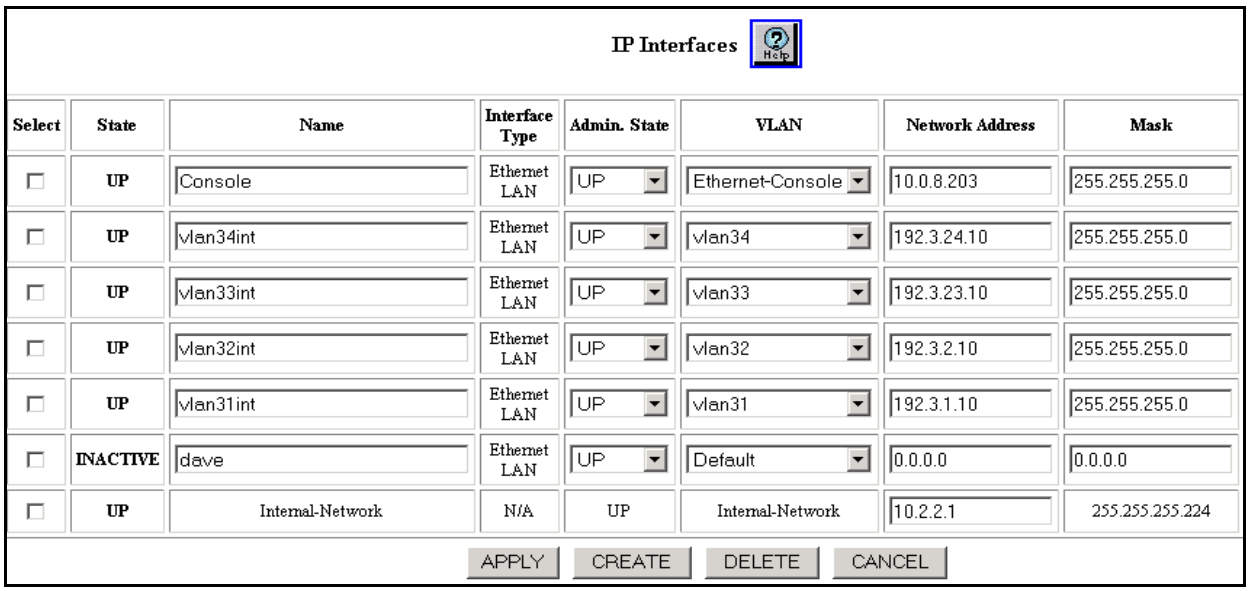
When you create an IP interface, you activate a location in the switch that communicates between IP and the embedded software of the switch. You can create an IP interface using either the Web Agent or the CLI.

Web Agent Procedure

To display, modify, delete a previously created IP interfaces or to create a new IP interface using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **Interfaces**. The IP Interfaces Web page is displayed in the content pane. See Figure 12-2.

Figure 12-2. IP Interfaces Web Page



2. To modify an IP interface, first select the specific interface, then modify the parameters, and then click **APPLY** or **CANCEL** to restore previous settings.

*** Note:** If you modify the ARP Timeout value, the new value does **not** take affect when you click **APPLY**. To modify the ARP Timeout setting, first change the interface **Admin State** to **Down,** then change the ARP Timeout value, and finally change the interface **Admin State** back to **Up**:

a. Select the **Select** checkbox for the interface.

b. Change the **Admin State** to **Down.**

c. Click **APPLY**.

d. Again select the interface, enter the new value of the **ARP Timeout** period.

e. Change the **Admin State** to **UP**.

f. Click **APPLY**.

3. To delete an IP interface first click on **Select,** to select the specific interface then click on **DELETE**.

4. See Table 12-1 for an explanation of the IP Interfaces Web page parameters.

Creating and Assigning IP Interfaces to the VLAN

Web Agent Procedure

After you have set up a VLAN, you must create the IP interface that your VLAN and your subnet uses to communicate with each other. While creating the IP interface, you must assign it to the VLAN. The last step is to enable IP forwarding (Routing) global parameters.

To create a new IP interface and assign it to a VLAN:

1. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **Interfaces**. The IP Interfaces Web page is displayed in the content pane. See Figure 12-2.
2. Click **CREATE**. The Add IP Interface Web page is displayed. See Figure 12-3.

Figure 12-3. Add IP Interface Web page

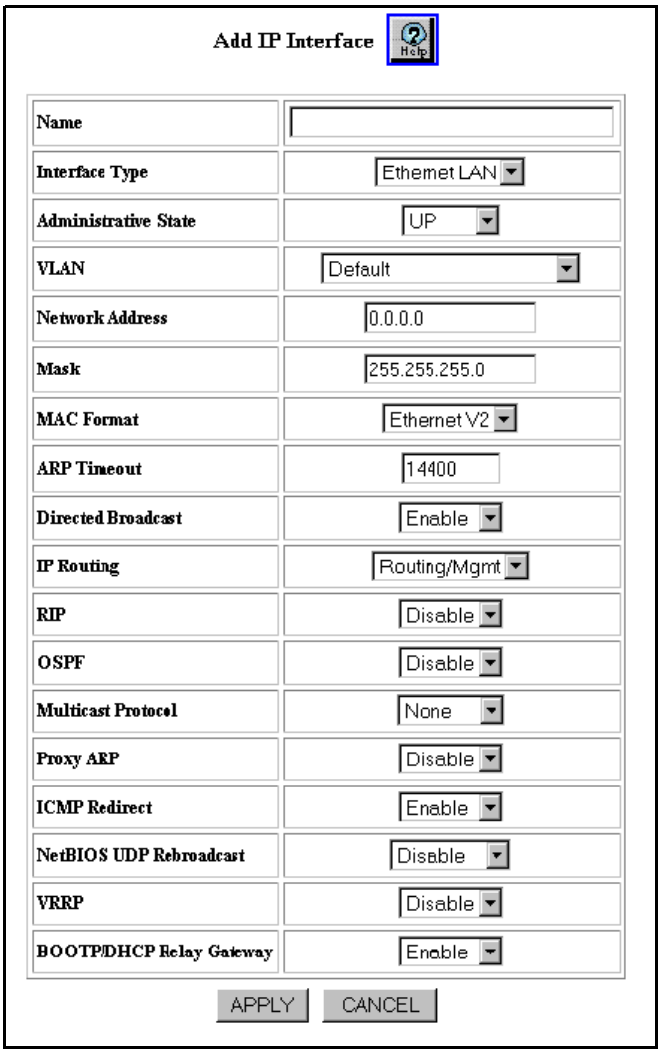


Table 12-1. Add IP Interface Web Page Fields

Parameter	Allows you to...
Name	Enter a unique, alphanumeric name for the interface (maximum 30 characters)
VLAN	A pull down menu that associates a VLAN with this IP Interface. **Note:** The way that you configure a port(s) to a VLAN determines the IP Routing option that you select for the interface. Use the following options to configure the appropriate interface: • **Default** - Selects the default VLAN.(VLAN ID=1) • **Discard** - Selects the VLAN to be discarded.(VLAN ID= 4097) • **VLANs-** A list of all VLANs created on the switch. *** Note:** The following two parameters, **Ethernet-Console** and **Serial-Console** are special IP interface settings, and are not part of the normal routed data path. Do not configure routing only or routing and management for the serial interface or the Ethernet Console, configure both for management only. • **Ethernet-Console** - Creates an IP interface and assigns the IP address to the Supervisor's Ethernet port. • **Serial Console** - Creates an IP interface and assigns the IP address to the Supervisor's serial port (RS232). This is necessary if you are using PPP to connect to the supervisor's serial port.
Interface Type	Ethernet LAN or NBMA. *Non-Broadcast Multi-Access* (NBMA) IP Interfaces. Makes it possible for the switch to exchange routing information over nonbridged connections (routed virtual switch ports (VSPs)) NBMA functionality was added to RIP and OSPF routing protocols on the Avaya Multiservice switch software.
Administrative State	Specify the administrative state of the interface. Options include: • **UP** - The interface is active. • **DOWN** - The interface is inactive.
Network Address	Enter the network IP address for the selected interface.
Mask	Enter the subnet mask for the interface.
	1 of 3

Table 12-1. Add IP Interface Web Page Fields

Parameter	Allows you to...
MAC Format	Select the MAC address format for the interface. Options include: • Ethernet V2 • Simple Network Access Protocol (SNAP)
ARP Timeout	Enter the timeout period for Address Resolution Protocol (ARP) in seconds.The ARP Timeout value is the period of time the switch will wait for data from a station, if no data is received, the station is deleted from the ARP Table. The default is **14,400** seconds (4 hours).
Directed Broadcast	When the IP Directed Broadcast feature is Enabled, it allows a Network Directed Broadcast (unicast IP address with the Host ID set to all One's) to be forwarded by the router on the selected interface. The default is **Enable.**
IP Routing	Select enable or disable IP routing on the interface. Options include: • **Routing/Mgmt** (Default) - Enables you to manage the switch, from the Command Line Interface (CLI) or the Web Agent, and configure IP routing for the switch. • **Mgmt Only** - Enables you to manage the switch, however, IP routing is disabled for this interface. **Note:** Do not enable routing protocols on an interface configured for Mgmt Only since the interface will act as an end point and will not pass traffic. • **Routing Only** - IP routing is enabled on the interface, but you cannot manage the switch through the interface. Routing Only interfaces do not permit management traffic destined for local interfaces but do allow all other traffic including management traffic destined for interfaces on other switches.
RIP	Enable or disable RIP. The default value is **Disable**.
OSPF	Enable or disable OSPF. The default value is **Disable.** **Note:** To Enable OSPF, The Admin State of the IP interface must be set to down. Enable OSPF and then set the Admin State to up.
	2 of 3

Table 12-1. Add IP Interface Web Page Fields

Parameter	Allows you to...
Multicast Protocol	Specify the multicast protocol for the interface. Options include: • None (Default) • DVMRP • IGMP *** Note:** By Default, multicast forwarding is disabled. To enable multicast forwarding go to the IP Global Configuration page.
Proxy ARP	Enable or disable Proxy ARP. The default value is **Disable**.
ICMP Redirect	Enable or disable ICMP Redirect (IDRP). The default value is **Enable**.
NetBIOS UDP Rebroadcast	Enable or disable NetBIOS UDP Rebroadcasts. The default value is **Disable.**
VRRP	Enable or disable Virtual Redundancy Router Protocol. The default value is **Disable**.
BOOTP/DHCP Relay Gateway	Enable or disable the BOOTP/DHCP Relay Agent. The default value is **Enable**.
	3 of 3

CLI Commands

To create an IP interface, you must use two CLI commands:

1. Create an Interface Name first.

   `(configure)#` **interface <*interface_name*> type <*ethernet* | *nbma*>**

   Example: `(configure)#` **interface myif type ethernet**

2. Assign an IP address to the interface.

   `(confg-if:myif)#` **ip address <*ip-address*> <*mask*>**

   Example: `(config-if:myif)#` **ip address 192.168.97.1 255.255.255.0**

Enabling IP Routing Global Parameters

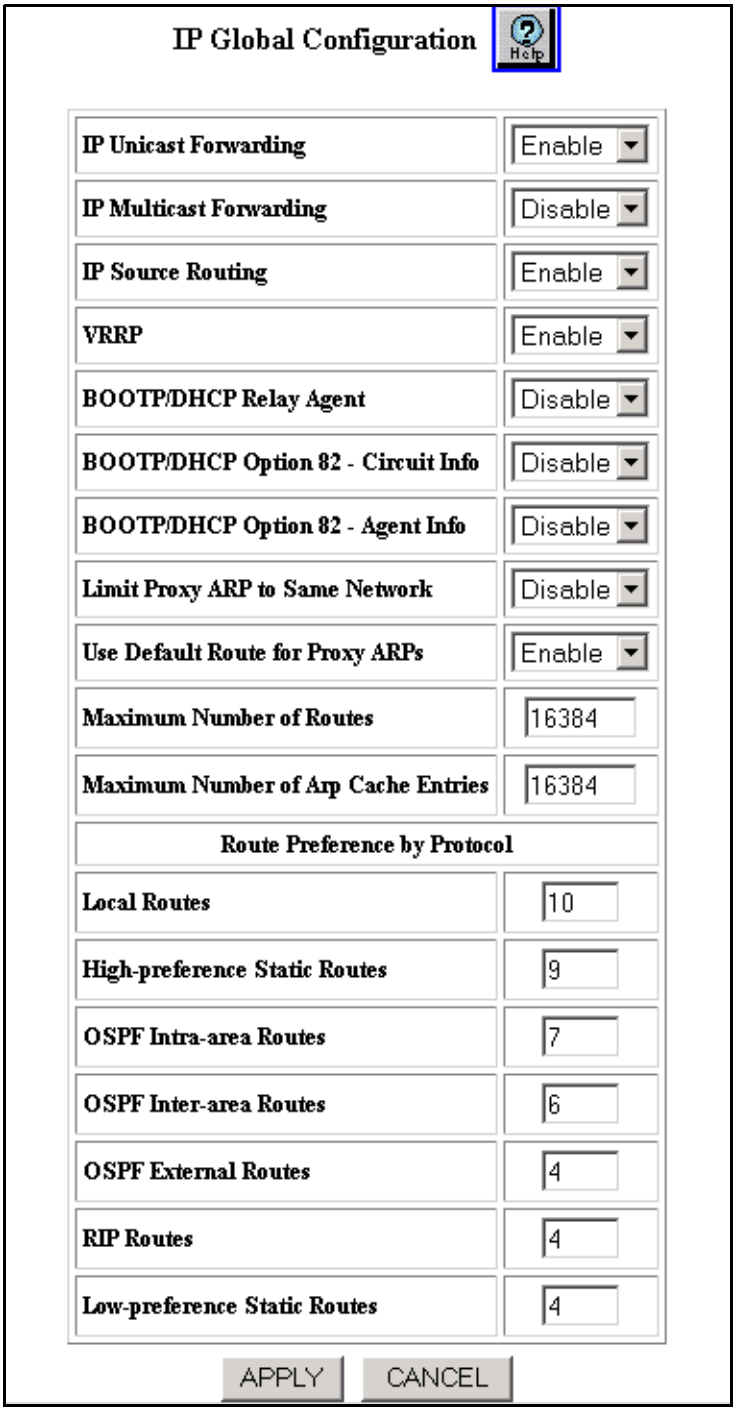
You can enable IP Unicast Forwarding (Routing) Globally using either the Web Agent or the CLI.

Web Agent Procedure

To enable IP routing using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **Global**. The IP Global Configuration Web page is displayed in the content pane. See Figure 12-4.

Figure 12-4. IP Global Configuration Web Page



2. See Table 12-2 to configure the IP Global Configuration Web page.
3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 12-2. IP Global Configuration Parameters

Parameter	Allows you to...
IP Unicast Forwarding	**Disable** IP Unicast Forwarding. Must be enabled for Unicast Forwarding. The default setting is **Enable**.
IP Multicast Forwarding	**Enable** IP multicast forwarding on a global basis. Must be enabled for Multicast Routing or Forwarding. The default setting is **Disable**.
IP Source Routing	**Disable** IP source routing globally. The default setting is **Enable**.
VRRP	**Disable** VRRP globally. The default setting is **Enable**.
BOOTP/DHCP Relay Agent	**Enable** Client requests for an IP address and forwards their requests to a server. This agent also relays responses from the server to the client. The default setting is **Disable**.
BOOTP/DHCP Option 82-Circuit Info	**Enable** This sub-option 82 identifies the slot and physical port number from which the DHCP request was received. The default setting is **Disable**. **Note:** **BOOTP/DHCP Relay Agent** must be set to enable, and must be enabled on the desired IP interface (enabled by default).
BOOTP/DHCP Option 82- Agent Info	**Enable** This sub-option 82 identifies the IP address and, if available, the system name of the switch. The Default setting is **Disable**. **Note:** **BOOTP/DHCP Relay Agent** must be set to enable, and must be enabled on the desired IP interface (enabled by default).
Limit Proxy ARP to Same Network	**Enable** Proxy ARP on the same network. When enabled, the router responds to ARP requests when the source and target IP address are in the same IP network and different IP subnets. When disabled, the router responds to ARP requests when the source and target IP address are in different networks. The default setting is **Disable**.
Use Default Route for Proxy ARPs	**Disable** the use of the default route on your switch as the route for Proxy ARPs. For example, if you have a default route configured to reach the 0.0.0.0 IP address, then any ARP request that does not match any of the other routes in your IP routing table will automatically go to this default route. The default setting is **Enable**.
	1 of 3

Table 12-2. IP Global Configuration Parameters

Parameter	Allows you to...
Maximum Number of Routes	Enter the maximum number of routes that you want added to the routing table. The default number of routes is **16384**. **Note:** These routes refer only to IP Unicast entries.
Maximum Number of ARP Cache Entries	Enter the maximum number of ARP cache entries. ARP entries refers to the space available for the ARP cache. When you increase the number of entries, it may cause the table to be relearned more frequently. Consequently, it will make address space usage more efficient. The default maximum number of entries is **16384**.
Route Preference by Protocol	This section describes the routing preferences set up according to the network administrator's preferences. These preferences are normally set up using the most to least trust. For example, local routes are normally considered to have *more trust* or a *higher preference*, while OSPF external routes are considered to have *less trust* or a *lower preference*. These preferences can be overridden, but careful consideration must be given when setting each. **Note:** Local Routes must always have the higher preference. **Note:** **Do not** change the Preferences unless you are certain of the consequences. Incorrect Preference settings can cause the switch to stop all routing.
Local Routes	Enter a preference value for local routes. **Note:** Local Routes must always have the higher preference.
High-Preference Static Routes	Enter a preference value for High preference static routes.
OSPF Intra-area Routes	Enter an OSPF intra-area route. A lower number indicates a lower preference for the path.
OSPF Inter-area Routes	Enter the inter-area paths to destinations in other OSPF areas. These are discovered through the examination of received summary Link State Advertisements (LSA). Enter a number to specify your path cost. A lower number indicates a lower preference for the path.
OSPF External Routes	Enter the number of external paths to destinations external to the Autonomous System (AS). These are detected through the examination of received AS external LSAs. Enter a number to specify your path cost (preference). A lower number indicates a lower preference for the path.
	2 of 3

Table 12-2. IP Global Configuration Parameters

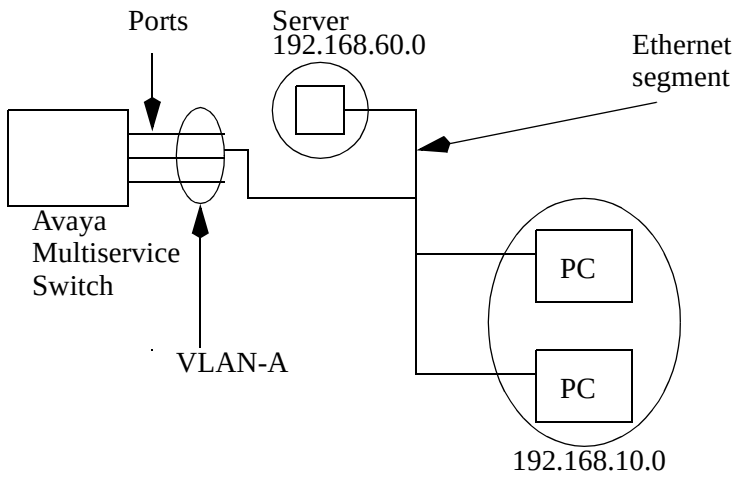
Parameter	Allows you to...
RIP Routes	Enter the number of RIP routes to use the hop count as a metric. Hence, to specify a preference for a RIP route, you need to enter a lower number (path cost).
Low-Preference Static Routes	Enter a preference value for Low preference static routes.
	3 of 3

IP Multinetting

You can set up a configuration of multiple subnets, and IP interfaces on a single VLAN. This is also referred to as **IP Multinetting**. A network is said to be multinetted when multiple logical groups of computers are brought together within a single broadcast domain. In the example (Figure 12-5), two separate networks are brought together in VLAN-A allowing computing resources on network 192.168.10.x to communicate with computing resources on network 192.168.60.x without having to go through a router. To create a multinetted network, you must assign multiple subnets to a VLAN.

*** Note:** Only one subnet can be BootP/DHCP Gateway. The second subnet must use Static IP addresses.

Figure 12-5. Multinet Diagram



Creating a Multinet Interface

To create a multinet interface:

1. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **Interfaces.** The IP Interfaces Web page is displayed in the content pane. See Figure 12-2.
2. Select **CREATE**. The Add IP Interface Web page is displayed. See Figure 12-3.
3. Enter a name for the interface in the **Interface** field.
4. Select a VLAN for this multinet interface from the **VLAN** field pull-down menu.
5. Enter the network address associated with this interface in the **Network Address** field.
6. Select **CREATE** from the bottom of the Add IP Address Web page.

   The IP Interfaces Web page is redisplayed and displays the new interface name listed in the **Interface** field.
7. Repeat steps 2 through 6 for any other interfaces that you want to associate with the same VLAN.

***Note:** You must have a different network address for each new multinetting interface that you create for the same VLAN.

***Note:** The first IP interface that has the BootP/DHCP Relay Gateway enabled will be the relay.

Configuring Short-Lived IP Protocol Filters

Overview

This feature helps conserve the forwarding engine resources of the switch. Depending on the configuration of your network, the forwarding cache may contain a high number of entries for short-lived protocols, for example, DNS and NTP. These short-lived protocol entries consume resources that could otherwise be used for data traffic of other protocols such as FTP, Telnet, HTTP, RealAudio, and online games.

If the forwarding cache contains a high number of short-lived protocol entries, you can set the switch to use the slow path to route packets of a specific short-lived protocol. By using the slow path to route short-lived protocol traffic, you conserve the resources of the forwarding engines, which provide the fast in band routing engine (FIRE) and fast out of band routing engine (FORE). For information about the slow path and FIRE and FORE paths, see "Routing Overview" in Chapter 1, "Introduction."

To configure short-lived protocol filters, your user account must be assigned one of the following access types:

- Read-write
- Administrator
- Custom access type that has read-write permission for the routing feature

You can configure short-lived IP protocol filters only by using the CLI. This feature is not available in the Web Agent or SNMP.

This section contains procedures for the following tasks:

- Enabling a Short-Lived IP Protocol Filter
- Disabling a Short-Lived IP Protocol Filter
- Displaying Short-Lived IP Protocol Filters

Enabling a Short-Lived IP Protocol Filter

Overview

If you enable a filter for a short-lived IP protocol, the switch uses the slow path to route packets of the protocol. For information about the slow path, see "Routing Overview" in Chapter 1, "Introduction."

By default a filter is enabled for DNS and NTP packets. To route the packets of other short-lived protocols by means of the slow path, you must enable a filter for the protocol.

CLI Command

To enable a filter for a short-lived IP protocol, enter Global Configuration mode and use the following command:

`(configure)#` **ip short-lived {tcp | udp}** <***port***>

Table 12-3. Keywords, Arguments, and Options

Keyword, Argument, or Option	Definition
{tcp | udp}	Enter **tcp** if the protocol that you want to filter uses a TCP port. Enter **udp** if the protocol that you want to filter uses a UDP port.
<*port*>	The TCP or UDP port number that the protocol uses. Enter a port number from 0 through 65535.

Examples

To send all SNMP packets to supervisor module for slow path routing, enter the following command:

ip short-lived udp 161

To send all BOOTP and DHCP packets to the supervisor for slow path routing, enter the following command:

ip short-lived udp 67

ip short-lived udp 68

Disabling a Short-Lived IP Protocol Filter

Overview

If you disable a filter for a short-lived IP protocol, the switch uses either the fast in band routing engine (FIRE) path or the fast out of band routing engine (FORE) path to route packets of the protocol. For information about the FIRE and FORE paths, see "Routing Overview" in Chapter 1, "Introduction.".

CLI Command

To disable the filter for a short-lived IP protocol, enter Global Configuration mode and use the following command:

`(configure)#` **no ip short-lived {tcp | udp}** <***port***>

Table 12-4. Keywords, Arguments, and Options

Keyword, Argument, or Option	Definition
{tcp | udp}	Enter **tcp** if the protocol that you want to filter uses a TCP port. Enter **udp** if the protocol that you want to filter uses a UDP port.
<*port*>	The TCP or UDP port number that the protocol uses.

Examples

To disable the short-lived IP protocol filter for DNS, enter the following command:

no ip short-lived udp 53

To disable the short-lived IP protocol filter for NTP, enter the following command:

no ip short-lived udp 123

Displaying Short-Lived IP Protocol Filters

CLI Command

To display the short-lived IP protocol filters that are currently enabled, enter Global Configuration mode and use the following command:

`(configure)#` **show ip short-lived**

Example

After you enter the **show ip short-lived** command, the switch displays the filters that are currently enabled. For example:

```
ip short-lived tcp 112
ip short-lived udp 53
ip short-lived udp 123
```

Creating IP Static Routes

You can create IP static routes using either the Web Agent or the CLI.

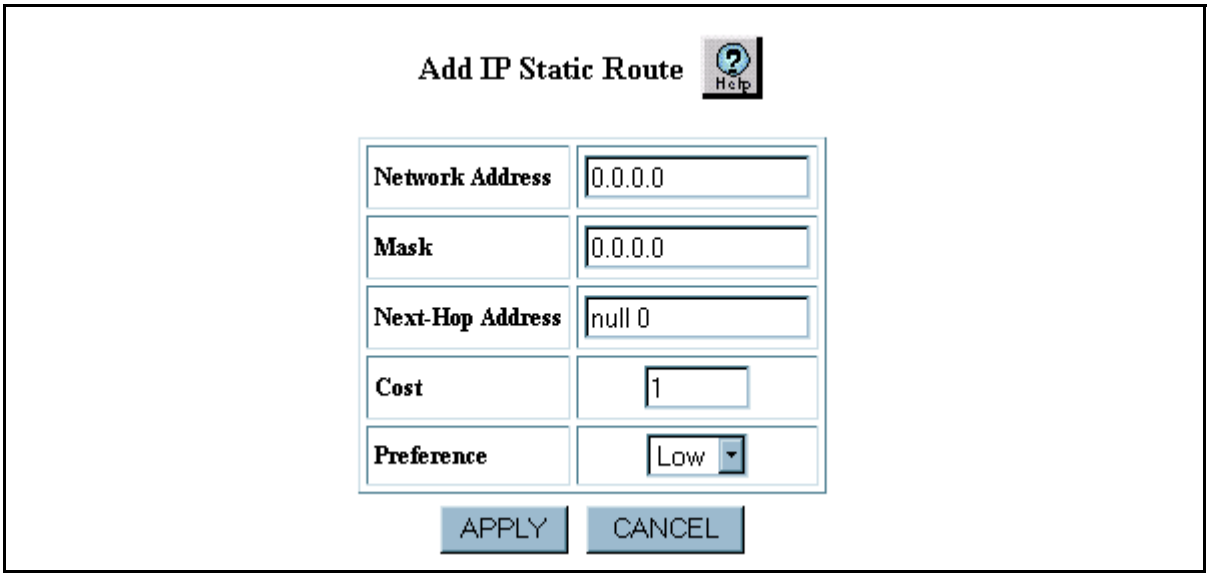
***Important:** Do not create a static route to a local interface. This configuration is not supported.

Web Agent Procedure

To create IP static routes using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders.
2. Click **Static Routes**. The IP Static Routes Web page is displayed in the content pane.
3. Click **CREATE**. The Add IP Static Route Web page is displayed in the content pane. See Figure 12-6.

Figure 12-6. Add IP Static Route Web Page



4. Configure the static route as appropriate. See Table 12-5 for an explanation of the Web page fields.
5. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 12-5. IP Static Route Parameters

Field	Explanation
Network Address	Enter an IP network address for your IP static route.
Mask	Enter an IP subnet mask for your IP static route.
Next-Hop Address	Enter an IP address for the gateway associated with the IP static route. The default setting is **null 0**, which creates a discard route. For information on creating a discard route, see "Creating a Static Route to a Null Interface."
Cost	Enter the metric between this router and the destination.
Preference	Select a *low* or *high* routing preference from the Preference field pull-down menu. Preference overrides cost. If two routes of the same preference are present, the switch uses the route that has the lower cost.

CLI Commands

To create and show IP static routes use the following CLI commands:

- To create an IP static route, `(configure)#` **ip route** <***ip-addr***> <***mask***> <***next hop IP address***> <***cost***> <***preference***>
- To view all IP static routes that are configured on the switch, `(configure)#` **show ip route static**

Creating a Static Route to a Null Interface

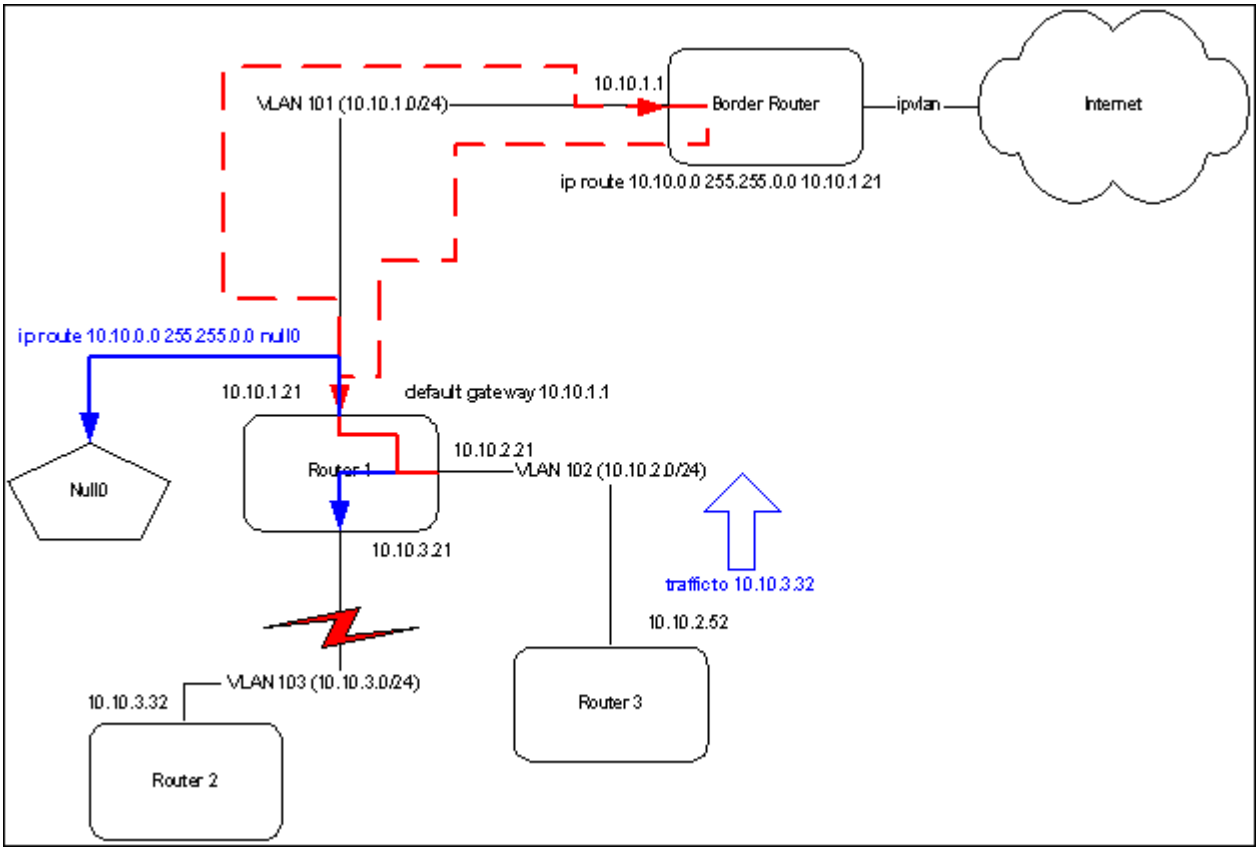
Overview

A null interface is a virtual interface that discards IP packets and is used to prevent routing loops from occurring in the network.

You do not assign the null interface an IP address. Instead, you create a static route for a network and set the next hop to the null interface (null0).

Figure 12-7 shows an example of a null interface preventing loops in the network.

Figure 12-7. Null Interface Example



In the preceding figure, Router 1 has a default route that points to the Border Router for addresses that are not known within the enterprise.

If a null interface is not configured:

1. Router 3 forwards packets that have a destination of 10.10.3.32 to Router 1.
2. Because VLAN 103 is down, Router 1 sends the packets to the Border Router.
3. The Border Router then sends the traffic back to Router 1 via the route 10.10.0.0/16.

The loop that occurs is shown as a red dotted line in Figure 12-7.

To prevent such a loop, you can create a static route from 10.10.0.0/16 to the null interface. Once this static route to the null interface is created, Router 1 drops all packets whose destination address is unreachable and within the 10.10.0.0/16 network.

When packets are sent to the null interface, the router sends an Internet Control Message Protocol (ICMP) destination unreachable message to the source of the packet (Router 3 in Figure 12-7).

Procedure

Web Agent Procedure

To create a static route to a null interface:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders.
2. Click **Static Routes**. The IP Static Routes Web page is displayed in the content pane.
3. Click **CREATE**. The Add IP Static Route Web page is displayed in the content pane. See Figure 12-6.
4. In the **Network Address** field, enter the IP address of the network for which you are creating a null interface.
5. In the **Mask** field, enter the subnet mask for the network.
6. In the **Next-Hop Address** field, enter **null 0**.
7. In the **Cost** field, enter the routing metric of the path. For an explanation of this field, see Table 12-5.
8. In the **Preference** field, select either **high** or **low**. For an explanation of this field, see Table 12-5.
9. Click **APPLY**.

CLI Command

To create a static route to a null interface, use the following CLI command:

`(configure)#` **ip route** <***route-addr***> <***mask***> **null 0** <***cost***>

For example, to create the static route to the null interface in Figure 12-7, you would enter **ip route 10.10.0.0 255.255.0.0 null 0** <***cost***>.

Creating IP Static ARP Entries

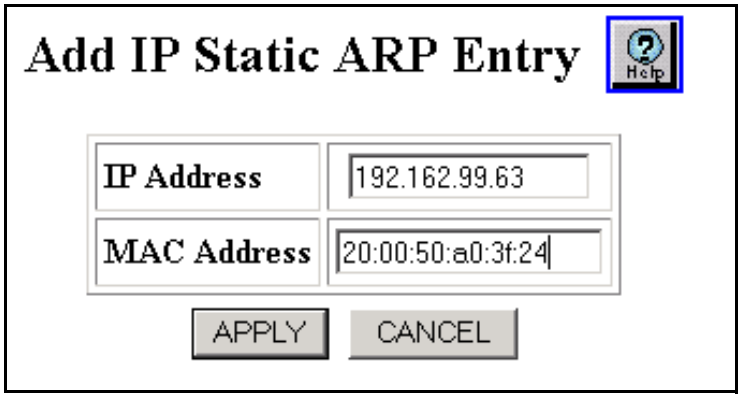
You can create static ARP entries on the switch using the Web Agent or the CLI.

Web Agent Procedure

To create a static ARP entry in your switch's ARP cache using the Web Agent:

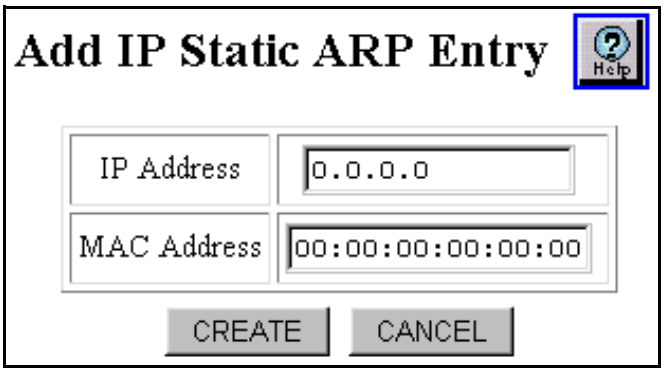
1. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **Static ARP**. The IP Static ARP Entries Web page is displayed in the content pane. See Figure 12-8.

Figure 12-8. IP Static ARP Entries Web Page



2. Select **CREATE**. The Add IP Static ARP Entry Web page is displayed. See Figure 12-9.

Figure 12-9. Add IP Static ARP Entry Web Page



3. See Table 12-6 to configure the Add IP Static ARP Entry Web page parameters:

Table 12-6. IP Static ARP Parameters

Parameter	Allows you to...
IP Address	Enter an IP address to associate with the Static ARP entry.
MAC Address	Enter the MAC address of a node to which you want to create a static ARP entry.

4. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To create IP static ARP entries, use the following CLI command:

`(configure)#` **arp** <***ip-address***> <***hardware-address***>

Creating a BOOTP/DHCP Server Entry

The BOOTP/DHCP Server Entry allows you to configure a Router as a BOOTP/DHCP Relay Agent between a BOOTP/DHCP server and the requesting client.

You can create a BOOTP/DHCP Server entry using either the Web Agent or the CLI.

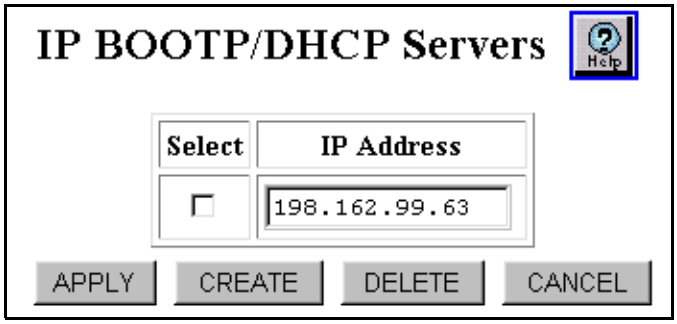
*** Note:** BootP/DHCP must first be enabled in the **IP Global Configuration.**

Web Agent Procedure

To create a BOOTP/DHCP entry using the Web Agent:

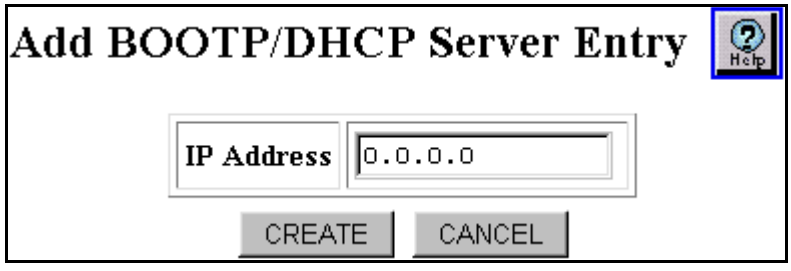
1. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **BOOTP/DHCP Servers**. The IP BOOTP/DHCP Servers Web page is displayed in the content pane. See Figure 12-10.

Figure 12-10. IP BOOTP/DHCP Server Web Page



2. Select **CREATE**. The Add BOOTP/DHCP Server Entry Web page is displayed. See Figure 12-11.

Figure 12-11. Add BOOTP/DHCP Server Entry Web Page



3. Enter the BOOTP/DHCP server IP address in the **IP Address** field.
4. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings.

*** Note:** It is possible to create multiple BOOTP/DHCP Server Entries if necessary.

CLI Command

To create a BOOTP/DHCP entry using the CLI, enter the following command in Configure mode:

`(configure)#` **ip boot-dhcp server** <***ip-address***>

Configuring Route Redistribution

Overview

Route redistribution makes it possible for different IP routing protocols to exchange routing information. To configure route redistribution, you create route redistribution entries that control the distribution of static, local, or dynamically learned routes from one protocol to another protocol. Route redistribution is supported only by dynamic routing protocols, such as RIP and OSPF.

When creating route redistribution entries, you set a source protocol and destination protocol. The *source protocol* is the protocol whose routes are redistributed, and the *destination protocol* is the protocol to which the routes are redistributed. You can set the entry to redistribute either all routes or specific routes. If you want to redistribute only specific routes, you must set up an access list to either permit or deny specific routes for redistribution. For more information about access lists, see Chapter 13, "Configuring Access Lists."

*** Note:** Avaya recommends that you do not globally enable an access list that you use to redistribute specific routes.

*** Note:** Route redistribution supports only standard access rules. You cannot use extended access rules to permit or deny specific routes for redistribution.

Table 12-7 shows the source and destination protocols between which you can set up route redistribution filters.

Table 12-7. Supported Source and Destination Protocols

Source Protocol	Destination Protocol
Local	OSPF and RIP
OSPF	RIP
Static	OSPF and RIP
RIP	OSPF

This section contains the following procedures:

- Creating an IP Redistribute List Entry
- Viewing IP Redistribute List Entries
- Modifying an IP Redistribute List Entry
- Deleting an IP Redistribute Entry

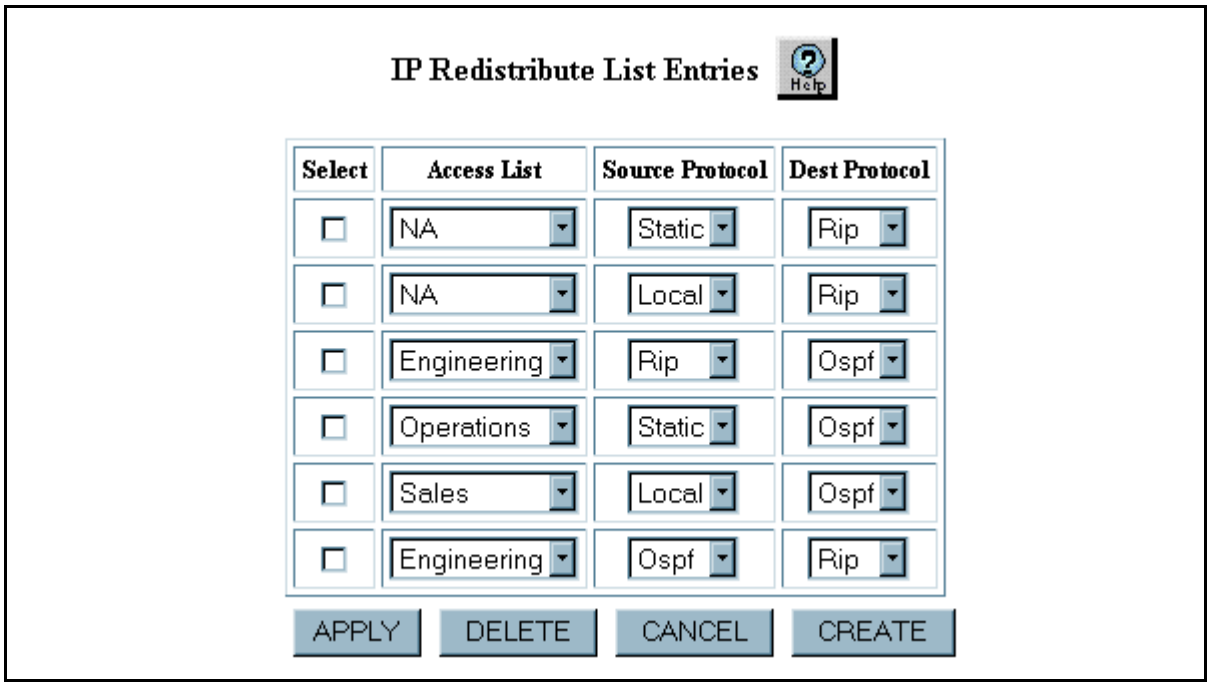
Creating an IP Redistribute List Entry

Web Agent Procedure

To create an IP redistribute list entry:

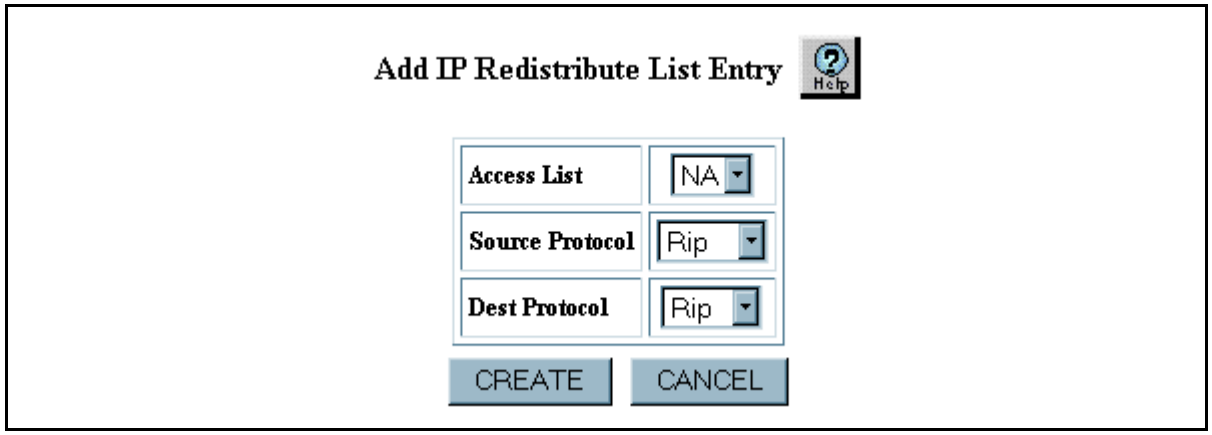
1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders.
2. Click **IP Redistribute List**. The IP Redistribute List Entries Web page is displayed in the content pane. See Figure 12-12.

Figure 12-12. IP Redistribute List Entries



3. Click **Create**.
4. The Add IP Redistribute List Entry Web page is displayed in the content pane. See Figure 12-13.

Figure 12-13. Add IP Redistribute List Entry



5. In the **Access List field**, select an access list to filter the routes that are redistributed. If you select **NA**, all routes are redistributed.

*** Note:** You can use different access lists for different IP redistribute list entries. All IP redistribute list entries do not have to use the same access list.

6. In the **Source Protocol** field, select the protocol whose routes you want to redistribute. Options are RIP, OSPF, static, and local.

7. In the **Destination Protocol** field, select the protocol to which you want the routes redistributed. Options are Rip and OSPF.

*** Note:** Selecting OSPF as the destination protocol causes OSPF adjacencies to be reestablished. During this reestablishment, a temporary loss of traffic occurs.

8. Click **Create**.

CLI Commands

To create an IP redistribute list entry, use the following CLI command:

- To redistribute routes to RIP, `(configure router:rip)#` **redistribute {ospf | local | static} [<*access-list-name*>]**
- To redistribute routes to OSPF, `(configure router:ospf)#` **redistribute {rip | local | static} [<*access-list-name*>]**

Viewing IP Redistribute List Entries

Web Agent Procedure

To view IP redistribute list entries:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders.
2. Click **IP Redistribute List**.

   The IP Redistribute List Entries Web page is displayed in the content pane. See Figure 12-12.

CLI Command

To view IP redistribute list entries, use the following CLI command:

> **show ip redistribute**

Modifying an IP Redistribute List Entry

*** Note:** Modifying an IP redistribute list entry that has OSPF for its destination protocol causes OSPF adjacencies to be reestablished. During this reestablishment, a temporary loss of traffic occurs.

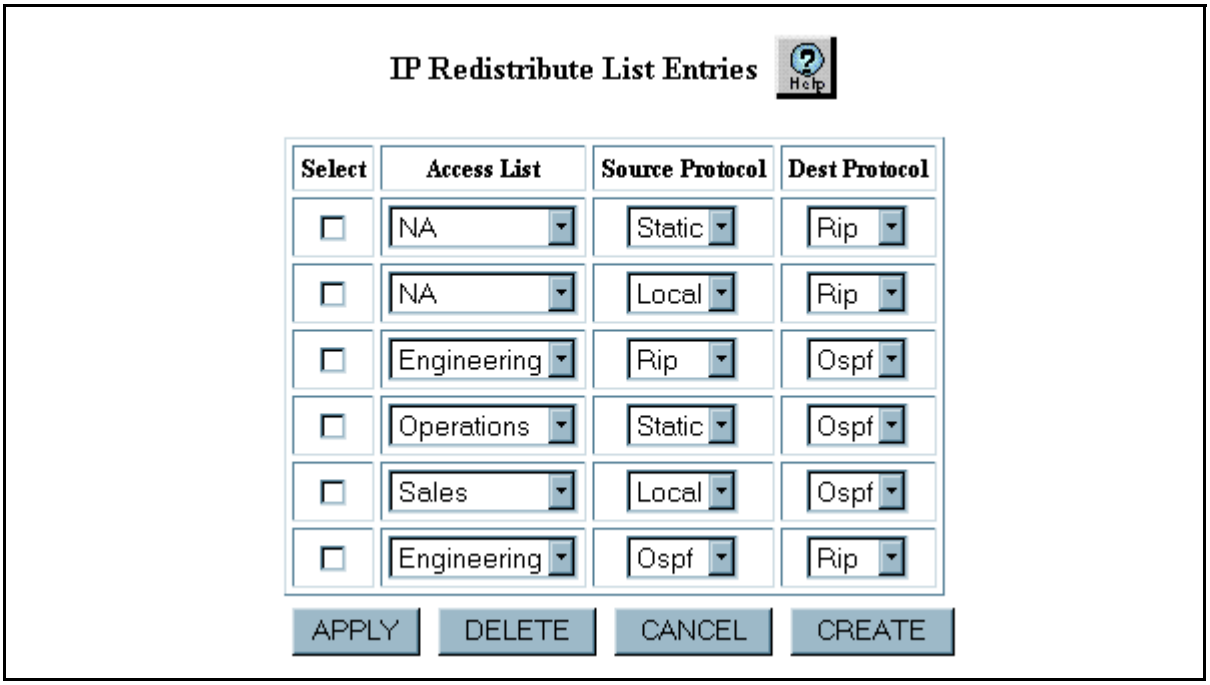
Web Agent Procedure

To modify an IP redistribute list entry:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **IP Redistribute List**.

   The IP Redistribute List Entries Web page is displayed in the content pane. See Figure 12-14.

Figure 12-14. IP Redistribute List Entries Web Page



2. Select the route redistribution entry that you want to modify.
3. Modify the entry as necessary.
4. Click **APPLY**.

CLI Command

To modify an IP redistribute list entry, use the same **redistribute** command that you use to create a route redistribution entry. For more information, see "Creating an IP Redistribute List Entry" earlier in this chapter.

Deleting an IP Redistribute Entry

*** Note:** Deleting an IP redistribute list entry that has OSPF for its destination protocol causes OSPF adjacencies to be reestablished. During this reestablishment, a temporary loss of traffic occurs.

Web Agent Procedure

To delete an IP redistribute list entry:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **IP Redistribute List**.

   The IP Redistribute List Entries Web page is displayed in the content pane. See Figure 12-12.

2. Select the route redistribution entry that you want to delete.

3. Click **Delete**.

CLI Commands

To delete an IP redistribute list entry, use the following CLI commands:

- To delete an entry that redistributes routes to RIP, `(configure router:rip)` **no redistribute {ospf | local | static}**
- To delete an entry that redistributes routes to OSPF, `(configure router:ospf)` **no redistribute {rip | local | static}**

IP Multicast

IP Multicast enables a single host to distribute information to multiple recipients. To do this, multicast protocols use class D IP addresses to specify specific multicast groups to which information is sent. The class D IP address used by multicast routing protocols, ranges from 224.0.0.1 to 224.0.0.255. The class D IP addresses available for general use are 224.0.1.0 to 239.255.255.255.

In addition, IP multicasting distributes information to multicast groups in two specific ways:

- **Multicast Forwarding** - allows a switch to forward multicast traffic from the local multicast server to group members on directly attached subnetworks. If a multicast packet is forwarded to multiple interfaces on one VLAN, only one Forwarding Entry is added on the VLAN for the packet. One copy of the packet is sent to the VLAN. To configure interfaces for multicast forwarding select Internet Group Management Protocol (IGMP) for the multicast protocol.

*** Note:** Interfaces configured for IGMP forwarding will not be able to participate in multicast sessions distributed through DVMRP Routing.

- **Multicast Routing** - allows a switch to deliver multicast traffic between neighboring routers and across the network using Distance Vector Multicast Routing Protocol (DVMRP).To configure interfaces for multicast routing select DVMRP for multicast protocol on all interfaces that are part of a multicast network.Since IGMP is required on all interfaces that have multicast clients, IGMP is enabled automatically on all interfaces configured for DVMRP.

*** Note:** You must globally enable IP multicast routing in order to successfully configure IGMP or DVMRP. See “Routing Function” earlier in this chapter for instructions on enabling IP multicast routing.

This section contains the following procedures:

- Configuring IGMP
- Globally Enabling IGMP
- Modifying IGMP Interfaces
- Configuring and Modifying DVMRP
- Configuring the DVMRP Global Configuration
- Modifying a DVMRP Interface

Configuring IGMP

IGMP enables hosts to inform routers when they join or leave groups. Routers periodically query hosts (query interval) for the groups in which the hosts are members. When there is more than one router in a broadcast domain (subnet), one of the routers becomes the designated querier. Only the designated router queries the hosts.

Both IGMP Version 1.0 and IGMP Version 2.0 are supported. After selecting the specific IGMP version for an interface, you can manually configure the Version 1.0 querier. The selection of the querier for Version 2.0 is dynamic but can be overridden.

Globally Enabling IGMP

Configuring IGMP includes enabling IGMP and the MTRACE globally on the switch. You can do this from either the Web Agent or the CLI.

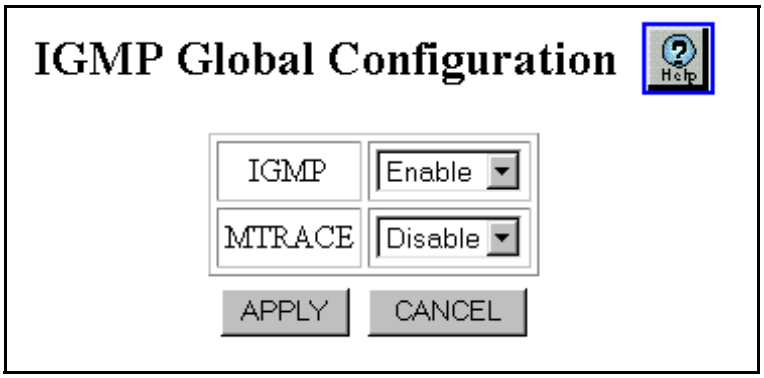
*** Note:** You also have to enable multicast forwarding from the Routing > IP > Configuration > Global

Web Agent Procedure

To globally enable IGMP from the Web Agent:

1. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **Global Configuration**. The IP Global Configuration Web page is displayed in the content pane. See Figure 12-4.
2. Select **Enable** from the **IP Multicast Forwarding** field pull-down menu to enable IP multicast globally.
3. In the navigation pane, expand the **Routing > IGMP** folders, and then click **Global Configuration**. The IGMP Global Configuration Web page is displayed in the content pane. See Figure 12-15.

Figure 12-15. IGMP Global Configuration Web Page



4. Select **Enable** from the **IGMP** field pull-down menu to enable IGMP.
5. Select **Enable** from the **MTRACE** field pull-down menu to enable MTRACE processing.
6. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To globally enable IGMP using the CLI, enter the following command from Configure mode:

```
(configure)# set router igmp
```

Modifying IGMP Interfaces

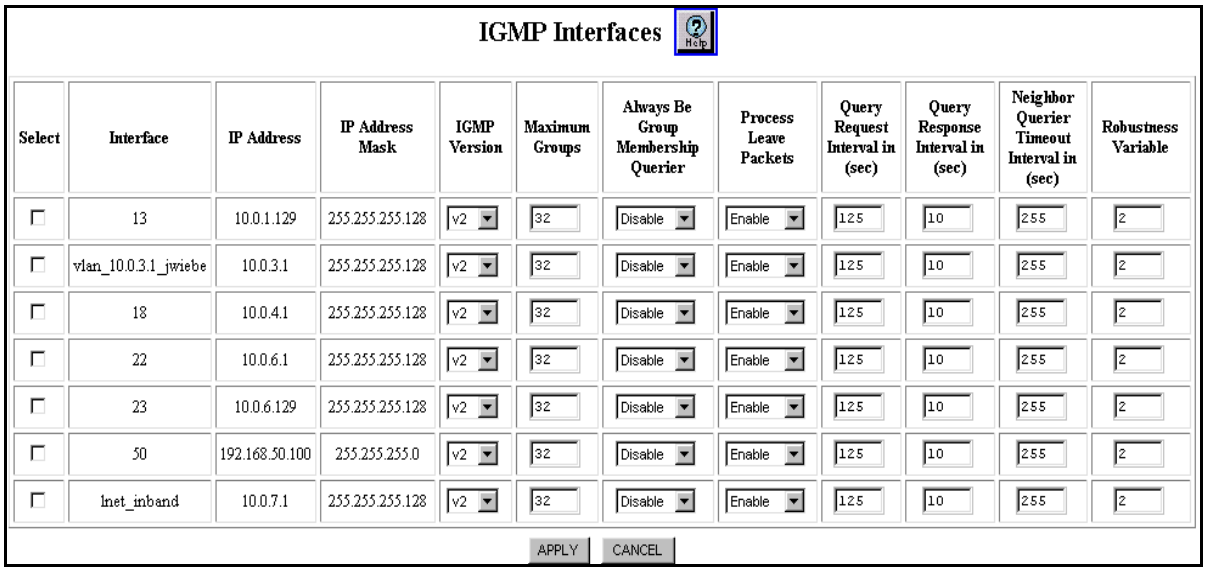
Web Agent Procedure

To modify IGMP interfaces from the Web Agent:

1. In the navigation pane, expand the **Routing** > **IGMP** folders, and then click **Interfaces**. The IGMP Interfaces Web page is displayed in the content pane. See Figure 12-16.

   ***Note:** You must first enable "multicast protocol"on an interface before you can configure IGMP. See "Displaying Existing IP Interfaces" earlier in this chapter for more information on how to select a multicast protocol.

Figure 12-16. IGMP Interfaces Web page



2. See Table 12-8 to modify an IGMP interface.

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 12-8. IGMP Interface Parameters

Parameter	Allows you to...
Select	Select the interface to be modified.
Interface	Displays the IP interface that is configured with IGMP.
IP Address	Displays the IP address associated with this interface. **Note:** You cannot configure this field from the IGMP Configuration Web page.
IP Address Mask	Displays the subnet mask associated with this interface. **Note:** You cannot configure this field from the IGMP Configuration Web page.
IGMP Version	Select the IGMP Version (1.0 or 2.0) to be associated with the IGMP interface.
Maximum Groups	Enter the number of IGMP Groups that can be active on this interface. The default setting is 32. The valid range for this field is 1 to 7,000.
	1 of 2

Table 12-8. IGMP Interface Parameters

Parameter	Allows you to...
Always be Group Membership Querier	**Version 1.0** - Select **Enable** to make this interface the designated querier. The default is **Disable**. **Version 2.0** - Select **Enable** to force this interface to send queries. The default is **Disable** which tells the interface to obey the designated querier election. **Note:** Only the designated router will query hosts on your network.
Process Leave Packets	turns off processing of IGMP leave messages when set to disable. The default value is **Enable**.
Query Request Interval in (sec)	Enter a time period between queries. The valid range for this field is 1 to 65,535. The default setting is 125 seconds.
Query Response Interval in (sec)	Enter a time (in seconds) to wait for a response from a host after a query is sent. If no response is received within this time, the host is removed from the group table. The valid range for this field is 1 to 25. The default setting is 10 seconds.
Neighbor Querier Timeout Interval in (sec)	Enter a time (in seconds) this interface should wait after hearing a neighbor's query before assuming the role of querier, if not already the querier. If no query is received from a neighbor with a lower IP address in the allotted time, this interface becomes the querier (IGMP Version 2.0 only). The valid range for this field is 30-600. The default setting is 255 seconds.
Robustness Variable	Modify this field for any expected packet loss on a subnet. If a subnet is expected to have more packet loss, the Robustness Variable should be increased. The Robustness Variable *must not* be set to 0 and *should not* be set to 1. The valid range for this field is 1 to 65,535. The default setting is 2.
	2 of 2

CLI Command

To modify an IGMP interface using the CLI, enter the following command from interface mode:

`(config-if:<`*`interface`*`>)#` **ip igmp** <***options***>

Configuring and Modifying DVMRP

Distance Vector Multicast Routing Protocol (DVMRP) uses IP packets with protocol type 2 (IGMP) to exchange routing datagrams. DVMRP enables multicast routers to exchange distance vector updates that contain multicast flow lists and their corresponding cost metrics. DVMRP may use tunneling between pairs of DVMRP routers when traffic must pass through one or more intermediary routers or gateways that do not implement DVMRP.This implementation adheres to the specification for DVMRP V3.

You can configure DVMRP Globally using either the Web Agent or the CLI.

*** Note:** When processing heavy traffic, the switch may lose a small number of DVMRP neighbor-to-neighbor probe messages. The loss of these messages may cause multicast routing to become unstable.

Configuring the DVMRP Global Configuration

Web Agent Procedure

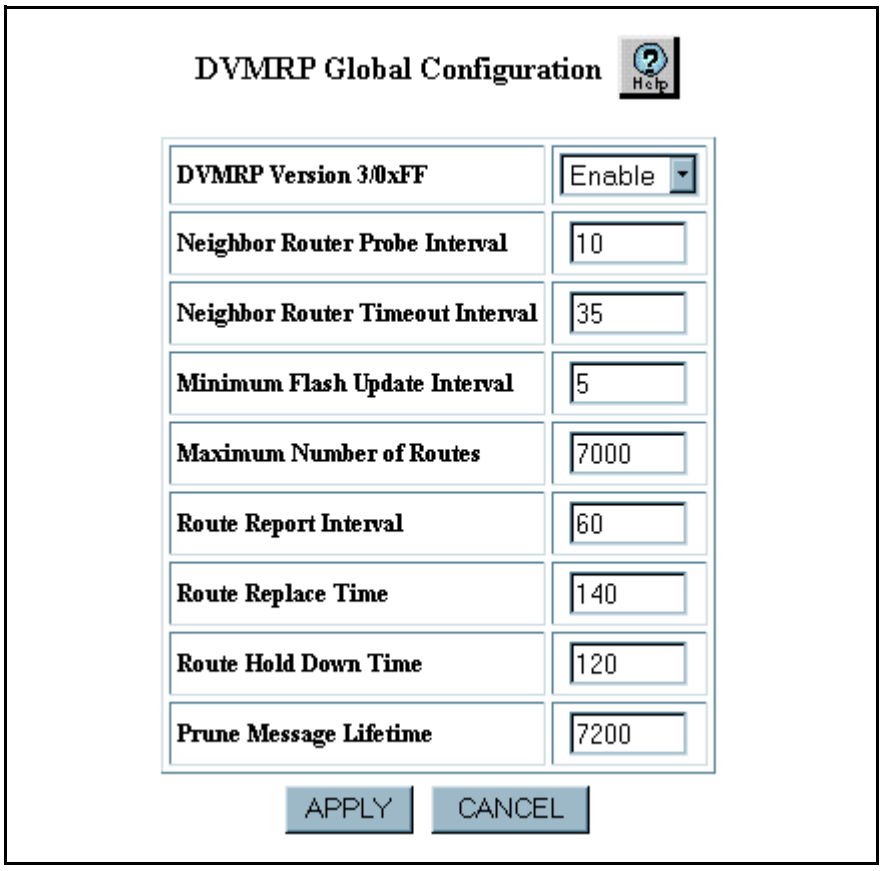
To configure DVMRP globally using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **Global Configuration**. The IP Global Configuration Web page is displayed in the content pane. See Figure 12-15.

2. Select **Enable** from the **IP Multicast Forwarding** field to enable IP multicast globally.

3. In the navigation pane, expand the **Routing > DVMRP** folders, and then click **Global Configuration**. The DVMRP Global Configuration Web page is displayed in the content pane. See Figure 12-17.

Figure 12-17. DVMRP Global Configuration Web Page



4. Select **Enable** from the **DVMRP Version 3/0xFF** field.
5. See Table 12-9 to configure the DVMRP Global Configuration Web page parameters.
6. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 12-9. DVMRP Global Configuration Web Page Parameters

Parameter	Allows you to...
DVMRP Version 3/xFF	Select **Disable** to globally disable DVMRP. The default value is **Enable**.
Neighbor Router Probe Interval	Enter **probe interval** (in seconds) for the switch to probe the network for available neighbor routers. The valid range for this field is 5 to 45 seconds. The default setting is 10 seconds.
	1 of 2

Table 12-9. DVMRP Global Configuration Web Page Parameters

Parameter	Allows you to...
Neighbor Router Timeout Interval	Enter the **time-out interval** (in seconds) that a neighbor stays up without confirmation. This is an important method used to time-out old routes. The valid range for this field is 10 to 50 seconds. The default setting is 35 seconds.
Minimum Flash Update Interval	Enter the **update interval** (in seconds) between flash updates. This represents the minimum time between advertisements of the same route. The valid range for this field is 5 to 20 seconds. The default setting is 5 seconds.
Maximum Number of Routes	Enter the maximum number of routes for this interface. The valid range for this field is 10 to 20,000. The default setting is 7,000.
Route Report Interval	Enter the **report interval** (in seconds) that elapses between delivery of DVMRP routing table updates. The valid range for this field is 30 to 90 seconds. The default setting is 60 seconds.
Route Replace Time	Enter the amount of **replace time** (in seconds) before which a route entry will be removed if it is not refreshed. The valid range for this field is 70 to 190 seconds. The default setting is 140 seconds.
Route Hold Down Time	Set the **hold down time** (in seconds) that the switch reports unavailable routes with a metric of infinity. The valid range for this field is 120 to 380 seconds. The default setting is 120 seconds.
Prune Message Lifetime	Enter the **lifetime** (in seconds) that a transmitted upstream prune message persists. The valid range for this field is 100 to 7,200 seconds. The default setting is 7,200 seconds.
	2 of 2

CLI Command

To configure DVMRP globally using the CLI:

```
(configure)# router dvmrp
```

Modifying a DVMRP Interface

You can modify a DVMRP interface using either the Web Agent or the CLI.

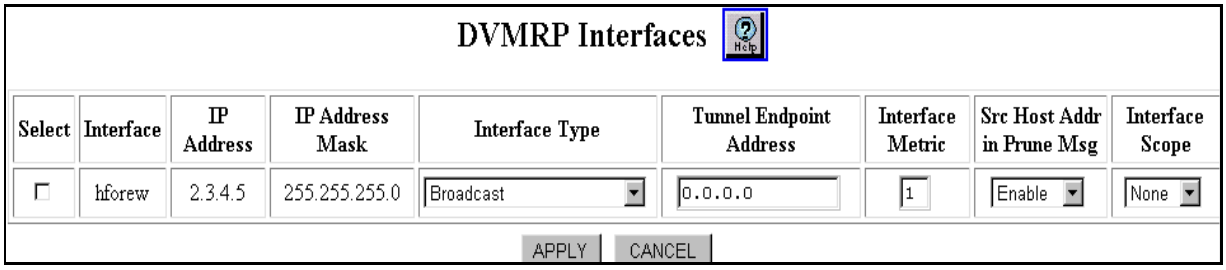
Web Agent Procedure

To modify a DMVRP interface using the Web Agent:

1. In the navigation pane, expand the **Routing > DVMRP** folders, and then click **Interfaces**.

   The DVMRP Interfaces Web page is displayed in the content pane. See Figure 12-18.

Figure 12-18. DVMRP Interfaces Web Page



*** Note:** You must first set "multicast protocol" on an interface to DVMRP before you can configure DVMRP. See "Displaying Existing IP Interfaces," earlier in this chapter, and enable a multicast protocol for this interface.

2. See Table 12-10 to complete your DVMRP configuration

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 12-10. DVMRP Interface Parameters

Parameter	Defines the...
Select	Select a DVMRP interface that you want to configure.
Interface	Displays the Interface that is configured with the DVMRP multicast protocol.
IP Address	Displays the IP address of each interface. You cannot modify this field.
IP Address Mask	Displays the Subnet mask for each interface. You cannot modify this field.
	1 of 2

Table 12-10. DVMRP Interface Parameters

Parameter	Defines the...
Interface Type	Select an Interface type. You can configure the interface type as: • **Broadcast** - All traffic is forwarded through the routers. This is not a tunnel and does not require encapsulation. • **Non-Encapsulated Tunnel** - All multicast data traffic is IPIP encapsulated, but the protocol messages are unicast. • **IPIP Tunnel** - All multicast traffic (data and protocol messages) are encapsulated in IP unicast packets with the protocol set to IPIP (IP in IP).
Tunnel Endpoint Address	Displays the Tunnel endpoint IP address of a router. You can modify this setting to represent the end router IP address to which you want to send datagrams through a tunnel. The origin and endpoint routers are separated by a gateway(s) or a router(s) that do not support DVMRP.
Interface Metric	Displays the Cost metric for the interface. The valid range for this field is 1 to 31. The default setting is 1.
Src Host Addr in Prune Msg	Select **Disable to s**end prune messages with only the subnet portion of the source address. The default is **Enable (**sends prune messages with the full source host address).
Interface Scope	Select the minimum TTL (time-to-live) required for a packet to leave this interface. The options are None, 127, and 255.
	2 of 2

CLI Command

To configure the DVMRP Interface using the CLI:

`(configure-if:<`*`interface`*`>)#` **ip dvmrp** <***options***>

Monitoring Switch Performance Using IP Statistics

This section provides detailed information on the analysis and use of IP and IP multicast statistics.

IP routing statistic options include:

- Displaying Global IP Routing Statistics
- Searching the IP Routing Table
- Displaying the IP Routing Table Statistics
- Searching the IP ARP Cache

Displaying Global IP Routing Statistics

You can monitor switch performance using either the Web Agent or the CLI.

Web Agent Procedure

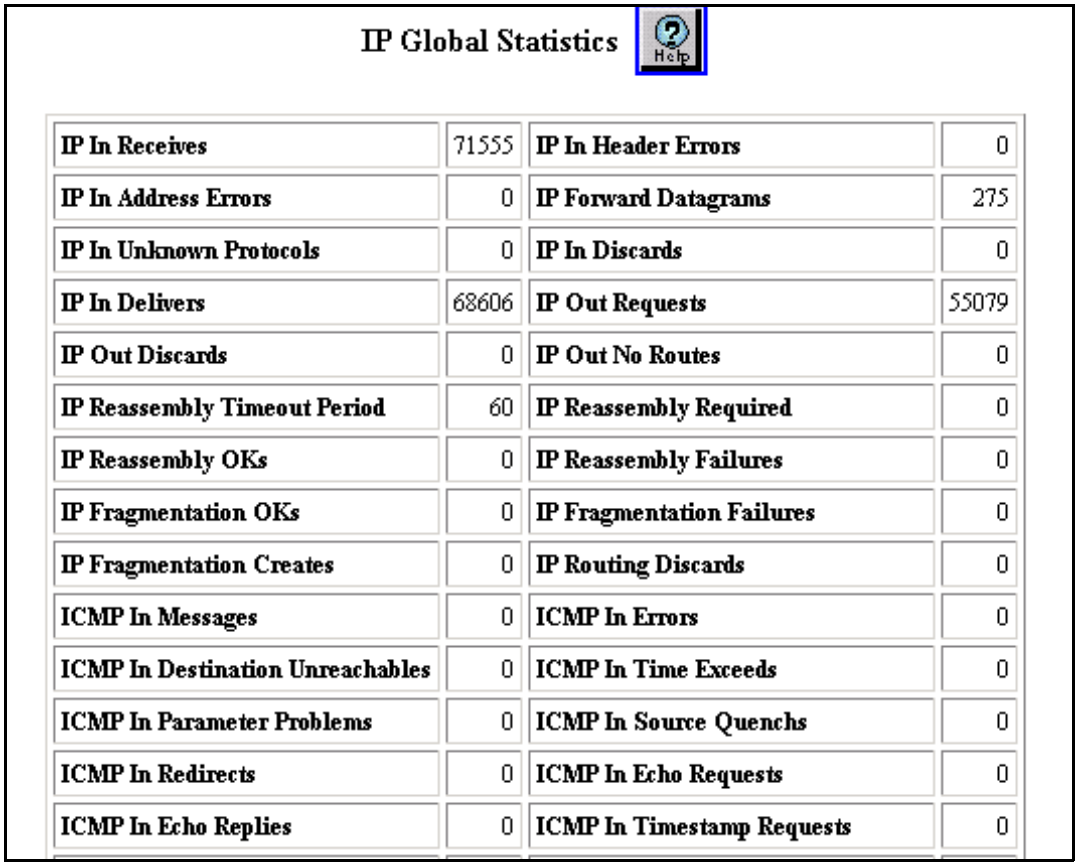
To display the global IP routing statistics using the Web Agent:

*** Note:** IP routing global statistics only represent traffic processed by the supervisor module software.

1. In the navigation pane, expand the **Routing** > **IP** > **Display** folders, and then click **Global Statistics**.

   The IP Routing Global Statistics Web page is displayed in the content pane. See Figure 12-19.

Figure 12-19. IP Routing Global Statistics Web Page



IP Global Statistics

IP In Receives	71555	IP In Header Errors	0
IP In Address Errors	0	IP Forward Datagrams	275
IP In Unknown Protocols	0	IP In Discards	0
IP In Delivers	68606	IP Out Requests	55079
IP Out Discards	0	IP Out No Routes	0
IP Reassembly Timeout Period	60	IP Reassembly Required	0
IP Reassembly OKs	0	IP Reassembly Failures	0
IP Fragmentation OKs	0	IP Fragmentation Failures	0
IP Fragmentation Creates	0	IP Routing Discards	0
ICMP In Messages	0	ICMP In Errors	0
ICMP In Destination Unreachables	0	ICMP In Time Exceeds	0
ICMP In Parameter Problems	0	ICMP In Source Quenchs	0
ICMP In Redirects	0	ICMP In Echo Requests	0
ICMP In Echo Replies	0	ICMP In Timestamp Requests	0

2. Click
   - **Clear** to set all statistics to zero.
   - **Refresh** to update all statistics.
3. See Table 12-11 to review the definition of each statistic:

Table 12-11. IP Routing Global Statistics

Statistic	Defines the...
IP In Receives	Total number of input datagrams received from interfaces, including those received in error.
IP In Header Errors	Number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options.
	1 of 6

Table 12-11. IP Routing Global Statistics

Statistic	Defines the...
IP In Address Errors	Number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For entities that are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
IP Forward Datagrams	Number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP Gateways, this counter will include only those packets which were Source-Routed via this entity, and the Source-Route option processing was successful. **Note:** This is routed by the supervisor in the software.
IP In Unknown Protocols	Number of input datagrams discarded due to errors in their IP headers. Such errors may include bad checksums, version number mismatches, other format errors, time-to-live exceeded, errors discovered in processing their IP options.
IP In Discards	Number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (for example, for lack of buffer space). **Note:** This counter does not include any datagrams discarded while awaiting re-assembly.
IP In Delivers	Total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
IP Out Requests	Total number of IP datagrams that local IP user-protocols (including ICMP) supplied to IP in requests for transmission. **Note:** This counter does not include any datagrams counted in ipForwDatagrams.
	2 of 6

Table 12-11. IP Routing Global Statistics

Statistic	Defines the...
IP Out Discards	Number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but were discarded (for example, for lack of buffer space). Note that this counter includes datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
IP Out No Routes	Number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this 'no-route' criterion. Note that this includes any Datagrams which a host cannot route because all of its default gateways are down.
IP Reassembly Timeout Period	Maximum number of seconds that received fragments are held while they are awaiting reassembly at this entity.
IP Reassembly Required	Number of IP fragments received that need to be reassembled.
IP Reassembly OKs	Number of IP datagrams successfully reassembled.
IP Reassembly Failures	Number of failures detected by the IP re-assembly algorithm (timeout errors). Note that this is not necessarily a count of discarded IP fragments since some algorithms can lose track of the number of fragments by combining them as they are received.
IP Fragmentation OKs	Number of IP datagrams that have been successfully fragmented at this entity.
IP Fragmentation Failures	Number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be.
IP Fragmentation Creates	Number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
IP Routing Discards	Number of routing entries that were chosen to be discarded even though they are valid. One possible reason for discarding such an entry could be to free-up buffer space for other routing entries.
ICMP In Messages	Total number of ICMP messages that the entity received. Note that this counter includes all those counted by icmpInErrors.
	3 of 6

Table 12-11. IP Routing Global Statistics

Statistic	Defines the...
ICMP In Errors	Number of ICMP messages that the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length).
ICMP In Destination Unreachables	Number of ICMP Destination Unreachable messages received.
ICMP In Time Exceeds	Number of ICMP Time Exceeded messages received.
ICMP In Parameter Problems	Number of ICMP Parameter Problem messages received.
ICMP In Source Quenchs	Number of ICMP Source Quench messages received.
ICMP In Redirects	Number of ICMP Redirect messages received.
ICMP In Echo Requests	Number of ICMP Echo (request) messages received.
ICMP In Echo Replies	Number of ICMP Echo Reply messages received.
ICMP In Timestamp Requests	Number of ICMP Timestamp (request) messages received.
ICMP In Timestamp Replies	Number of ICMP Timestamp Reply messages received.
ICMP In Address Mask Requests	Number of ICMP Address Mask Request messages received.
ICMP In Address Mask Replies	Number of ICMP Address Mask Reply messages received.
ICMP Out Messages	Total number of ICMP messages that this entity attempted to send. Note that this counter includes all those counted by icmpOutErrors.
ICMP Out Errors	Number of ICMP messages that this entity did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations, there may be no types of error which contribute to this counter's value.
ICMP Out Destination Unreachables	Number of ICMP Destination Unreachable messages sent.
ICMP Out Time Exceeds	Number of ICMP Time Exceeded messages sent.
ICMP Out Parameter Problems	Number of ICMP Parameter Problem messages sent.
	4 of 6

Table 12-11. IP Routing Global Statistics

Statistic	Defines the...
ICMP Out Source Quenchs	Number of ICMP Source Quench messages sent.
ICMP Out Redirects	Number of ICMP Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects.
ICMP Out Echo Requests	Number of ICMP Echo (request) messages sent.
ICMP Out Echo Replies	Number of ICMP Echo Reply messages sent.
ICMP Out Timestamp Requests	Number of ICMP Timestamp (request) messages sent.
ICMP Out Timestamp Replies	Number of ICMP Timestamp Reply messages sent.
ICMP Out Address Mask Requests	Number of ICMP Address Mask Request messages sent.
ICMP Out Address Mask Replies	Number of ICMP Address Mask Reply messages sent.
UDP In Datagrams	Total number of UDP datagrams delivered to User Datagram Protocol (UDP) users.
UDP In No Ports	Total number of received UDP datagrams for which there was no application at the destination port.
UDP In Errors	Number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
UDP Out Datagrams	Total number of UDP datagrams sent from this entity.
IP Multicast Forward Datagrams	Number of input multicast datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination.
IP Multicast In Discard	Number of input IP multicast datagrams for which no problems were encountered to prevent their continued processing, but were discarded (for example, for lack of buffer space). **Note:** This counter does not include any datagrams discarded while awaiting re-assembly.
IP Multicast In Receives	Total number of input multicast datagrams received from interfaces, including those received in error.
BOOTP/DHCP In Requests	Number of requests received by the BOOTP/DHCP Relay Agent.
	5 of 6

Table 12-11. IP Routing Global Statistics

Statistic	Defines the...
BOOTP/DHCP In Responses	Total number of BOOTP/DHCP response datagrams received by the BOOTP/DHCP Relay Agent.
BOOTP/DHCP In Discards	Number of BOOTP/DHCP requests discarded.Incremented when an IP interface receives a DHCP/BootP request, but the IP interface does not have the BooTP/DHCP Relay Gateway Enabled.
BOOTP/DHCP In Hops Exceeds	Number of BOOTP/DHCP requests not forwarded due to number of hops exceeds.
BOOTP/DHCP Out Requests	Total number of BOOTP/DHCP requests forwarded by the BOOTP/DHCP Relay Agent.
BOOTP/DHCP Out Responses	Total number of BOOTP/DHCP responses forwarded by the BOOTP/DHCP Relay Agent.
	6 of 6

CLI Command

To display the global IP routing statistics using the CLI, enter the following command from the enable mode or configuration mode prompt:

> show ip traffic

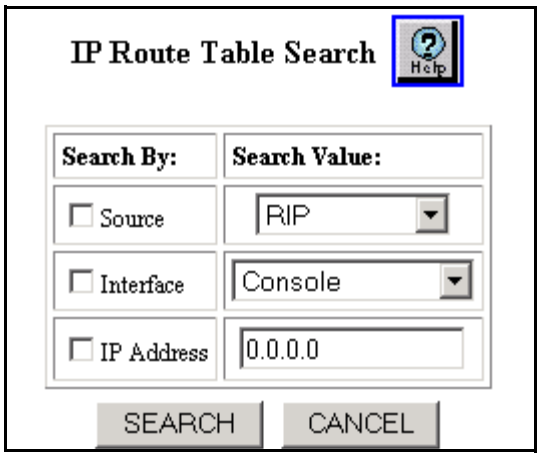
Searching the IP Routing Table

To use the IP routing table:

1. In the navigation pane, expand the **Routing** > **IP** > **Display** folders, and then click **Route Table Search**.

   The IP Route Table Search Web is displayed in the content pane. See Figure 12-20.

Figure 12-20. IP Route Table Search Web Page



2. Select the search criteria you want to use to find more specific information on available routes. For example, if you want to find all static routes that are presently configured on your switch, search by *source* and specify *static* as your search value.

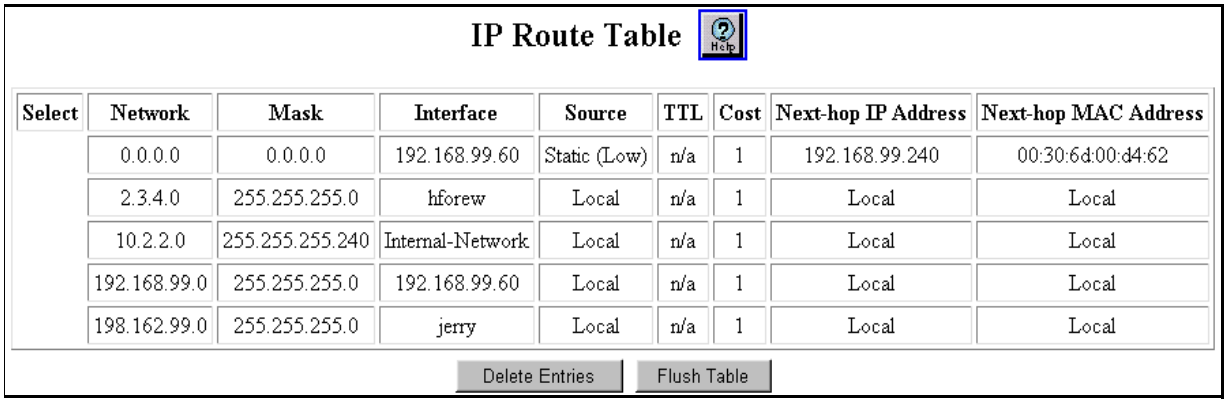
3. See Table 12-12 to determine your search parameters:

Table 12-12. IP Route Table Search Parameters

Parameter	Allows you to search...
Source	Your IP routing table using one of the following parameters: • RIP • OSPF • Static • Local Once you select one of these parameters, the search attempts to find routes associated with that parameter.
Interface	The interface you selected. System default entries include: • **Default** • **Discard** • Ethernet Console • Configured Interface
IP Address	The IP address you entered.

4. Select **SEARCH**. If routes are available, they are displayed in the IP Routing Table Web page. See Figure 12-21.

Figure 12-21. IP Route Table Web Page



*** Note:** To delete a local entry from your IP routing table, you must delete the local IP interface associated with that entry.

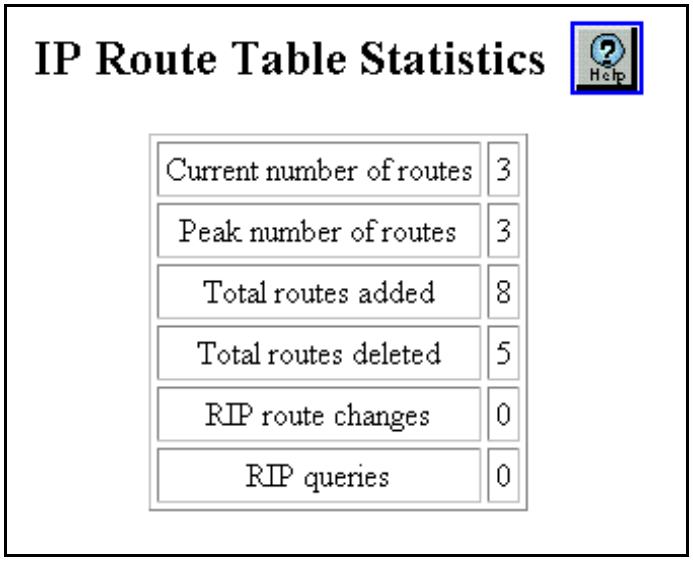
Displaying the IP Routing Table Statistics

To display the IP Routing Table Statistics:

1. In the navigation pane, expand the **Routing > IP > Display** folders, and then click **Route Table Statistics**.

   The IP Routing Table Statistics Web page is displayed in the content pane. See Figure 12-22.

Figure 12-22. IP Route Table Statistics Web Page



See Table 12-13 for a definition of the IP Routing Table Statistics Web page parameters:

Table 12-13. IP Routing Table Statistics Web Page Parameters

Parameter	Definition
Current Number of Routes	Displays the total number of active routes.
Peak Number of Routes	Displays the peak number of routes.
Total Routes Added	Displays the total number of routes added.
Total Routes Deleted	Displays the total number of routes deleted.
RIP Route Changes	Displays the number of changes to the IP route database made by RIP.
RIP Queries	Displays the number of RIP queries sent to the network.

CLI Command

To display the IP routing table statistics, enter the following command from the prompt:

> show ip route summary

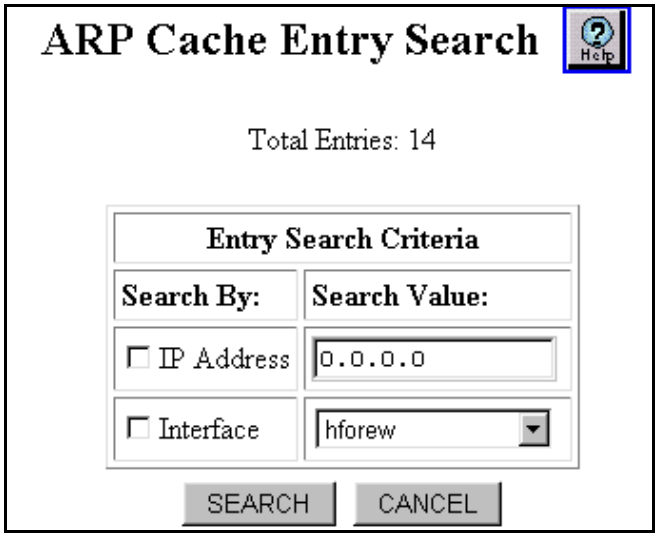
Searching the IP ARP Cache

To search the IP ARP Cache entries:

1. In the navigation pane, expand the **Routing > IP > Display** folders, and then click **ARP Cache Search**.

   The ARP Cache Entry Search Web page is displayed in the content pane. See Figure 12-23.

Figure 12-23. ARP Cache Entry Search Web Page



2. Select the search criteria you want to use to find more specific information on your switch's current ARP cache. For example, if you want to find all of the IP ARP cache entries associated with your out-of-band connection on your switch, search by **VLAN** and specify **Ethernet Console** as your search value.

3. See Table 12-14 for an explanation of the ARP Cache Search Web page parameters:

Table 12-14. ARP Cache Search Web Page Parameters

Parameter	Allows you to do a search based on...
IP Address	The IP address you entered.
Interface	The interface you selected. System entries include all interfaces that you previously configured.

4. Click SEARCH to start the search. If matching entries are found they are displayed in IP ARP Cache Web page.

IP Multicast Statistics

You can view IP Multicast statistics through by using either the Web Agent or the CLI. IP multicast statistics include:

- Displaying IGMP Global Statistics
- Displaying IGMP Interface Statistics
- Displaying the IGMP Group Membership Table
- Displaying the IGMP Local Multicast Forwarding Cache
- Displaying DVMRP Global Statistics
- Displaying DVMRP Interface Statistics
- Displaying DVMRP Neighbor Routers
- Displaying DVMRP Routing Table Statistics
- Displaying the DVMRP Route Table
- Displaying the DVMRP Upstream Routers
- Displaying the DVMRP Designated Forwarder(s) Table
- Displaying the DVMRP Downstream Dependent Routers
- Displaying the DVMRP Multicast Forwarding Cache

*** Note:** It is possible to use **access rules** to filter and prioritize multicast traffic.

Displaying IGMP Global Statistics

Web Agent Procedure

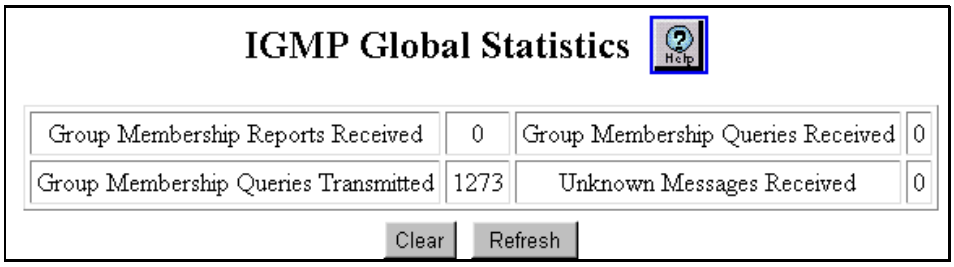
IGMP global statistics provides membership reports, membership queries transmitted and received, and unknown messages.

To display IGMP global statistics from the Web Agent:

1. In the navigation pane, expand the **Routing** > **IGMP** folders, and then click **Global Statistics**.

   The IGMP Global Statistics Web page is displayed in the content pane. See Figure 12-24.

Figure 12-24. IGMP Global Statistics Web Page



2. To modify your global statistics, perform one of the following:
   - Click **CLEAR** to reset all statistics to zero
   - Click **REFRESH** to view the latest statistics.
3. See Table 12-15 for an explanation of the IGMP Global Statistics Web page parameters:

Table 12-15. IGMP Global Statistics Web Page Parameters

Parameter	Defines the...
Group Membership Reports Received	Number of reports received in response to a group membership query. Hosts respond to a Query by generating Host Membership Reports reporting each host group to which they belong on the network interface from which the Query was received.
Group Membership Queries Transmitted	Number of query messages sent by all local IGMP interfaces. These messages are sent to discover which host groups have members on their attached local networks. Queries are addressed to the all-hosts group (address 224.0.0.1), and carry an IP time-to-live of 1.
Group Membership Queries Received	Number of query messages received by all local IGMP interfaces.
Unknown Messages Received	IGMP message of a type *other* than Group Membership Query, Group Membership Report, or Leave group.

Displaying IGMP Interface Statistics

You can view IGMP interface statistics for each configured IP interface that has multicast protocol enabled.

You can display IGMP interface statistics using either the Web Agent or the CLI.

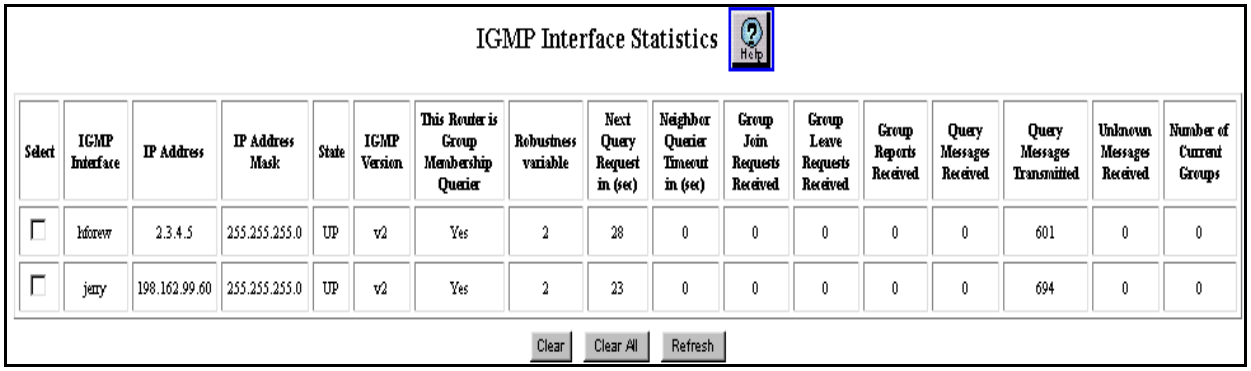
Web Agent Procedure

To display the IGMP Interface Statistics using the Web Agent:

1. In the navigation pane, expand the **Routing > IGMP** folders, and then click **Interface Statistics**.

   The IGMP Interface Statistics Web page is displayed in the content pane. See Figure 12-25.

Figure 12-25. IGMP Interface Statistics Web Page



2. To modify the interface statistics, perform one of the following:

   — Select one or more interfaces, and then *c*lick **CLEAR** to reset statistics on those interfaces to zero.

   — Click **CLEAR ALL** to reset all statistics to zero.

   — *C*lick **REFRESH** to view the latest interface statistics.

3. See Table 12-16 for an explanation of the IGMP Interface Statistics Web page parameters:

Table 12-16. IGMP Interface Statistics Web Page Parameters

Parameter	Defines the...
IGMP Interface	IP interface for these statistics.
IP Address	IP address associated with the interface.
IP Address Mask	Subnet mask associated with each listed interface.
State	Current state of the interface. For example, if the interface is enabled and operating properly, UP is displayed.
IGMP Version	Version of IGMP enabled on each interface.The default version is V2
	1 of 2

Table 12-16. IGMP Interface Statistics Web Page Parameters

Parameter	Defines the...
This Router is Group Membership Querier	Router that was configured or elected to be the designated **group membership querier**. The switch queries hosts on each interface only when it is acting as the designated querier on that interface.
Robustness Variable	Setting for the expected packet loss on a subnet. If a subnet is expected to have more packet loss, the Robustness Variable should be increased. The Robustness Variable must not be set to 0 and should not be set to 1. The default value is 2
Next Query Request (sec)	Remaining amount of time (in seconds) before the next group membership query is transmitted.
Neighbor Querier Timeout (sec)	Amount of time (in seconds) remaining before this interface assumes the role of designated querier. This timer is reset to the value entered for the Neighbor Querier Timeout Interval each time an IGMP query is received from a neighbor with a lower IP address. If no response is received in the allowed time, the switch will become the designated querier on this interface. Applicable only if IGMP V2 is used.
Group Join Requests Received	Number of new groups on this interface.
Group Leave Requests Received	Number of leave requests received on this interface.
Group Reports Received	Number of reports received on this interface in response to a group membership query. Hosts respond to a Query by generating Host Membership Reports reporting each host group to which they belong on the network interface from which the Query was received.
Query Messages Received	Number of query messages received from other multicast routers.
Query Messages Transmitted	Number of query messages sent by a multicast router. These messages are sent to discover which host groups have members on their attached local networks. Queries are addressed to the all-hosts group (address 224.0.0.1), and carry an IP time-to-live of 1.
Unknown Messages Received	IGMP messages received with an unsupported type.
Number of Current Groups	Number of groups on each interface for which there are entries in the Group Membership Table.
	2 of 2

CLI Command

To display the IGMP Interface Statistics using the CLI, enter the following command from configuration mode:

```
(configure)# show ip igmp statistics
```

Displaying the IGMP Group Membership Table

The multicast group table provides information on interfaces that are members of an IGMP group and contains an expiry time for the entry, IP address of the group, and the group reporter address.

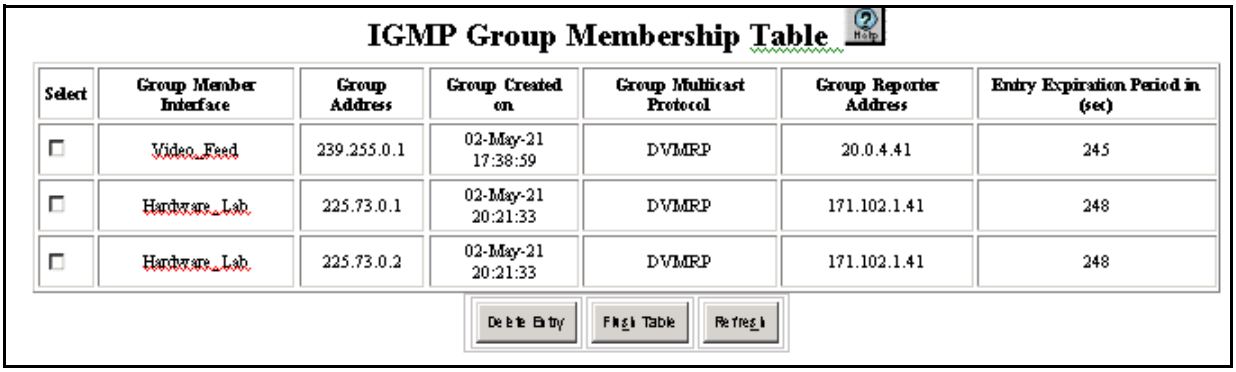
Web Agent Procedure

To display the multicast group table using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IGMP** folders, and then click **Group Membership Table**.

   The IGMP Group Membership Table Web page is displayed in the content pane. See Figure 12-26.

Figure 12-26. IGMP Group Membership Table Web Page



2. To modify the **IGMP Group Membership Table**, perform one of the following steps:

   — Select the entry and click **Delete Entry** to delete one or more entries.

   — Click **Flush Table** to clear the entire table.

   — Click **REFRESH** to receive the most up-to-date information on the entries in the table.

3. See Table 12-17 for an definition of the IGMP Group Membership Table Web page parameters:

Table 12-17. IGMP Group Membership Table Web Page Fields

Parameter	Defines the...
Group Member Interface	Interface that is connected to a member of an IGMP group.
Group Address	Group address that has members on this interface.
Group Created On	Time at which the group was created on the router.
Group Multicast Protocol	Routing protocol being used for the group. If no routing protocol is being used on the interface a group is on, this column displays IGMP.
Group Reporter Address	IP address of the host that sent the most recent host membership report for this group.
Entry Expiration Period in (sec)	Expiration time (in seconds) of the group that is being displayed.

CLI Command

To display the multicast group table using the CLI, enter the following command from configuration mode:

```
(configure)# show ip igmp groups
```

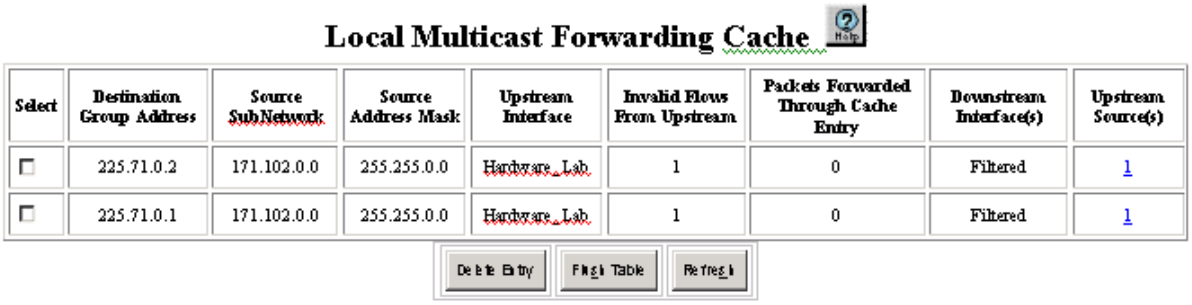
Displaying the IGMP Local Multicast Forwarding Cache

To display the multicast forwarding cache information (IGMP only interfaces):

1. In the navigation pane, expand the **Routing** > **IGMP** folders, and then click **Local Multicast Forwarding Cache**.

   The Local Multicast Forwarding Cache Web page is displayed in the content pane. See Figure 12-27.

Figure 12-27. Local Multicast Forwarding Cache Web Page



2. To modify the Local Multicast Forwarding Cache, select an entry and:
   - Click **Delete Entry** to delete one or more entries.
   - Click **Flush Table** to clear the entire table.
   - Click **REFRESH** to receive the most up-to-date information on the entries in the table.
3. See Table 12-18 to review the IGMP Multicast Forwarding Cache information:

Table 12-18. IGMP Local Multicast Forwarding Cache Parameters

Parameters	Defines the...
Destination Group Address	Destination group address of the multicast transmission.
Source SubNetwork	Subnet on which the IGMP interface(s) exist.
Source Address Mask	Subnet mask associated with the IGMP source subnetwork.
Upstream Interface	IP interface configured on the upstream interface.
Invalid Flows From Upstream	Number of invalid flows received from the upstream neighbor.
Packets Forwarded Through Cache Entry	Number of packets successfully forwarded in the CPU (supervisor module).
Downstream Interface(s)	Number of downstream interfaces and provides a link to the IGMP Downstream Interfaces Web page.
Upstream Source(s)	Number of upstream interfaces and provides a link to the IGMP Upstream Interfaces Web page.

Displaying DVMRP Global Statistics

You can view the DVMRP Global Statistics from by using either the Web Agent or the CLI.

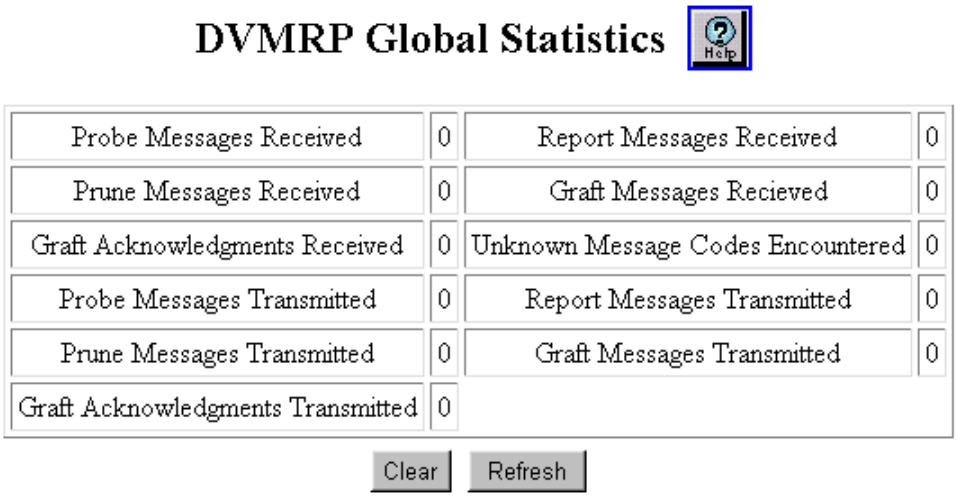
Web Agent Procedure

To display the DVMRP global statistics:

1. In the navigation pane, expand the **Routing** > **DVMRP** folders, and then click **Global Statistics**.

   The DVMRP Global Statistics Web page is displayed in the content pane. See Figure 12-28.

Figure 12-28. DVMRP Global Statistics Web Page



2. Click:
   - **CLEAR** to reset all statistics to zero
   - **REFRESH** to view the latest statistics.
3. See Table 12-19 for an explanation of the DVMRP Global Statistics Web page parameters:

Table 12-19. DVMRP Global Statistics Web Page Parameters

Statistic	Defines the number of...
Probe Messages Received	Probe messages received on this switch. DVMRP routers exchange probes and routing updates so they each have a picture of their neighbors' capabilities and the DVMRP network topology.
Report Messages Received	Route Report messages received on this switch.
Prune Messages Received	Prune messages received on this switch. This indicates the number of old branches removed from the multicast distribution tree.
Graft Messages Received	Graft messages received. This indicates the number of attempts at adding a new branch to the multicast distribution tree.
Graft Acknowledgments Received	Graft acknowledgments received.
Unknown Message Codes Encountered	Messages received that are not graft, report, or probe messages.
	1 of 2

Table 12-19. DVMRP Global Statistics Web Page Parameters

Statistic	Defines the number of...
Probe Messages Transmitted	Probe messages transmitted to the network.
Report Messages Transmitted	Report messages transmitted on this switch.
Prune Messages Transmitted	Prune messages transmitted upstream on this switch. This indicates the number of old branches removed from the multicast distribution tree.
Graft Messages Transmitted	Graft messages transmitted upstream from this switch. This indicates the number of new upstream branches added to the multicast distribution tree.
Graft Acknowledgments Transmitted	Graft acknowledgments sent downstream from this switch. This indicates the number of new downstream branches added to the multicast distribution tree.
	2 of 2

Displaying DVMRP Interface Statistics

DVMRP interface statistics list active DVMRP interfaces and provide specific information on each interface.

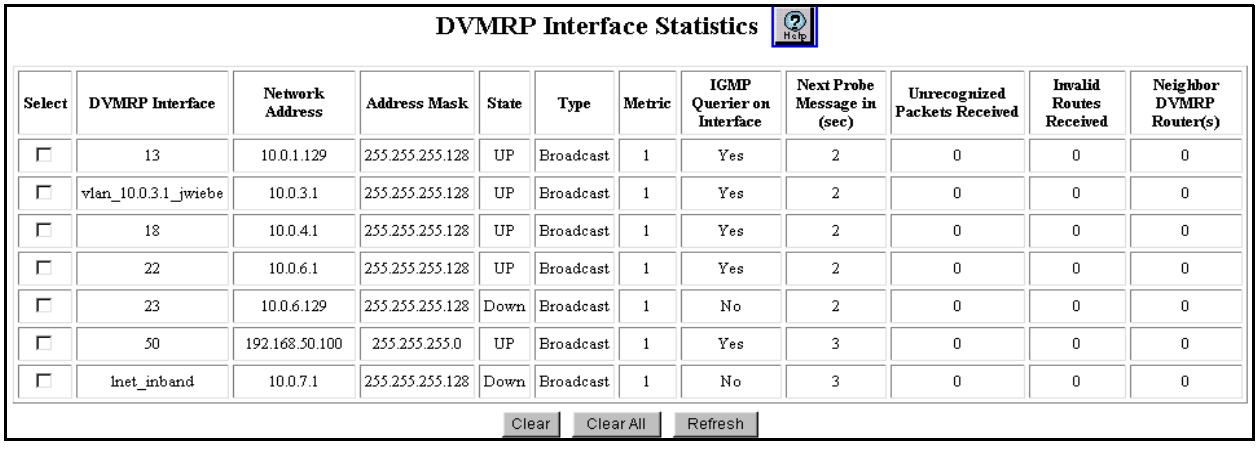
Web Agent Procedure

To display DVMRP interface statistics using the Web Agent:

1. In the navigation pane, expand the **Routing > DVMRP** folders, and then click **Interface Statistics**.

   The DVMRP Interface Statistics Web page is displayed in the content pane. See Figure 12-29.

Figure 12-29. DVMRP Interface Statistics Web Page



2. Select an interface and perform one of the following:
   - Click **REFRESH** to view the latest interface statistics.
   - Click **CLEAR All** to reset all statistics to zero.
   - Click **CLEAR** to reset the selected interface.
3. See Table 12-20 for an explanation of the DVMRP Interface Statistics Web page parameters:

Table 12-20. DVMRP Interface Statistics Parameters

Parameter	Defines the...
DVMRP Interface	IP interface configured with the DVMRP multicast routing protocol.
Network Address	IP address of the interface configured with the DVMRP multicast.
Address Mask	IP subnet mask associated with the interface.
State	Current status of the interface. Possible status indications include: • **UP** - The interface is active. • **DOWN** - The interface is inactive.
Type	Type of interface configured. Possible values include: • **Broadcast** - All traffic is forwarded through the routers. This is not a tunnel and does not require encapsulation. • **IPIP Tunnel** - All multicast traffic (data and protocol messages) on this interface is encapsulated in IP unicast packets with the protocol set to IPIP (IP in IP). • **Non-Encapsulated Tunnel** - All multicast data traffic on this interface is IPIP encapsulated, but the protocol messages are simple unicast.
Metric	Interface cost.
IGMP Querier on Interface	Router is the IGMP querier. The purpose of the IGMP querier is to periodically poll hosts on your network to trigger group membership reports.
Next Probe Message in (sec)	Time (in seconds) remaining until the next probe message is sent.
Unrecognized Packets Received	Number of unknown DVMRP messages.
	1 of 2

Table 12-20. DVMRP Interface Statistics Parameters

Parameter	Defines the...
Invalid Routes Received	Number of invalid routes received on this interface.
Neighbor DVMRP Router(s)	Number of (neighbor) routers that are also running DVMRP. **Note:** This number is a hypertext link that provides additional information on the DVMRP neighbor router(s).
	2 of 2

CLI Command

To display the DVMRP interface statistics using the CLI, enter the following command from configuration mode:

```
(configure)# show ip dvmrp interface
```

Displaying DVMRP Neighbor Routers

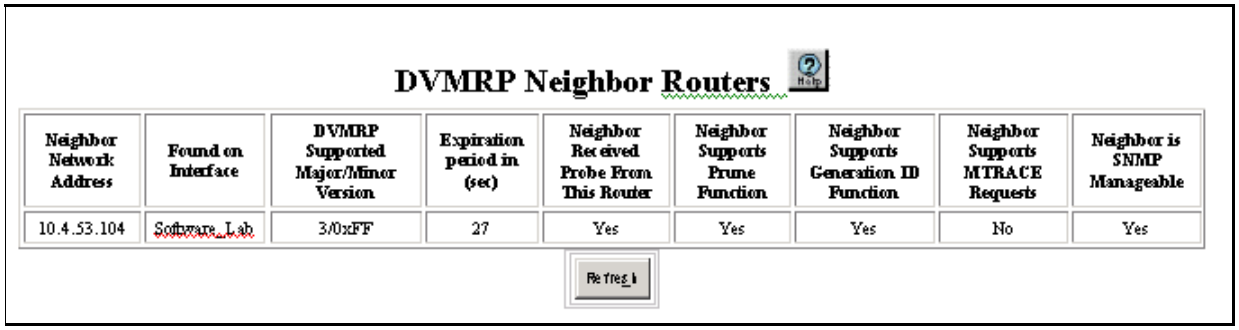
Web Agent Procedure

To view the DVMRP neighbor routers using the Web Agent:

1. In the navigation pane, expand the **Routing** > **DVMRP** folders, and then click **Interface Statistics**.

   The DVMRP Interface Statistics Web page is displayed in the content pane. See Figure 12-29.

Figure 12-30. DVMRP Neighbor Routes Web Page



2. Select the number in the **Neighbor DVMRP Router(s)** column, if it is 1 or more. The DVMRP neighbor routers Web page is displayed.

3. See Table 12-21 to view more information on DVMRP neighbor routers.

Table 12-21. DVMRP Neighbor Routers

Parameter	Displays...
Neighbor Network Address	The neighbor router's IP address.
Found on Interface	The neighbor routers found on this interface.
DVMRP Supported Major/Minor Version	The DVMRP version supported by the neighbor router.
Expiration period in (sec)	The time (in seconds) before the neighbor router times out.
Neighbor Received Probe From This Router	Whether the neighbor router received this router's probe message.
Neighbor Supports Prune Function	Whether the neighbor router supports prune functionality.
Neighbor Supports Generation ID Function	Whether the neighbor router supports generation of IDs.
Neighbor Supports MTRACE Requests	Whether the neighbor router supports MTRACE requests.
Neighbor is SNMP Manageable	Whether the neighbor router can be managed by SNMP.

CLI Command

To display the DVMRP neighbor routers using the CLI, enter the following command from configuration mode:

```
(configure)# show ip dvmrp interface neighbors
```

Displaying DVMRP Routing Table Statistics

DVMRP routing table statistics provides information on the current number of valid routes, the number of total routes present (both valid and invalid), and the number of triggered routes.

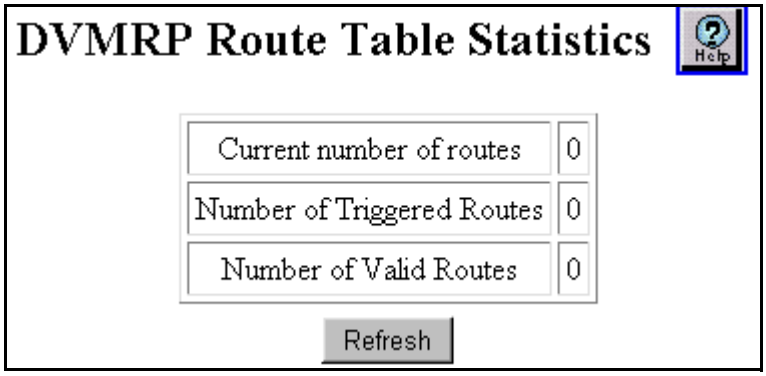
Web Agent Procedure

To view the DVMRP routing table statistics using the Web Agent:

1. In the navigation pane, expand the **Routing** > **DVMRP** folders, and then click **Route Table Statistics**.

   The DVMRP Routing Table Statistics Web page is displayed in the content pane. See Figure 12-31.

Figure 12-31. DVMRP Route Table Statistics Web Page



2. See Table 12-22 for an explanation of the DVMRP Route Table Statistics Web page parameters:

Table 12-22. DVMRP Routing Statistical Parameters

Parameter	Defines the...
Current Number of Routes	Total number of routes present in the routing database. This number includes both valid and invalid routes.
Number of Triggered Routes	Total number of routes added to the routing table that were triggered by a topology change in the network.
Number of Valid Routes	Total number of *valid routes* present in the routing database.

Displaying the DVMRP Route Table

The DVMRP route table contains information on valid DVMRP routes, the expiry for those routes, and additional next-hop information.

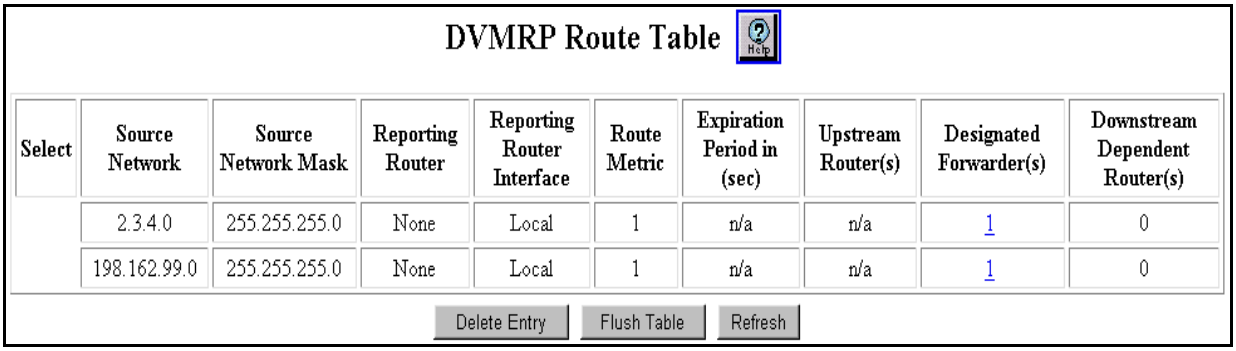
Web Agent Procedure

To view the DVMRP route table using the Web Agent:

1. In the navigation pane, expand the **Routing > DVMRP** folders, and then click **Route Table**.

The DVMRP Route Table Web page is displayed in the content pane. See Figure 12-32.

Figure 12-32. DVMRP Route Table Web Page



2. To modify your DVMRP Route table, do one of the following:

   — To delete one or more entries., select the entry and click **Delete Entry**

   — To clear the entire table, click **Flush Table**.

   — To receive the most up-to-date information on the entries in the table, Click **REFRESH**.

3. See Table 12-23 for an explanation of the **DVMRP Route Table** parameters:

Table 12-23. DVMRP Route Table Parameters

Parameter	Defines the...
Source Network	Network from which a multicast flow may originate.
Source Network Mask	Source network mask.
Reporting Router	IP address of the router reporting this route to its neighbors.
Reporting Router Interface	IP interface configured, which leads to the upstream neighbor (DVMRP router).
Route Metric	Router's cost to source network.
Expiration Period in (sec)	Time (in seconds) remaining before the source network is removed from the DVMRP routing table.
Upstream Router(s)	IP address of the DVMRP router that is the upstream neighbor to the local router. The local DVMRP router must know which DVMRP router is its upstream neighbor to determine how packets from a given source will be transmitted to a given multicast group. Displays the DVMRP Upstream Router(s) Web page.
	1 of 2

Table 12-23. DVMRP Route Table Parameters

Parameter	Defines the...
Designated Forwarder(s)	Network router(s) responsible for forwarding from the source network onto the downstream interface. Displays the Designated Forwarders Web page.
Downstream Dependent Router(s)	Number of downstream DVMRP routers that are dependent on this router for this particular route. Displays the DVMRP Downstream Dependent Router(s) Web page.
	2 of 2

CLI Command

To display the DVMRP route table using the CLI, enter the following command from configuration mode:

```
(configure)# show ip dvmrp routes
```

Displaying the DVMRP Upstream Routers

Web Agent Procedure

To view the DVMRP upstream routers:

1. In the navigation pane, expand the **Routing > DVMRP** folders, and then click **Route Table.**

   The DVMRP Route Table Web page is displayed in the content pane. See Figure 12-32.

2. Select the number in the **Upstream Router(s)** column to view information on the upstream routers. The Upstream Router(s) Web page is displayed.

3. See Table 12-24 for an explanation of the DVMRP Upstream Router(s) Web page parameters.

Table 12-24. DVMRP Upstream Router(s) Web Page Parameters

Parameter	Definition
Router Network Address	Displays the router network address.
Router Cost to Source Network	Displays the cost metric.
Found on Interface	Displays the interface on which the upstream router was found.

Displaying the DVMRP Designated Forwarder(s) Table

Web Agent Procedure

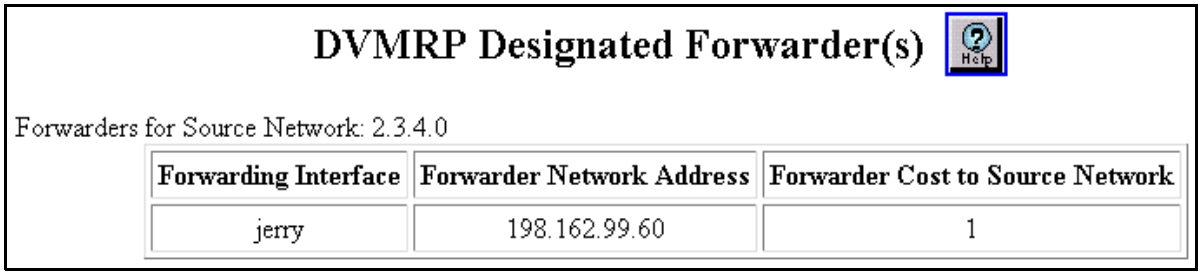
To view the DVMRP Designated Forwarder table:

1. In the navigation pane, expand the **Routing > DVMRP** folders, and then click **Route Table**.

   The DVMRP Route Table Web page is displayed in the content pane. See Figure 12-32.

2. Select the number from the **Designated Forwarder** column for the appropriate source network. The Designated Forwarder(s) Table Web page is displayed. See Figure 12-33.

Figure 12-33. Designated Forwarder(s) Table Web Page



3. See Table 12-25 for an explanation of the DVMRP Designated Forwarder(s) Web page parameters:

Table 12-25. DVMRP Designated Forwarder(s) Table Web Page Parameters

Parameter	Defines the...
Forwarding Interface	Local interface which leads to the network where the Designated Forwarder resides.
Forwarder Network Address	Designated Forwarder for the given source network on the indicated Forwarding Interface.
Forwarder Cost to Source Network	Cost reported by the Designated Forwarder for the given source network.

Displaying the DVMRP Downstream Dependent Routers

Web Agent Procedure

To view the DVMRP downstream dependent routers:

1. In the navigation pane, expand the **Routing** > **DVMRP** folders, and then click **Route Table**.

   The DVMRP Route Table Web page is displayed in the content pane. See Figure 12-32.

2. Select the number from the **Downstream Dependent Router(s)** column for the appropriate source network. The Downstream Dependent Router(s) Web page is displayed.

3. See Table 12-26 for an explanation of the DVMRP Downstream Dependent Router(s) Web page parameters.

Table 12-26. DVMRP Downstream Dependent Router(s) Web Page Parameters

Parameter	Definition
Router Network Address	Displays the router network address of the downstream dependent router.
Found on Interface	Displays the name of the interface on which the downstream router was found.
DVMRP Supported Major/ Minor Version	Displays the DVMRP version supported.
Router Received Probe from This Router	Displays whether the router received a probe from this router.
Router Supports Prune Function	Displays whether this router supports prune functionality.
Router Supports Generation ID Function	Displays whether the router supports generation ID function.
Router is SNMP Manageable	Displays whether the router can be managed by SNMP.

Displaying the DVMRP Multicast Forwarding Cache

The DVMRP Multicast Forwarding Cache Web page provides detailed information on the multicast forwarding attributes including information on downstream interfaces and upstream sources. DVMRP allows the switch to construct paths from the hosts that are sending to a multicast group to the hosts that are receiving it.

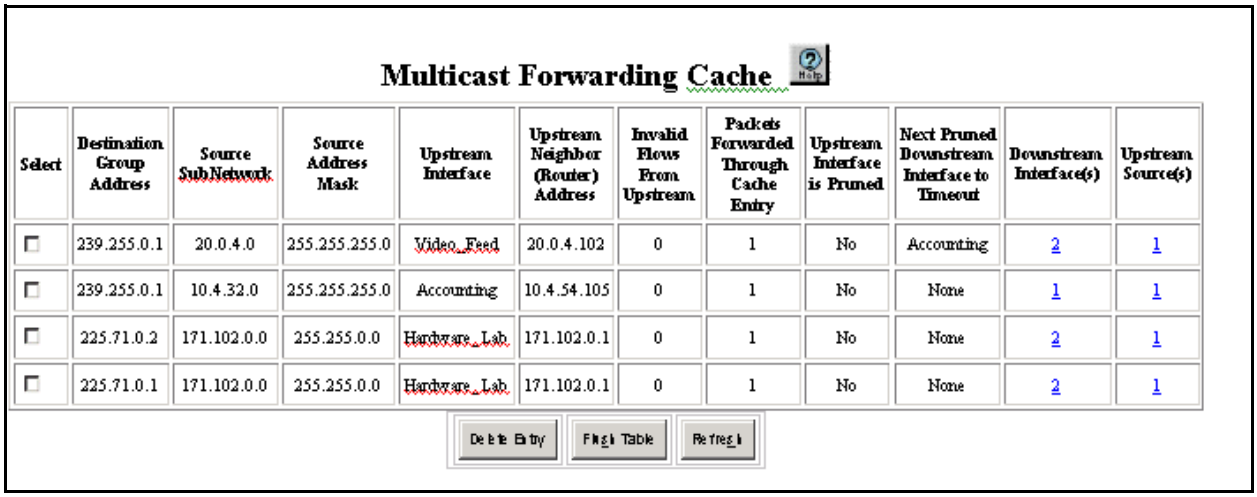
Web Agent Procedure

To display the multicast forwarding cache Web page:

1. In the navigation pane, expand the **Routing > DVMRP** folders, and then click **Multicast Forwarding Cache**.

   The Multicast Forwarding Cache Web page is displayed in the content pane. See Figure 12-34.

Figure 12-34. Multicast Forwarding Cache Web Page



2. To modify the Multicast Forwarding Table:

   — Select the entry and click **Delete Entry** to delete one or more entries.

   — Click **Flush Table** to clear the entire table.

   — Click **REFRESH** to receive the most up-to-date information on the entries in the table.

3. See Table 12-27 for an explanation of the Multicast Forwarding Cache Web page parameters:

Table 12-27. Multicast Forwarding Cache Web Page Parameters

Parameter	Defines the...
Select	Selection of the multicast forwarding cache.
Destination Group Address	Destination group address of the multicast transmission.
Source SubNetwork	Subnet from which the flow is coming.
	1 of 2

Table 12-27. Multicast Forwarding Cache Web Page Parameters

Parameter	Defines the...
Source Address Mask	Subnet mask associated with the DVMRP source subnetwork.
Upstream Interface	Local interface which is receiving this flow.
Upstream Neighbor (Router) Address	IP address of the upstream neighbor (router).
Invalid Flows From Upstream	Number of invalid flows received from the upstream neighbor.
Packets Forwarded Through Cache Entry	Number of packets successfully forwarded in the CPU (supervisor module) for this flow.
Upstream Interface is Pruned	Router that is sending prunes to the upstream neighbor. Allows you to open the DVMRP Upstream Prune Information Web page.
Next Pruned Downstream Interface to Timeout	Next interface that is currently pruned which will be grafted back.
Downstream Interface(s)	Number of downstream interfaces. Allows you to open the DVMRP Downstream Links Web page.
Upstream Source(s)	Number of upstream interfaces. Allows you to open the Upstream Sources Web page.
	2 of 2

4. Select the number in the **Upstream Interface is Pruned** field. The Upstream Prune Information Web page is displayed.

5. See Table 12-28 for an explanation of the Upstream Prune Information Web page parameters.

Table 12-28. Upstream Prune Information Web Page Parameter

Parameter	Displays the...
Destination Group Address	Destination group address of the multicast session.
Source SubNetwork	Subnet on which the DVMRP interface exists.
DVMRP Upstream Interface	Name of the upstream interface.
	1 of 2

Table 12-28. Upstream Prune Information Web Page Parameter

Parameter	Displays the...
Interface Type	The interface type. Types include: • **Broadcast** - All traffic is forwarded through the routers. This is not a tunnel and does not require encapsulation. • **IPIP Tunnel** - All multicast traffic (data and protocol messages) on this interface is encapsulated in IP unicast packets with the protocol set to IPIP (IP in IP). • **Non-Encapsulated Tunnel** - All multicast data traffic on this interface is IPIP encapsulated, but the protocol messages are simple unicast.
Interface is Pruned	Status of whether the interface has been pruned.
Prune Expiration Time in sec	Time (in seconds) that the interface times out waiting for the prune message to expire.
	2 of 2

6. Select the number in the **Downstream Interface(s)** field. The DVMRP Downstream Link(s) Web page is displayed.

7. See Table 12-29 for an explanation of the DVMRP Downstream Links Web page parameters.

Table 12-29. DVMRP Downstream Links Web Page Parameters

Parameter	Displays the...
Destination Group Address	Destination group address.
Source SubNetwork	Source subnetwork.
DVMRP Downstream Interface	DVMRP Downstream interface.
	1 of 2

Table 12-29. DVMRP Downstream Links Web Page Parameters

Parameter	Displays the...
Interface Type	Interface type. Types include: • **Broadcast** - All traffic is forwarded through the routers. This is not a tunnel and does not require encapsulation. • **IPIP Tunnel** - All multicast traffic (data and protocol messages) on this interface is encapsulated in IP unicast packets with the protocol set to IPIP (IP in IP). • **Non-Encapsulated Tunnel** - All multicast data traffic on this interface is IPIP encapsulated, but the protocol messages are simple unicast.
Interface is Pruned	Status of whether the interface has been pruned.
Prune Expiration in (sec)	Time (in seconds) that the interface times out waiting for the prune message to expire.
	2 of 2

8. Select the number in the **Upstream Source(s)** field. The DVMRP Upstream Source(s) Web page is displayed.

9. See Table 12-30 for an explanation of the DVMRP Upstream Source(s) Web page parameters.

Table 12-30. DVMRP Upstream Source(s) Web Page Parameters

Parameter	Displays the...
Destination Group Address	Destination group address for the upstream interface.
Flow Source Address	Host source address for the upstream flow.
Flow Upstream Interface	Name of the flow source interface.
Payload Protocol Type	Protocol type for the payload.
Source Port Number	Source port number.
Destination Port Number	Destination port number.

Configuring VRRP

Virtual Router Redundancy Protocol (VRRP) is used to provide fast-fail over for hosts if the default gateway fails. This eliminates the single point of failure inherent in a network with statically configured default routes. The VRRP protocol defines an election process that will determine a Master and a Backup router. The Master router will forward all packets destined for the IP Address associated with the virtual router. The Backup router monitors the availability of the Master router and will assume Mastership in the event that the Master router fails.

The VRRP protocol is described in detail in RFC 2338.

This section contains procedures for the following tasks:

- Globally Enabling VRRP
- Enabling VRRP on an Interface
- Creating a VRRP Virtual Router
- Displaying VRRP Statistics
- VRRP Configuration Considerations

Globally Enabling VRRP

VRRP can be enabled (or disabled) globally using the Web or CLI. Disabling VRRP globally will cause all virtual routers to transition to the Initialize state. VRRP is globally enabled by default.

You can globally enable VRRP from either the Web Agent or the CLI.

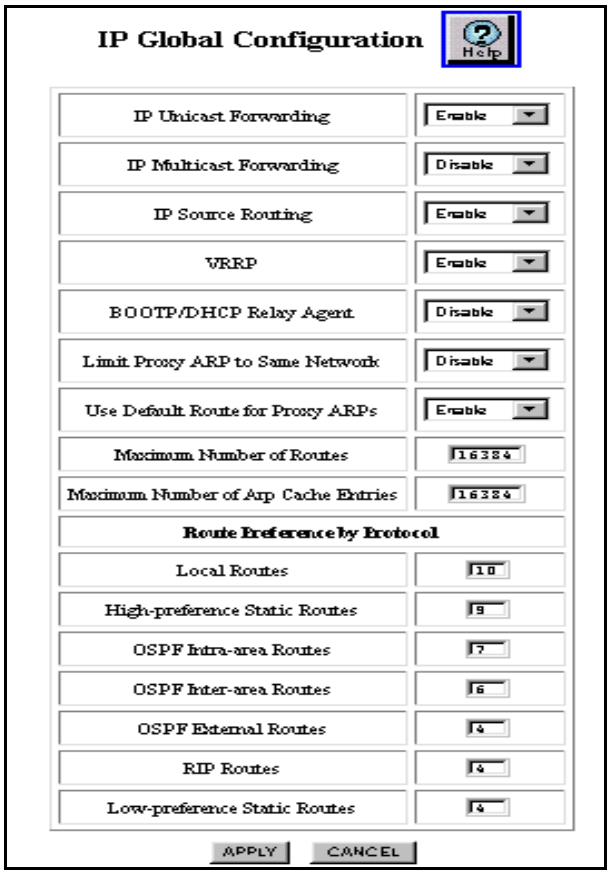
Web Agent Procedure

To globally enable VRRP using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Global Configuration**.

   The IP Global Configuration Web page is displayed in the content pane. See Figure 12-35.

Figure 12-35. IP Global Configuration Web Page



2. Select **Enable** from the **VRRP** field pull-down menu.

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To globally enable VRRP from the CLI, enter the following command in Configure mode:

```
(configure)# router vrrp
```

Enabling VRRP on an Interface

VRRP can be enabled on an interface using the Web Agent or the CLI.

*** Note:** For VRRP to function correctly, it must be enabled globally and enabled on the interfaces on which any virtual routers are to be created. Otherwise, any virtual routers created will be in the Initialize state (see Table 12-32).

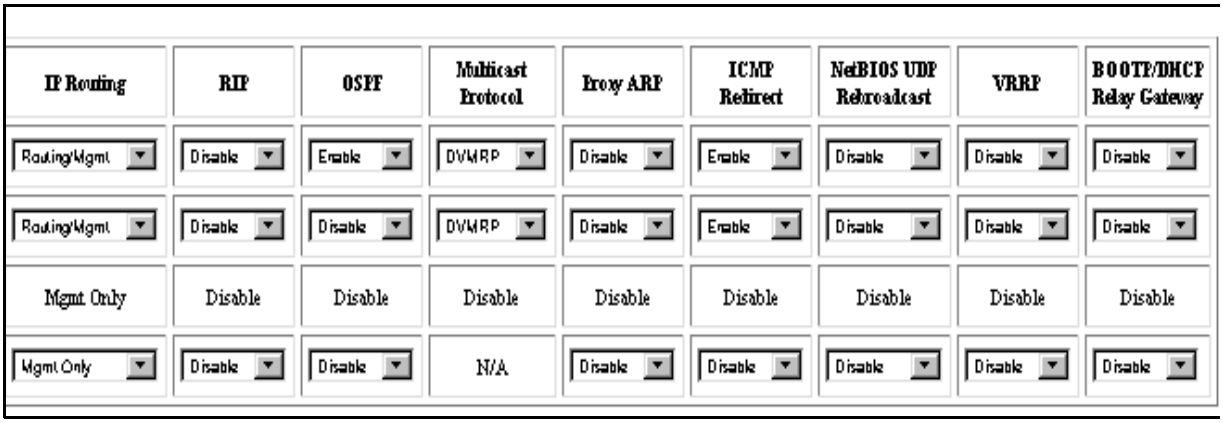
Web Agent Procedure

To enable an VRRP on an interface using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Interfaces**.

   The IP Interfaces Web page is displayed in the content pane. See Figure 12-36.

Figure 12-36. IP Interface Web Page



2. Select the interface on which you want VRRP to be enabled.
3. Select **Enable** from the **VRRP** field pull-down menu for the interface you selected.
4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To enable an VRRP on an interface using the CLI, enter the following command in Configure/Interface mode:

`(config-if:<if name>)#` **ip vrrp**

Creating a VRRP Virtual Router

You can create a VRRP router using either the Web Agent or the CLI.

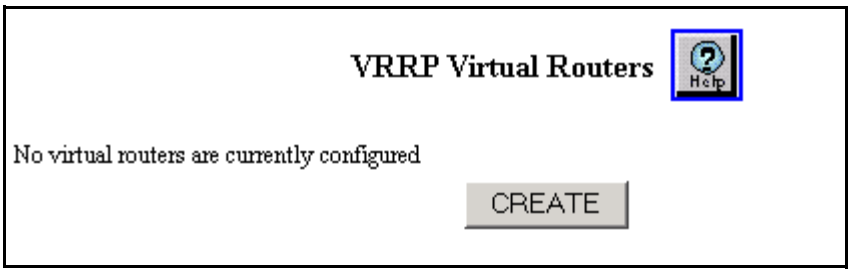
Web Agent Procedure

To create a VRRP router from the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **VRRP**.

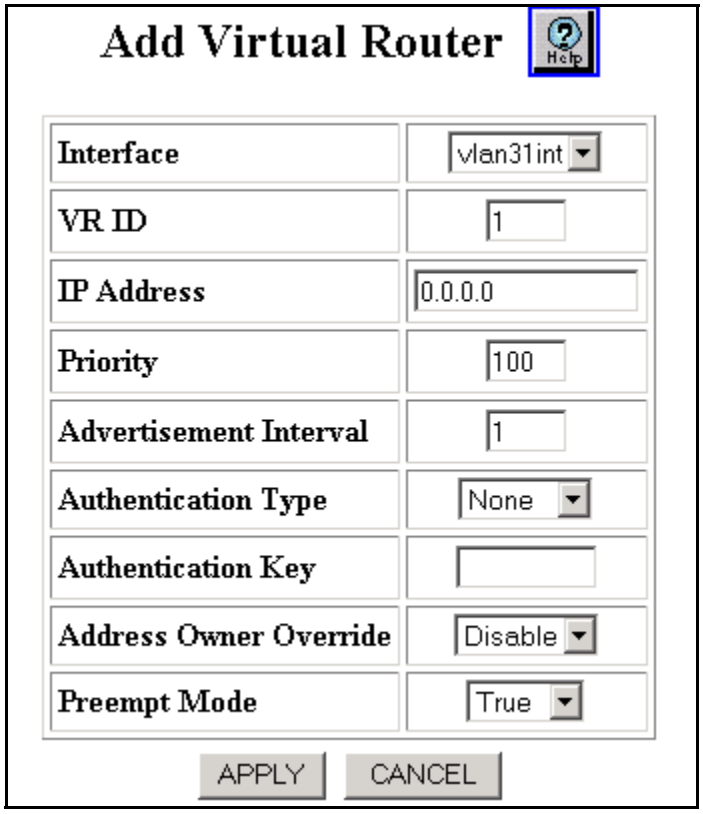
   The VRRP Virtual Routers Web page is displayed in the content pane. See Figure 12-37.

Figure 12-37. VRRP Virtual Routers Web Page



2. Select **CREATE**. The Add VRRP Virtual Router Web page opens (Figure 12-38).

Figure 12-38. Add VRRP Virtual Router Web Page



3. Table 12-31 lists the parameters and describes the type of information that should be entered in the Add VRRP Virtual Router Web page fields to create a VRRP virtual router.

Table 12-31. Add VRRP Virtual Router Web Page Parameters

Parameter	Definition
Interface	This is the IP Interface that the virtual router will be associated with. The drop-down menu displays all the IP Interfaces currently created. Select the Interface that the virtual router will be associated with. Be sure to enable VRRP on the particular Interface selected. Multiple virtual routers may be created on a single interface but they must have unique IP Addresses and VRID's. The default setting is the top interface in the Routing > IP > Configuration > Interface Web page.
VR ID	This is the virtual router identification number. The range is 1 - 255. The same VR ID can be used for multiple virtual routers as long as the associated IP Interfaces are on different VLAN's (each VR ID number can be used once per VLAN). The default value is 1. The range is 1-255 (decimal).
IP Address	Enter the IP Address for the virtual router. This is known as the Virtual IP Address (VIP). The VIP can be the same as the IP Address of the associated Interface. This is known as IP Address Owner. The VIP can also be a unique address. This is called non-Address Owner. IP Address Owner and non-Address Owner are discussed in more detail below under VRRP Considerations. The VIP must be on the same subnet as the corresponding IP Interface. The default value is 0.0.0.0.
Priority	Enter the Priority of the virtual router. If the router is the IP Address Owner of the VIP (as described above), then the Priority will automatically be set to 255 regardless of the value entered. If the router is not the owner of the VIP, than any value in the range of 1 - 254 can be entered. Note that the higher the value, the higher Priority the virtual router will have. The default value is 100.
	1 of 3

Table 12-31. Add VRRP Virtual Router Web Page Parameters

Parameter	Definition
Advertisement Timer	Enter the value of the Advertisement Interval in seconds. This is the time interval that the router will send out advertisements if it is the Master. The range is from 1 - 255 seconds. The Advertisement Interval also defines the time that a Backup will wait until becoming Master. A Backup will become Master if it does not receive an advertisement in approximately three times the Advertisement Interval value. An Advertisement Interval will allow for the fastest fail over time (approximately three seconds). The default value is 1.
Authorization Type	Select the Authentication Type for this virtual router. If None is selected, then no Authentication Key will be used. If Simple is selected, then the virtual router will employ Simple Text Password authentication and use the password in the Authentication Key field (described below). The default value is None.
Authorization Key	Enter the Authentication Key in this field. If the Authentication Type was set to None, then a password will not be used in the VRRP election process regardless if one is entered in this field. If the Authentication Type was set to Simple, then the password entered here will be used to verify correct configuration of the corresponding virtual routers involved in the VRRP election process. The password must be an alphanumeric string from 0 - 8 characters and is case-sensitive. Leaving this field blank is also a valid password. If a blank field password is desired, set the Authentication Type to None and leave the Authentication field blank. After the virtual router is created, select and modify it and set the Authentication Type to simple. A password may also be entered without selecting Simple in the Authentication Type field. In this case, the password will be stored but not used until the Authentication Type is set to Simple. This field is blank by default.
	2 of 3

Table 12-31. Add VRRP Virtual Router Web Page Parameters

Parameter	Definition
Address Owner Override	Enable or Disable Address Owner Override using the drop-down menu. Enabling Address Owner Override allows the VIP to reply to ICMP requests if the router is not the IP Address owner of the virtual router's IP Address (the VIP and IP Address are different). Address Owner Override helps to ensure connectivity and availability of all virtual routers. **Note:** This parameter is not defined in the VRRP RFC 2338. The default value is Disable
Preempt Mode	Enable or disable Preempt Mode using the drop-down menu. Preempt Mode will allow a Backup virtual router with a higher Priority to preempt a Master virtual router with a lower Priority. Note that the router that owns the IP Address associated with the virtual router will always preempt regardless if this is enabled or disabled. To disable this feature, set this field to False. The default value is True.
	3 of 3

4. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To create a VRRP router from the CLI, enter the following command in Configure mode:

`(config-if:<interface name>)#` **ip vrrp <*vr-id*> address <*ip-address*>**

Displaying VRRP Statistics

You can display VRRP router statistics by using either the Web Agent or the CLI.

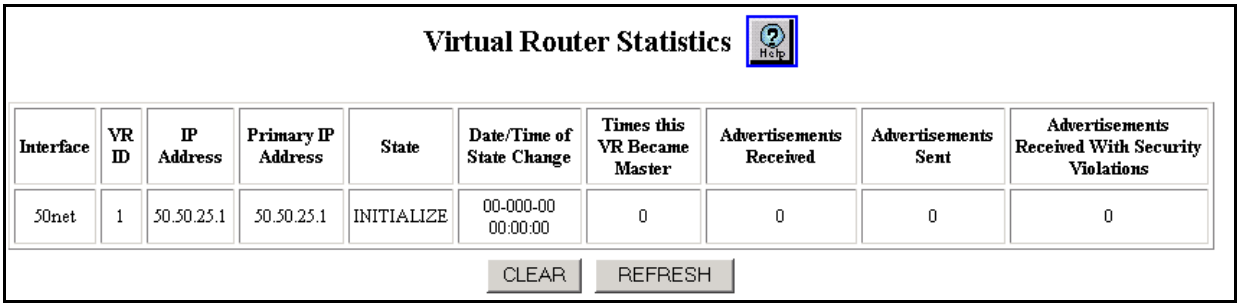
Web Agent Procedure

To display VRRP virtual router statistics using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **Display** folders, and then click **VRRP Statistics**.

   The VRRP Virtual Router Statistics Web page is displayed in the content pane. See Figure 12-39.

Figure 12-39. VRRP Virtual Router Statistics Web Page



2. To modify the VRRP statistics:

- Click **CLEAR** to reset all the entries.
- Click **REFRESH** to receive the most up-to-date information on the entries.

3. See Table 12-32 for an explanation of the VRRP Virtual Router Statistics Web page parameters:

Table 12-32. VRRP Virtual Router Statistics Web Page Fields

Parameter	Definition...
Interface	Displays the IP Interface name that the virtual router is associated with.
VR ID	Displays the virtual router identification number.
IP Address	Displays the IP Address of the virtual router. It is also known as the virtual IP Address (VIP).
Primary IP Address	The Primary IP address is the real IP address of the IP interface that a virtual router is associated with.This address can be viewed in the VRRP statistics page of the Web Agent. **Note:** For more information about the Primary IP Address see "Changing the Primary IP address" later in this chapter.
	1 of 2

Table 12-32. VRRP Virtual Router Statistics Web Page Fields

Parameter	Definition...
State	Displays the current state of the virtual router. ·**Initialize** - Indicates that the virtual router is waiting for a Startup Event. A virtual router could be in this state due to VRRP being disabled Globally, VRRP being disabled on the corresponding IP Interface, or that the associated IP Interface is Down. ·**Backup** - Indicates that the virtual router is in the Backup state. A virtual router in this state monitors the availability of the Master router. ·**Master** - Indicates that the virtual router is in the Master state. A virtual router in this state functions as the forwarding router for the associated IP Address
Date/Time of State Change	Displays the date and time when the last state change occurred. This value is displayed in year-month-day and hh:mm:ss.
Times this VR Became Master	Displays the number of times this virtual router became the Master router.
Advertisements Received	Displays the number of advertisements received by this virtual router.
Advertisements Sent	Displays the number of advertisements sent by this virtual router.
Advertisements Received with Security Violations	Displays the number of advertisements that were discarded by this virtual router. An advertisement is discarded if it contains incorrect parameters.
	2 of 2

CLI Command

To view VRRP virtual router statistics using the CLI, enter the following command in User mode:

> show ip vrrp [cr | <*if-name*> | <*detail*>]

VRRP Configuration Considerations

Selecting VR ID numbers for Virtual Routers

The allowable range for the VR ID of a virtual router is 1-255. For easiest troubleshooting and identification, it is recommended that each virtual router have a unique VR ID. Duplicate VR ID's are allowed, however, as long the IP Interfaces they are associated with are on different VLAN's. Because of this, VRRP cannot be configured in a multinetted network.

Assigning the IP Address of a Virtual Router

When creating a virtual router, it must be first decided whether to make the IP Address of the virtual router (the VIP) unique or the same as the Primary IP Address of the associated IP Interface address. If the VIP is unique, this is known as non-IP Address Owner. This means the VIP is not "owned" by the router as a Primary IP Address. If the VIP is not unique, meaning it matches the Primary IP Address of the associated IP Interface, then this is known as IP Address Owner. The VIP Address is "owned" by the associated IP Interface. Each method will be discussed below in more detail.

IP Address Owner Features

IP address owner allows for the highest possible priority. A Priority of 255 is reserved for IP Address Owner. In a properly configured network, there will never be a virtual router with a higher priority (if there is another virtual router on the same subnet with a priority of 255, then the network is misconfigured with duplicate IP Addresses). This means that this virtual router will always assume Mastership if it is available to do so.

IP address owner features are:

- Requires less configuration. The Priorities are automatically assigned. Backup routers can use the default Priority settings and correct Master-Backup election will be achieved.
- Another IP Address does not have to be reserved for the virtual router. This is helpful if IP Addresses are scarce or limited.
- Preempt mode cannot be disabled. A virtual router that is IP Address Owner will always preempt and assume Mastership from a Backup router.
- A virtual router that is IP Address Owner will “overwrite” the MAC Address of the associated IP Address. In the ARP Cache Table, the MAC Address of the IP Address will be the VRRP MAC.

Non-Address Owner Features

Non-address owner features are:

- Requires more configuration but allows for more control and customization of the Priorities of each virtual router.
- The virtual router IP Address and associated IP Address will both be displayed in the ARP Cache Table. This is helpful for troubleshooting and identification.
- Preempt Mode can be used. As described in “IP Address Owner Features,” if the Preempt Mode is set to False, then a higher Priority virtual router will not assume Mastership from a virtual router that is currently Master. This is helpful when a router goes down and the Backup router becomes Master. When the original router comes back online, it will not take Mastership from the current Master. This will allow the current traffic to remain unchanged. Another state change will only occur if the current Master router goes down.

This is especially helpful in a network where a routing protocol is used (RIP, OSPF). In most cases, the VRRP election process will take place before the routing table has been updated. This may cause a longer period of traffic loss.

Assigning Priorities to Virtual Routers

The allowed configurable range for the Priority of a virtual router is 1-254. 255 is reserved for a virtual router that is IP Address Owner (as described above). If two routers are participating in VRRP, making the desired Master router 254 (or 255 if IP Address Owner) and the Backup router 100 (the default Priority) is a common configuration. If multiple routers will be backing up the Master router, then the Priorities can be assigned in descending order (100, 99, 98, etc., for example).

Equal Priorities are also allowed but may be more difficult to troubleshoot. If two Master routers have the same Priority (and the same VRRP settings), the VRRP election process then checks the associated IP Address of each virtual router. In that case, the associated IP Address with a higher value becomes Master. However, if a Master and Backup virtual router have the same Priority, the Backup will not become Master even if it has a higher associated IP Address. In other words, equal priorities are not enough for a backup router that is already in the Primary state to become Backup.

Selecting an Advertisement Interval

In most cases, leaving the Advertisement Interval to its default value of 1 is adequate. Usually, this is also desired because it will allow for the fastest fail over time. As described above, the time a Backup will become Master if it stops receiving advertisements is approximately three times this value in seconds. A higher value may be selected if the network is very stable and less advertisement traffic is desired. In any case, it is not recommended that the interval be increased to values above five or six.

Deciding to Use Authentication

In networks where there is little or no security risk and a minimal chance of miscommunication, authentication is probably not needed. Using Simple Text Password authentication can provide protection against accidental misconfiguration. The correct password must be entered for a new virtual router before it will begin the VRRP election process with another virtual router. Simple Text Password authentication does not provide protection from hostile attacks.

Changing the Primary IP address

If there is a virtual router that is an IP Address Owner created on an IP interface, the priority will be changed to 100 (default) if the Primary IP Address is changed, because a Priority of 255 is reserved for the IP Address Owner only.

If the virtual router is a non-Address Owner and the associated Primary IP Address is changed to the same as the virtual router IP Address, the Priority of that virtual router is changed to 255, since it is now the IP Address Owner.

If the Primary IP Address is changed so that it is no longer on the same subnet as an associated virtual router, the virtual router will be deleted

Configuring IRDP

ICMP Router Discovery Protocol (IRDP) is an alternative router discovery protocol using ICMP messages on multicast links. ICMP uses router discovery messages, known as router advertisements and router solicitations.

Each router periodically multicasts a router advertisement from each of its multicast interfaces, announcing the IP address of that interface, and other router IP addresses. Hosts discover the addresses of their neighbor routers by listening for the advertisements. When a host attached to a link starts up, it may multicast a router solicitation to ask for immediate advertisements, rather than waiting for the next periodic one to arrive.

If no advertisements arrive, the host re-transmits the solicitation, but does not send additional solicitations. Routers that subsequently start up, or were not discovered because of packet loss or temporary link partitioning, are eventually discovered by reception of their periodic (unsolicited) advertisements.

Enabling IRDP on an Interface

You can enable IRDP on an interface using either the Web Agent or the CLI.

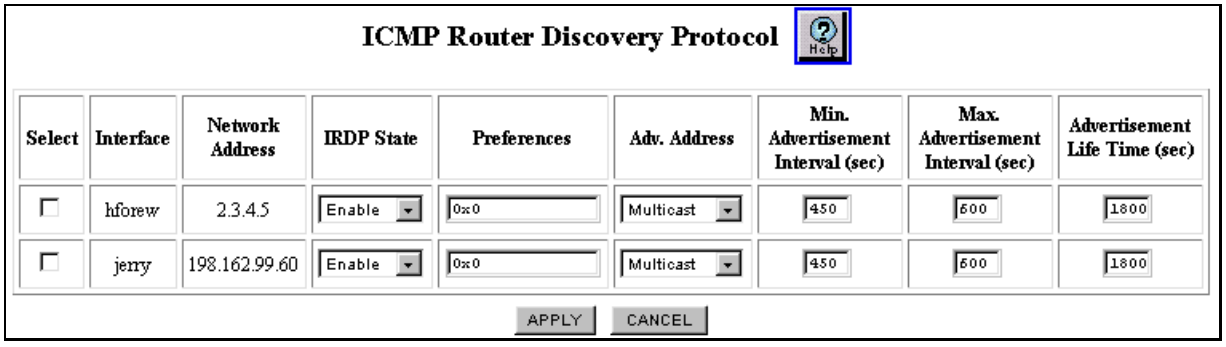
Web Agent Procedure

To enable IRDP on an interface using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **IRDP**.

   The ICMP Router Discovery Protocol Web page is displayed in the content pane. See Figure 12-40.

Figure 12-40. ICMP Router Discovery Protocol Web Page



2. Select the interface on which to enable IRDP in the Select column. A checkmark displays.

3. See Table 12-33 to configure the ICMP Router Discovery Protocol Web page parameters.

Table 12-33. ICMP Router Discovery Protocol Web Page Fields

Parameter	Allows You to...
Select	Select the interface to be configured.
Interface	View the IRDP interface name.
Network Address	View the network IP address of the IRDP interface.
IRDP State	Select Enable IRDP on the selected interface. The default value is **Disable**.
Preferences	Enter the preference of the address as a default router address, relative to other router addresses on the same subnet. The minimum value (80000000 hex) is a signed 32-bit value used to indicate that the address should not be used by neighboring hosts as a default router address, even though it may be advertised. The default value is **0**.
Adv. Address	Select an IP destination address used for multicast router advertisements sent from the interface. Options include: • **Multicast** - Used on any link where listening hosts support IP multicast. The default value is 224.0.0.1. • **Broadcast** - Used on any link where listening hosts support IP unicast. The default value is 255.255.255.255.
Min. Advertisement Interval (sec)	Enter the minimum time (in seconds) that is allowed between sending unsolicited multicast router advertisements from the interface. This value must be no less than three seconds and no greater than the Max. Advertisement Interval. The default value 450 is **0.75 times the maximum interval**.
Max. Advertisement Interval (sec)	Enter the maximum time (in seconds) allowed between sending multicast router advertisements from the interface. This value must be no less than four seconds and no greater than **1800** seconds. The default value is **600 seconds**.
Advertisement Life Time (sec)	Enter the time (in seconds) of the life of a router advertisement that is sent from the interface. This value must be no less than the maximum advertisement interval and no greater than **9000** seconds. The default value is **1800**.

4. Select **Enable** from the **IRDP State** pull-down menu.

5. Select **Multicast** from the **Adv. Address** pull-down menu.

6. Modify the default value in the **Min. Advertisement Interval (sec.)** field and enter the minimum time interval that passes before the host contacts the switch.

7. Modify the default value in the **Max. Advertisement Interval (sec.)** field and enter the maximum time interval that passes before the host contacts the switch.

8. Modify the default value in the **Advertisement Life Time (sec.)** field, and enter a duration, in seconds, of the IRDP advertisement.

9. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

10. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **Global Configuration**.

    The IP Global Configuration Web page is displayed in the content pane. See Figure 12-35.

11. Select **Enable** the **IP Multicast Forwarding** field pull-down menu.

12. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To enable IRDP on an interface using the CLI, enter the following command in **Enable/Configure/Interface** mode:

`(config-if:<`*`if name`*`>)#` **ip irdp**

Configuring LDAP

Lightweight Directory Access Protocol (LDAP) allows you to access the "Access Control Lists"(ACLs), retrieve them from a database on an LDAP server, and apply them to the Avaya P580 or P882 Multiservice switches.

You can also:

- View LDAP statistics for the switch
- Configure a primary and secondary LDAP server
- Configure a search base for the switch to contact in response to an LDAP request.

Configuring a secondary server ensures that LDAP requests can be fulfilled if a primary server fails. The LDAP client sends a search for access control lists to the primary server if the client finds the primary server. The primary server retrieves the access lists from the LDAP database and returns them to the switch's LDAP client. If the client does not find the primary server and if the primary server does not respond after three retries, the client tries to connect to the secondary server. If the secondary server fails after three retries, the client times out. The LDAP client applies the access lists to manage the way traffic is forwarded.

This section contains procedures for the following tasks:

- Configuring LDAP Settings
- Viewing LDAP Statistics
- Configuring a Static Route for the PPP Console
- Configuring the IP Interface for the PPP Console

Configuring LDAP Settings

You can configure LDAP settings from either the Web Agent or the CLI.

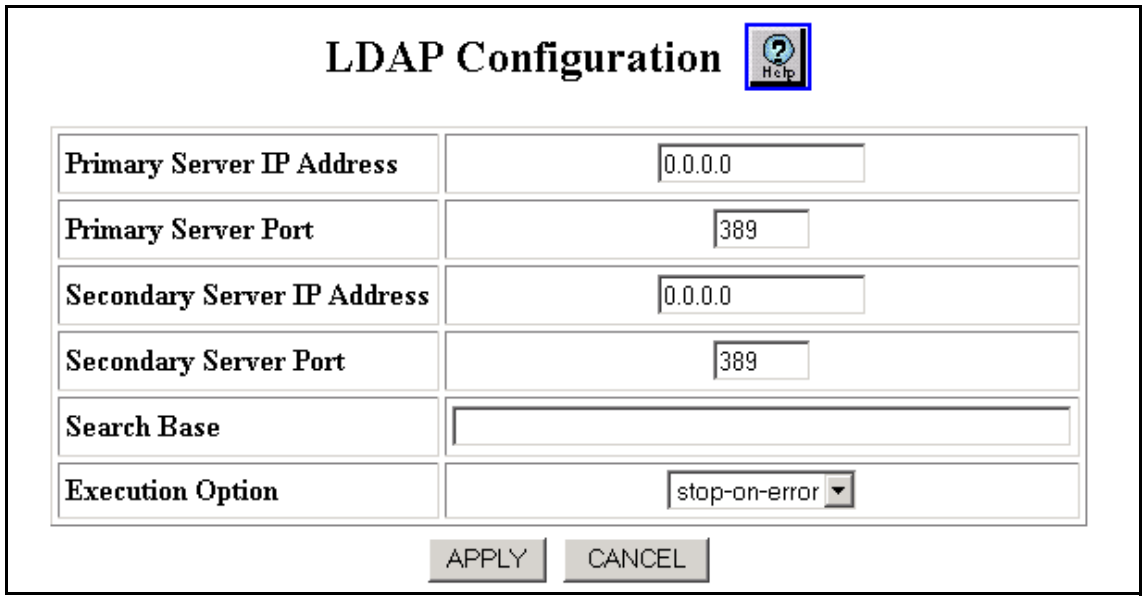
Web Agent Procedure

To configure LDAP settings from the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **LDAP Configuration**.

   The LDAP Configuration Web page is displayed in the content pane. See Figure 12-41.

Figure 12-41. LDAP Configuration Web Page



2. See Table 12-34 to configure the LDAP Configuration Web page parameters.

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 12-34. LDAP Configuration Web Page Parameters

Parameter	Allows You To...
Primary Server IP Address	Enter the IP address of your primary LDAP server for the access control list domain. This address is used first when connecting to and downloading access lists from an LDAP server. The default value of 0.0.0.0 indicates to the client that there is no primary LDAP server.
Primary Server Port	Enter the **port number** of the primary LDAP server for the access control list domain. The port number is used in conjunction with the primary server IP address. There are no special overload values. The default port is 389.
Secondary Server IP Address	Enter the backup LDAP server IP address for the access control list domain. This address is used as a backup when connecting to and downloading access lists from an LDAP server. If the LDAP client is unsuccessful in connecting to or downloading access lists from the primary server, the secondary server IP address is used. If the primary server IP address has a value of 0.0.0.0, the secondary server IP address is used. The default value is 0.0.0.0. **Note:** Setting the IP address of the secondary server to 0.0.0.0 indicates to the LDAP client that there is no secondary server.
	1 of 2

Table 12-34. LDAP Configuration Web Page Parameters

Parameter	Allows You To...
Secondary Server Port	Enter the backup LDAP server port number for the access control list domain. The port number is used in conjunction with the secondary server IP address. There are no special overload values. The default port number is 389.
Search Base	Enter the search criteria that will be sent to the LDAP server. The default value is "ou=Devices, ou=AvayaPolicyManager, o=Avaya". **Note:** No default for the Search Base, this field is empty if the LDAP server has not been installed:
Execution Option	Open drop-down menu and select **stop-on-error** or **ignore-errors**. This option lets you decide if you want the policy to continue being sent to the device, if LDAP encounters any errors. Select **stop-on-error** if you want execution of the policy to stop on the first error encountered.Select **ignore-errors** if you want execution of the policy to continue even if errors are encountered. In this case, any commands containing errors are ignored. The default value is **stop-on-error.**
	2 of 2

CLI Command

To configure LDAP settings on a primary and secondary server, use the following commands:

- To change the LDAP search base,
  `(configure)#` **ldap search-base <*search-base-DN*>**
- To change the primary LDAP server's primary ip address and port,
  `(configure)#` **ldap server primary <*ip-addr*> [<*port-num*>]**
- To change the secondary LDAP server's primary ip address and port,
  `(configure)#` **ldap server secondary <*ip-addr*> [<*port-num*>]**
- To configure policy retrieval to stop on the first error, or ignore errors,
  `(configure)#` **ldap execution-option {stop-on-error | ignore-errors}**

Viewing LDAP Statistics

You can view LDAP statistics from either the Web Agent or the CLI.

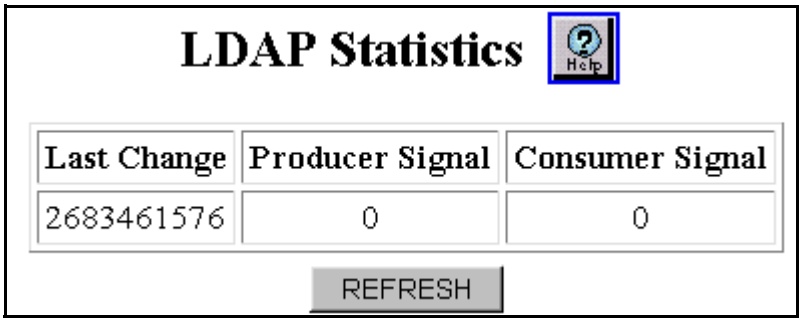
Web Agent Procedure

To view LDAP statistics using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > Display** folders, and then click **LDAP Statistics**.

   The LDAP Statistics Web page is displayed in the content pane. See Figure 12-42.

Figure 12-42. LDAP Statistics Web Page



2. Click **Refresh** to dynamically update **LDAP Statistics** parameters.

3. See Table 12-35 for an explanation of the LDAP Statistics Web page parameters:

Table 12-35. LDAP Statistics Web Page Parameters

Parameter	Definition
Last Change	Displays the **sysUpTime** since this device was last modified. You can detect a change in the Policy Capabilities by polling a single object using this information.
Producer Signal	The sequence number that, when modified, triggers the LDAP client to download the latest policy from the LDAP server. Typically, Avaya Policy Manager will set this value whenever there is a new policy to download. If this value is a non-zero value, the LDAP client will compare it to the producer signal on the LDAP server. No comparison is made if the value is zero.
	1 of 2

Table 12-35. LDAP Statistics Web Page Parameters

Parameter	Definition
Consumer Signal	Indicates the success of the LDAP client when downloading a policy. If the consumer signal matches the producer signal, downloading LDAP to a policy was successful. If the consumer signal is -1, then either the LDAP client had a problem processing the access lists or the consumer signal set on the LDAP client did not match the signal configured on the LDAP server. If the consumer signal is not -1 and does not match the producer signal, then the LDAP client was unable to connect to the LDAP server(s).
	2 of 2

CLI Command

To view configuration and LDAP statistics using the CLI, enter the following command:

> **show ldap**

Configuring a Static Route for the PPP Console

Web Agent Procedure

To configure a PPP Console static route:

1. Configure your console serial port as a PPP Console. See the "Connecting a Modem" section in Chapter 2, "Setting Up the Switch".
2. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Static Routes**.

   The IP Static Routes Web page is displayed in the content pane.
3. Select **CREATE**. The Add IP Static Routes Web page opens (Figure 12-6).
4. Enter the IP address that is to be associated with the PPP console port in the Network Address field.
5. Enter the network mask IP address in the **Mask** field.
6. Enter the IP address of the gateway associated with this static route in the **Next-Hop Address** field.
7. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings. The IP Static Routes Web page reopens.
8. Click **APPLY** to save your changes, or **CANCEL** to clear your selection.

Configuring the IP Interface for the PPP Console

You can configure an IP interface for the PPP console from either the Web Agent or the CLI.

Web Agent Procedure

To configure the PPP console with an IP address and mask using the Web Agent:

1. Configure your console serial port as a PPP Console. See "Connecting a Modem", in Chapter 2, "Setting Up the Switch".
2. Connect your modem cable to the switch's serial port.
3. In the navigation pane, expand the **Routing > IP > Configuration** folders, and then click **Interfaces**.

   The IP Interfaces Web page is displayed in the content pane. See Figure 12-36.
4. Select **CREATE**. The Add IP Interface Web page opens (Figure 12-3).
5. Select **Serial-Console** from the **VLAN** field pull-down menu. This indicates the interface for the PPP console.
6. Enter the IP address in the **Network Address** field that is to be associated with the PPP console port. If you do not enter a name for this interface, the IP address is used.
7. Enter the network mask IP address in the **Mask** field (for example, 255.255.255.0).
8. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings. The IP Interfaces Web page redisplays.
9. Click **APPLY** to save your changes, or **CANCEL** to clear your selection.

CLI Command

To configure the PPP console after you have made all the applicable cable connections with an IP address using the CLI, enter the following command in Interface mode:

`(config if <interface-name>)#` **ip vlan {<*vlan-id*> | name <*vlan-name*> | ethernet-console | serial-console}**

13 Configuring Access Lists

Overview

Contents

This chapter contains the following sections:

- Creating Access Rules
- Enabling an Access List
- Example: Using an Access Control Rule to Filter Web Traffic
- Logging ACL Activity
- Optimizing Switch Performance

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

What is an Access Control List?

An Access Control List (ACL) is a group of Access Control Rules. Access Control Rules are used to describe *how* to forward (route) packets, as opposed to where to forward them. The *how* can be to forward the packet with a specific priority(0-7), forward the traffic with an un-changed priority, or filter packets (drop).

ACLs provide the mechanism to prioritize traffic flows through the router and the network. This traffic flow management is commonly referred to as Quality of Service (QoS). See Chapter 25 for more information on QoS.

Standard vs. Extended ACLs

A *standard* ACL allows you to prioritize traffic by the Source IP address.

An *extended* ACL provides greater control over what traffic is prioritized. Extended ACLs can use any or all of the following parameters:

- Source IP address
- Destination IP address
- TCP/UDP Source port
- TCP/UDP Destination port.
- Protocol ID (RFC1700)

These five parameters are referred to as a “5 Tuple”. Source or destination addresses allow you to prioritize traffic between any combination of host, subnet, and network addresses.

Further, if you specify a source or destination TCP/UDP port, you can prioritize specific traffic between hosts on two different networks. For example, you could increase the priority of Voice Over IP (VoIP) traffic between two subnets in a call center by specifying the UDP port number used by the phone system.

*** Note:** If you specify TCP or UDP port numbers, you must specify Protocol ID 6 or 17 respectively. Failure to specify the Protocol ID number will result in an error message and the ACL will not be created.

Using a Protocol ID, you can prioritize traffic based on the Protocol ID number. RFC 1700 lists Protocol ID numbers. You could, for example, block all RSVP traffic (Protocol ID 46) through the router.

Naming Conventions for ACLs

The Avaya Multiservice switch supports ACL names up to 32 characters (Alpha-Numeric). Spaces are allowed in the names but are not recommended. Instead of spaces, use the underscore (_) character.

ACL Names are entered when you create a rule. By using the same ACL Name for multiple rules, you are effectively creating a list of rules. The Access List Index determines the order of rules in the list. For that reason, when you create a rule, you give the rule the Access List Name and the Access List Index. No two rules in a list may have the same index.

If you create rules with ***different*** ACL Names, you have created different lists. This is useful when you need more than one ACL. For example, you may have one ACL for normal working hours, one ACL for nighttime hours, and a third list for Holidays. When you enable (activate) an ACL, the currently enabled (active) ACL is automatically disabled (de-activated). In other words, there can only be one active ACL.

The only restriction on the amount of ACLs you can create is that the total number of rules cannot exceed 512. So, you could create a single ACL with 512 rules; or one ACL with 300 rules and a second with 212 rules, etc. See “Creating Access Rules” for guidelines to consider when designing ACLs.

An ACL name can be a number. For example, you can create a list whose name is 1 or 151. If you chose to use numbers, keep in mind the following restrictions:

- Numbers 1 through 99 are reserved for Standard type Rules ***ONLY***. So, for example, if you try to create an Extended Rule whose ACL Name is 1, it will be rejected.
- Numbers 100 through 199 are reserved for Extended type Rules ***ONLY.*** So, for example, if you try to create a Standard Rule whose ACL Name is 100, it will be rejected.
- ACL names that contain any letter (Alphabetic) character, can be either Standard or Extended.
- You cannot mix ACL types in a list. This means that if you create an ACL with a Standard Rule with ACL Name Test1, you cannot create an Extended Rule in ACL Test1. If you do, you will receive the following message:

```
Access Rule Name is already being used by the other
type.
Choose a different name and try again
```

How Packets are Processed

Assuming an ACL is active, when a packet arrives on the Avaya Multiservice switch, the parameters in the packet are compared to the parameters in the Access Rule starting with the lowest index number. If there is a match, that rule is applied to the packet and the search stops.

If the 5-tuple's of the packet and rule do not match, the next (higher index) rule is compared. This process continues until a match is found, or there are no more rules. There is an implied **permit all** at the end of every list. Therefore, if no match is found, the packet is forwarded with the priority un-changed.

What are Wildcards?

Wildcards are a template that govern which part of an IP address is significant when evaluating a rule. When you create a rule based on source or destination IP address, you must also specify the Wildcard.

Wildcards are in principal, the same as a subnet mask. The differences are you invert the mask's bits and there is no requirement of contiguous bits. For Example: a decimal wildcard of 0.255.0.255 is allowed.

For example: If you want to create a rule that blocks all traffic on the 192.168.24.0 (subnet mask 255.255.255.0) network, you would specify a Wildcard of 0.0.0.255 in the rule.

If you wanted to block traffic from a specific host whose IP address was 192.168.24.143 (subnet mask 255.255.255.0) you would specify a Wildcard of 0.0.0.0. This mask "tells" the supervisor to evaluate the entire IP address when evaluating a packet against the rule.

What is TCP Established?

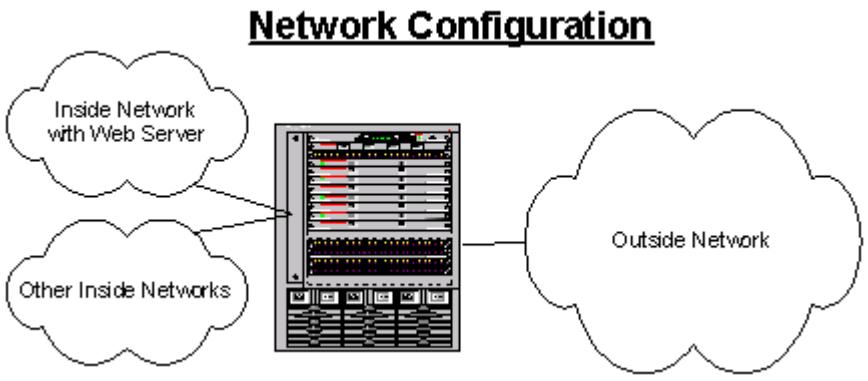
TCP Established is a criteria applied by a rule where the “Acknowledge” bit in a TCP header is examined. If this option is not “checked”, the rule will apply to the packets whose Acknowledge bit is clear (0). If the option is checked, packets that have the Acknowledge bit set will be affected by the rule.

The Acknowledge bit in the TCP header, when 0, indicates that the packet is an initial “call” to the destination. The host sending the message will clear the bit (0). The host that responds to the message will set the bit(1) indicating this message is a response. Effectively the call is now “Established”. All subsequent packets between these two hosts for this session will have the Acknowledge bit set.

For example: Suppose the Avaya Multiservice switch has interfaces to “Outside” networks as well as interfaces to “Inside” networks. The Outside networks need access to a Web server and should be denied access to any other resource within the Inside network.

Hosts on the Inside networks should have full access to all other resources on the Inside.

The Web server itself should not be able to establish any new connections to the Outside but should be able to pass traffic to the other Inside networks.



Assume the following abbreviations:

WS = Web Server.

IN = Inside Networks

ON = Outside Networks

Any = Both Inside and Outside Networks

The rules for implementing the above restrictions would be as follows:

Rule #	Rule	TCP Established
1	allow Any to call WS dest port 80	Un-Checked
2	allow WS to respond	Checked
3	allow WS to IN	Un-Checked
4	deny WS to Any (Outside)	Un-Checked
5	allow IN to Any	Un-Checked
6	deny remaining traffic from Outside to Any	Un-Checked

Rules 1 and 2 collectively manage Web traffic ***to and from*** the Web server (WS). Rule 1 says that any source address can get to the Web server's IP address using destination port 80. Because the TCP Established criteria is unchecked, hosts from any network can send a TCP "call setup" message as a first step in requesting a Web page.

Rule 2 says that the Web server may respond to any (TCP) Web request. Although it can send a message back to **any** address from any source port, only messages that are in response to a Web request will be forwarded because TCP Established is checked and the source port criteria is specified.

Rules 3 and 4 handle traffic from the Web server that is not in response to a Web request. Rule 3 gives the Web server access to the rest of the Inside networks. And Rule 4 blocks the Web server from getting to the rest of the networks (Outside).

Rule 5 gives the hosts on the Inside network access to any network.

Rule 6 blocks any other host from using resources on any of the Inside networks.

*** Note:** In this simple example, pseudo-rules are used. In practice, the pseudo-rule "allow WS to IN" would require that you create rules that forward traffic from the Web server's IP address to ***each*** network on the Inside explicitly. If you had 30 inside networks, you'd create 30 rules. This is where a subnetted network would be powerful; because you could summarize subnets into a few rules.

Creating Access Rules

This section contains the following procedures:

- Creating Standard Access Rules
- Creating Extended Access Rules

Creating Standard Access Rules

To create standard access rules:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Access Lists**.

   The IP Access List Web page is displayed in the content pane. See Figure 13-1.

   ***Note:** The IP Access List displays all standard and extended access rules that have been created. If no rules have been created, the following statement displays: `No IP Access Rules are currently configured`.

   Due to its size, Figure 13-1 shows the Access List Web page split in two separate sections.

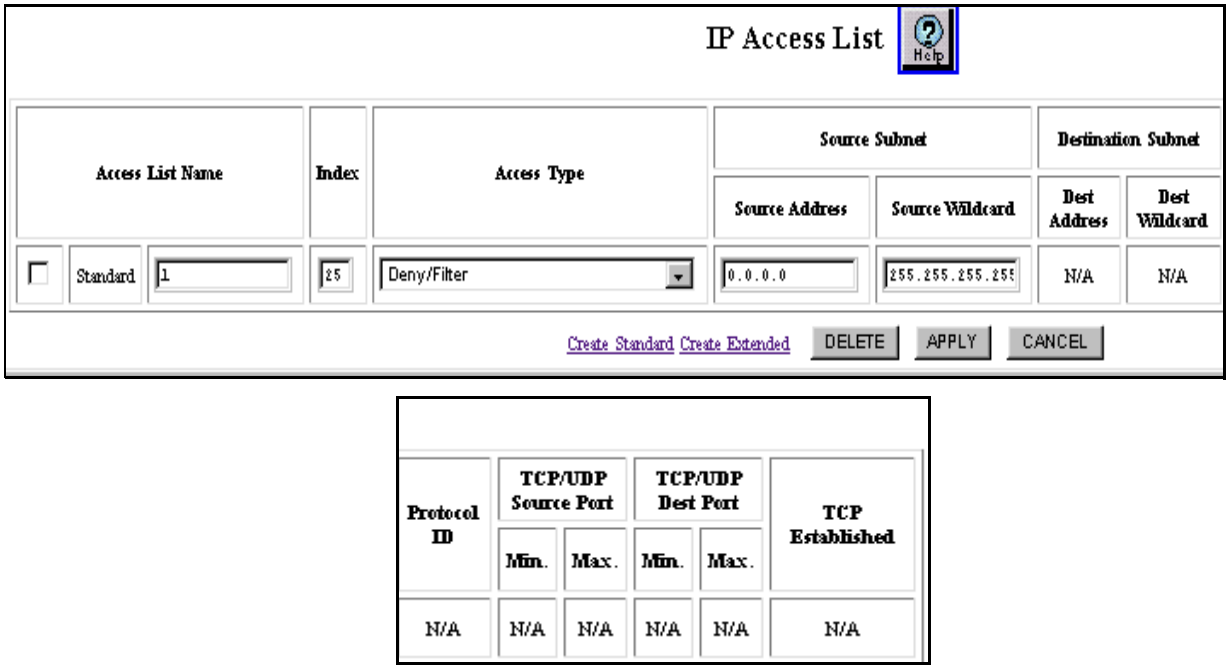
   The switch supports a maximum total of 512 access control rules, regardless of the number of access lists. For example, you could create the following three ACLs:

   - ACL A with 100 rules
   - ACL B with 200 rules
   - ACL C with 212 rules

   The combined total of rules cannot exceed 512. Only one access control list can be enabled at a time.

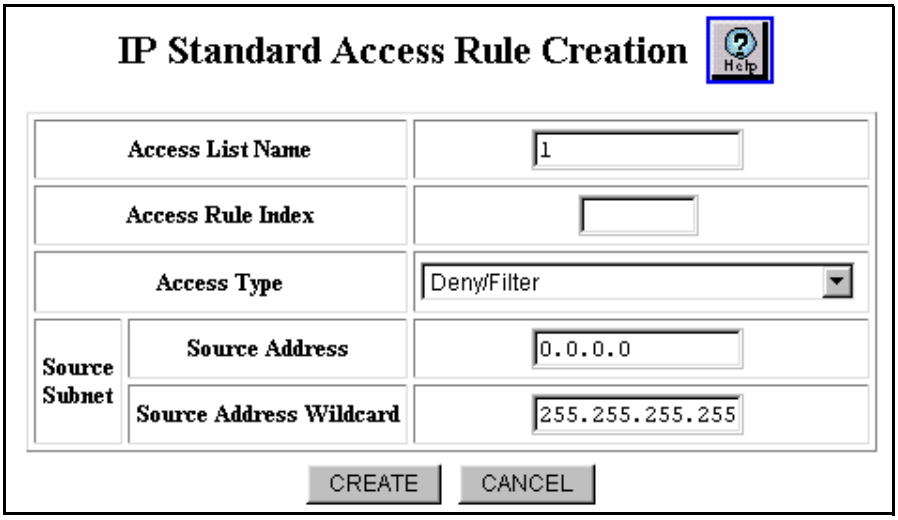
   Index numbers of access rules can range from 1 to 512. Do not use index numbers higher than 512.

Figure 13-1. IP Access List Web Page



2. Click **Create Standard**. The Create Standard Access List Web page is displayed. See Figure 13-2.

Figure 13-2. Create IP Standard Access List Web Page



3. See Table 13-1 to configure the Create Standard IP Access List Web page to filter or prioritize traffic:

Table 13-1. Create Standard IP Access List Web Page Parameters

Parameter	Allows you to...
Access List Name	The Avaya Multiservice switch supports ACL names up to 32 characters (Alpha-Numeric). Spaces are allowed in the names but are not recommended. Instead of spaces, use the underscore "_" character. See "Naming Conventions for ACLs" for more information.
Access Rule Index	Enter the sequence number for each new rule you create. Index numbers can be 1 through 512. Packets are compared against rules in ascending index order. **Note:** Entering a new rule may override other rules. Review your current configuration prior to creating new access list rules.
Access Type	Select the method of handling incoming datagrams based on the IP access type from the following pull-down menu options: • **Deny/Filter** - Allows you to filter out (drop) packets based on the specified configuration. • **Permit/Fwd pri8 (high) to pri1 (low)** - Allows you to prioritize traffic based on the specified configuration. • **Permit/Fwd with no change in priority** - Allows you to forward traffic with no change in priority.
Source Subnet	• **Source Address** - Enter the IP address that you want to deny or grant access to the switch. The Wildcard will determine how the address is evaluated. • **Source Address Wildcard** - Enter the Wildcard for this address. For more information on wildcards, see "What are Wildcards?" earlier in this chapter.

4. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings.

Creating Extended Access Rules

To create Extended Access Rules:

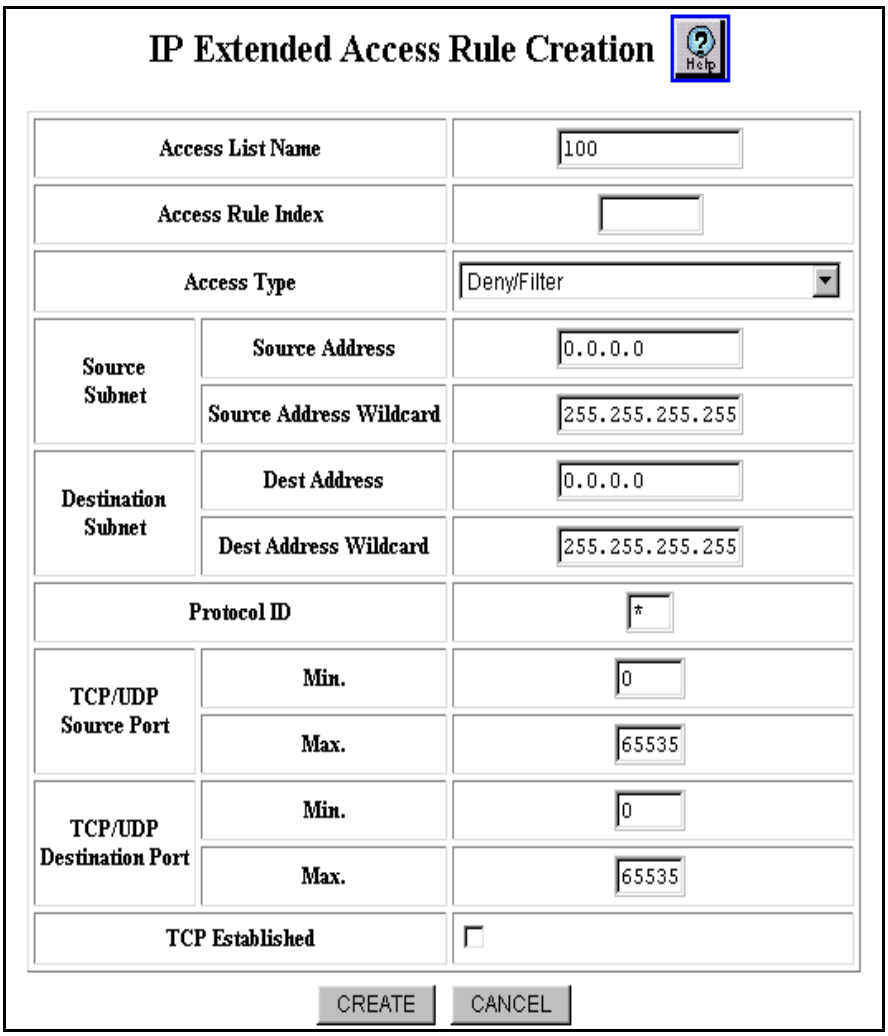
1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Access Lists**.

   The IP Access List Web page is displayed in the content pane. See Figure 13-1.

***Note:** The IP Access List Web page displays all standard and extended access rules that have been created. If no rules have been created, the following statement displays: `No IP Access Rules are currently configured`.

2. Select **Create Extended**. The IP Extended Access Rule Creation Web page is displayed. See Figure 13-3.

Figure 13-3. IP Extended Access Rule Creation Web Page



3. See Table 13-2 to configure the IP Extended Access Rule Creation Web page parameters to filter or prioritize traffic.

4. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings.

Table 13-2. IP Extended Access Rule Creation Parameters

Parameter	Allows you to...
Access List Name	Enter the alphanumeric name of the access list this rule will be added to. See "Naming Conventions for ACLs" for more information
Access Rule Index	Enter the sequence number for each new rule you create. Index numbers can be 1 through 512. Packets are compared against rules in ascending index order. **Note:** Entering a new rule may override other rules. Review your current configuration prior to creating new access list rules:
Access Type	Select the method of handling incoming datagrams based on the IP access type you set from the following options: • **Deny/Filter** - Allows you to filter out traffic based on the specified configuration. • **Permit/Fwd pri8 (high) to pri1 (low)** - Allows you to prioritize traffic based on the specified configuration. • **Permit/Fwd with no change in priority** - Allows you to forward traffic with no change in priority.
Source Subnet	• **Source Address** - Enter the IP address that you want to deny or grant access to the switch. The Wildcard will determine how the address is evaluated. • **Source Address Wildcard** - Enter the Wildcard for this address. For more information on wildcards, see "What are Wildcards?" earlier in this chapter.
Destination Subnet	• **Dest Address** - Enter the IP address that you want to deny or grant access to the switch. The Wildcard will determine how the address is evaluated. • **Dest Address Wildcard** - Enter the Wildcard for this address. For more information on wildcards, see "What are Wildcards?" earlier in this chapter.
Protocol ID	Specify a protocol ID to be filtered. (For example, ICMP=1, IGMP=2). A single asterisk (*) indicates all protocols. RFC 1700 defines the protocol IDs. To see the complete list of protocol numbers, see http://www.iana.org/assignments/protocol-numbers.
	1 of 2

Table 13-2. IP Extended Access Rule Creation Parameters

Parameter	Allows you to...
TCP/UDP Source Port	Specify a range of source ports that pass between two hosts or switches using the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP). Options include: • **Min.** - The lowest numbered port in the range. The default is 0. • **Max.** - The highest numbered port in the range. The default is 65,535. *** Note:** The protocol ID parameter must first be configured with either 6 for TCP or 17 for UDP, to enable the TCP/UDP Source port parameter. To see the complete list of well-known port numbers (specifically in relation to the destination port), see: http://www.iana.org/assignments/port-numbers.
TCP/UDP Destination Port	Specify a range of destination ports that pass data between two hosts or switches using the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP). Options include: • **Min.** - The lowest numbered port in the range. The default is 0. • **Max.** - The highest numbered port in the range. The default is 65,535. **Note:** The protocol ID parameter must first be configured with either 6 for TCP or 17 for UDP, to enable the TCP/UDP Destination port parameter. To see the complete list of well-known port numbers (specifically in relation to the destination port), see: http://www.iana.org/assignments/port-numbers.
TCP Established	Criteria for matching TCP packets of established (connected) or not established (initial call) sessions.
	2 of 2

Enabling an Access List

You can enable an access list on the switch through the Web Agent and the CLI.

When you enable an ACL, the switch:

- May change the maximum number of forwarding entries for IP unicast traffic to improve the performance of the switch. If the maximum number of entries is set to:
  - Less than 15,000, the switch automatically doubles the setting.
  - Between 15,000 and 30,000, the switch changes the setting to 30,000.
  - Greater than 30,000, the switch does not change the setting.
- Automatically sets hash mode for IP unicast traffic to **SA-DA**. This setting improves the performance of the switch when an ACL is enabled.

Web Agent Procedure

To enable an access list:

*** Note:** Before you can enable an access list, you must first configure the access list. For information about configuring access lists, see "Configuring Access Lists" earlier in this chapter.

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Access Control**.

   The IP Access Control Web page is displayed in the content pane. See Figure 13-4.

Figure 13-4. IP Access Control Web Page



2. Select **Enable** from the **Enable** field pull-down menu to filter inbound traffic.

3. Select the name of the access list to be used for filtering when IP access control is enabled from the **IP Access List** field pull-down menu.

4. Click **APPLY** to save your changes.

CLI Command

To activate an access list using the CLI, enter the following command in configure mode:

`(configure)#` **access-group** <***access-list-name***>

To disable an access list using the CLI command, enter the following command in configuration mode:

`(configure)#` **no access-group** <***access-list-name***>

Example: Using an Access Control Rule to Filter Web Traffic

To configure your switch to filter Web traffic to a particular Web server:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Access Lists**.

   The IP Access List Web page is displayed in the content pane. See Figure 13-1.

2. Select **Create Extended**. The IP Extended Access Rule Creation Web page displays (Figure 13-3).

3. Enter a number between 100 and 199 (or Alphanumeric) (for extended ACL) in the **Access List Name** field to identify your new access control list.

4. Enter a number in the **Access Rule Index** field to identify the access rule.

5. Select **Deny/Filter** from the **Access Type** field pull-down menu.

6. Leave 0.0.0.0 and 255.255.255.255 as the Source Address and Source Address Wildcard for the Source Subnet parameter settings. All source traffic will match.

7. Enter the IP address in the **Dest Address** field that represents the destination address of the Web server. Enter a wildcard of 0.0.0.0 to identify the specific IP address of the destination Web server.

*** Note:** To deny/filter traffic to a specific address and not to an entire subnet, you must specify the destination IP address of the network node, and use a subnet wildcard of 0.0.0.0.

8. Specify the TCP protocol ID = 6. If you specify TCP or UDP port numbers, you must specify Protocol ID 6 or 17 respectively.

9. Enter the following in the **TCP/UDP Destination Port** field:

   — a Min. of 80 (HTTP).

   — a Max. of 80 (HTTP).

10. Leave the **TCP/UDP Source Port** field alone:

   — a Min. of 0 (HTTP).

   — a Max. of 65536 (HTTP).

11. Select **TCP Established**. A check mark displays in the check box.

12. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings.

Each time any end user attempts to access the Web server specified by the destination TCP/UDP ports, all Web requests are filtered.

Second Example

Filtering Traffic Between 10.1.1.0 and 10.1.2.0 Subnets. It is possible to filter traffic to either a specific address or to an entire subnet. In this example, all traffic between the two subnets is filtered:

This example also assumes that the network is a Class C sub network (255.255.255.0).

To filter traffic between both subnets, you must create two access rules.

To create the Extended Access list and rules:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Access Lists**.

   The IP Access List Web page is displayed in the content pane. See Figure 13-1.

2. Select **Create Extended**. The IP Extended Access Rule Creation Web page is displayed. See Figure 13-3.

3. Enter a number between 100 and 199 (or Alphanumeric) (extended ACL) in the **Access List Name** field to identify your new access control list.

4. Enter a number in the **Access Rule Index** file to identify the access rule.

5. Select **Deny/Filter** from the **Access Type** field pull-down menu.

6. Enter the source address (10.1.2.0) and the source address wildcard (0.0.0.255), respectively, in the **Source Address** field.

7. Enter the destination address (10.1.1.0) and the destination address wildcard (0.0.0.255), respectively, in the **Dest Address** field.

8. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings. Once you create both access rules, all traffic between subnet 10.1.1.0 and 10.1.2.0 is deny/filtered.

***Note:** Traffic between any other 10.1.x.0 subnets are not filtered because the access rules only deny/filter traffic between subnets 10.1.1.0 and 10.1.2.0.

To deny/filter traffic to a specific address and not to an entire subnet, you must specify the destination IP address of the network node, and use a subnet wildcard of 0.0.0.0.

To deny/filter all traffic, you must specify a destination address of 0.0.0.0 and a wildcard of 255.255.255.255. This is useful if you want to filter all traffic except traffic that matches a previous rule. Ensure that you *do not* make this your first rule, since ACL rules are read from the top down and stop after the first rule match, which ignores all subsequent rules.

Logging ACL Activity

Overview

You can log information about packets that match specific rules in the active access control list (ACL). Each log entry contains the following information:

- Source and destination IP address.
- Protocol ID (RFC 1700 defines these ID numbers).
- Time that the match occurred.
- Index number of the access rule that was matched.

- Forwarding information:
  - The priority of the packet if the ACL assigns it a priority and forwards it.
  - `Forward with no change in priority` if the ACL forwards the packet without changing its priority.
  - `FILTER` if the ACL blocks the packet.

To avoid excessive consumption of the switch resources, ACL matches are logged for slow path traffic at specific time intervals. You can set the interval at which the matches are logged. The default setting is two seconds.

ACL matches are not logged for FIRE path or FORE path traffic or for the supervisor ethernet or console port. For more information about the slow path and FIRE and FORE paths, see "Routing Overview" in Chapter 1, "Introduction."

This section includes procedures for the following tasks:

- Enabling ACL Logging
- Setting a Logging Interval

Enabling ACL Logging

The default setting for ACL logging is disabled.

Web Agent Procedure

To enable ACL logging by using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Access Lists**.

   The IP Access List Web page is displayed in the content pane. See Figure 13-1.

2. Select the check box next to each rule that you want to enable ACL logging for.

3. In the **Rule Match Logging** column, select **Enable** for each access rule that you want to enable ACL logging for.

   The default setting for ACL logging is disabled.

4. Click **Apply**.

CLI Command

To *enable* ACL logging, use the following CLI command:

`(configure)#` **ip acl-logging enable <*access-list-name*> <*rule-number*>**

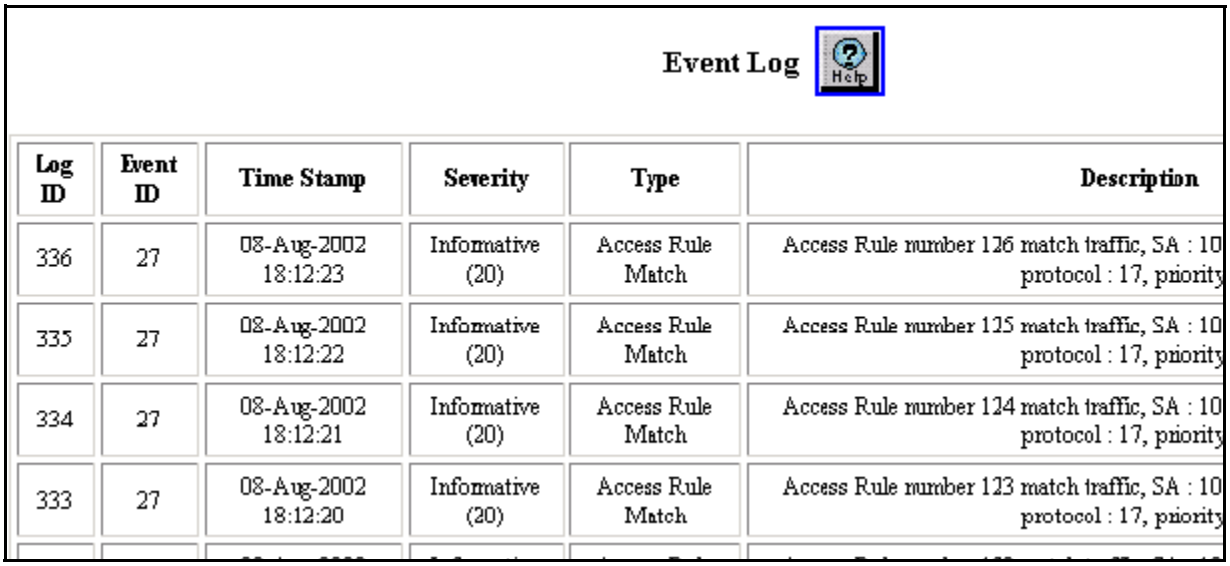
To *disable* ACL logging, use the following CLI command:

`(configure)#` **ip acl-logging disable <*access-list-name*> <*rule-number*>**

Sample Event Log Entries

Figure 13-5 displays sample ACL matches in the event log.

Figure 13-5. Sample ACL matches in the event log



Event Log

Log ID	Event ID	Time Stamp	Severity	Type	Description
336	27	08-Aug-2002 18:12:23	Informative (20)	Access Rule Match	Access Rule number 126 match traffic, SA : 10 protocol : 17, priority
335	27	08-Aug-2002 18:12:22	Informative (20)	Access Rule Match	Access Rule number 125 match traffic, SA : 10 protocol : 17, priority
334	27	08-Aug-2002 18:12:21	Informative (20)	Access Rule Match	Access Rule number 124 match traffic, SA : 10 protocol : 17, priority
333	27	08-Aug-2002 18:12:20	Informative (20)	Access Rule Match	Access Rule number 123 match traffic, SA : 10 protocol : 17, priority

Setting a Logging Interval

Web Agent Procedure

To set the interval for ACL logging:

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Access Control**.

   The IP Access Control Web page is displayed in the content pane. See Figure 13-4.

2. In the **Interval between logging of Access Rule Matches (in seconds)** field, enter the interval at which you want ACL matches logged. Enter an interval from 1 to 60 seconds.

   The default setting is 2 seconds.

3. Click **Apply**.

CLI Command

To set the interval for ACL logging, use the following CLI command:

`(configure)#` **ip acl-logging logging-interval <*time-in-seconds*>**

Optimizing Switch Performance

Overview

*** Note:** This section provides a detailed discussion of the architecture and functionality of the Avaya Multiservice switch with respect to ACLs. This material goes well beyond standard configuration issues by addressing system performance, memory management, and optimization.

Purpose

The purpose of this section is to explain the configuration options when using Access Lists. Deploying an Access List affects the use of hardware and software resources and may impact system performance. An Access Control List (ACL), also referred to as an Access List, is a tool for associating rules (permit, deny, prioritize for Quality of Service (QoS)) with identified IP traffic through the switch. This section will show how to monitor performance and adjust configurations to optimize performance.

Terminology

The following terms are used extensively in this section:

- **5-tuple**: The five elements that fully describe the criteria of the ACL rule: Source IP/ Mask, Destination IP/Mask, Protocol, Source Port, Destination Port. The masks allow the user to specify a narrow or wide range of matches. All elements are optional, but the Ports are only valid if TCP or UDP is the selected protocol and can be expressed as a single port or range of ports. The protocol ID for TCP and UDP is 6 and 17 respectively.
- **Access List/Access Control List** (ACL): an ordered list of ACL Rules.
- **ACL Rule**: An element of an ACL that identifies traffic based upon a 5-tuple (condition), and specifies a queue (0-7), permit, or deny action for packets matching the condition.
- **CPU**: The general processor for the P580 and P882 that resides on the Supervisor module. The CPU determines whether to forward or filter packets. It identifies Flows by resolving IP-to-MAC addressing and matching ACL Rules. It updates the F-chip's forwarding cache for future Fast Pathing of packets that match this Flow. The process of examining Flows and updating all of the F-chips' forwarding caches is called Slow Path.

- **DA**: Destination IP Address.
- **F-chip**: Forwarding-chip (F-chip) is an application specific integrated circuit (ASIC) that forwards recognized packets via Fast Path or unrecognized packets via Slow Path. It learns or flushes L3 forwarding entries (L3FE) as directed by the CPU. In earlier versions of the hardware, the F-Chip was referred to as the Packet Routing Engine (PRE).
- **Fast Path**: When an F-chip is able to recognize and forward a packet to the destination port without CPU intervention. Both FIRE (media modules) and FORE (supervisor module) are Fast Path forwarding mechanisms.
- **FIRE (Fast In-Band Routing Engine)**: When an 80-series Ethernet media module has been licensed for routing, the F-Chips on that module provide fast, direct forwarding of L3 packets to their destination ports. The destination ports can be on the same module or a different module.
- **Flow**: A set of routed packets that get the same Layer 3 Forwarding Entry (L3FE) due to equivalent address characteristics. The complexity of this equivalence will affect the frequency that packets match a single Flow. In the simplest case, such as no ACL present, a Flow is identified only by the destination address of the packet, for a high rate of L3FE re-use. In the most complex case, such as an ACL that specifies the entire 5-tuple, the granularity of identification will cause a very low rate of re-use.
- **FORE (Fast Out-of-Band Routing Engine)**: When an 80-series Ethernet media module is not licensed for routing, the media modules forward all routed packets to the supervisor. The F-Chip on the supervisor module provides fast, direct forwarding of L3 packets to their destination ports.
- **Hash Mode**: The F-chip uses a 12-bit key to initially locate available memory for Flows. In DA-only mode (default) it builds the key from a sample of the Flow's destination address (DA). In SA-DA mode, it uses both the source and destination addresses. The F-chip automatically selects which 12 bits to use, and when to re-hash with a different key.
- **Ingress F-chip**: The F-chip that receives a packet from an external source.
- **Layer 3 Forwarding Entry (L3FE)**: When a packet arrives on a media module, the F-Chip needs to know where to forward that packet. The L3 forwarding cache is a list that identifies the path taken by a data packet through the switch. Each entry in this list is a L3FE. Each entry is identified by a combination of the Flow's 5-tuple and the corresponding ACL rule. If the F-chip does not find a match in the L3 forwarding cache, the packet is sent Slow Path to

the CPU for processing and forwarding. Once the CPU has determined the destination, it updates the L3 forwarding cache on the F-chips with the L3FE. Once updated, the F-chip can forward future packets via Fast Path.

- **SA**: Source IP Address.
- **Slow Path**: When an ingress F-chip does not recognize a packet compared to its cache of known Flows, the packet is forwarded to the CPU to determine proper destination and ACL Rule assignment.

Techniques

You can use several techniques to optimize the switch performance when an access list is enabled. The techniques are related and must be considered together.

- Recognizing Performance Issues
- Evaluating System Performance
- Enabling Routing at the Module
- Designing Safe, Efficient ACLs
- Identifying the Ports
- Configuring Hash Mode
- Using Protocol or Port IDs in Access Rules
- Managing F-chip Memory

Recognizing Performance Issues

When the ACL is the root of a performance problem, it shows as the Slow Path becoming overused. The Slow Path is not designed to handle significant traffic levels since the single CPU also handles all other management functions. There are several ways to determine if the CPU is overloaded:

- Continuous PING to the supervisor: timeouts or inconsistent timing of echo responses.
- Slow Scrolling LED Marquee: This is good visual sign that the CPU is busy.
- Slow Management response: If Avaya Multiservice Network Manager (MSNM), Avaya Policy Manager (APM), HPOV, or a MIB browser get slow updates, this can signify a busy CPU or saturated network.
- Slow network response: This can be measured in a variety of ways.

Once you have determined that the ACL is the root cause, and have optimized the rules, you need to tune system parameters.

Evaluating System Performance

The first step is to determine whether the use of an ACL affects system performance. This requires an analysis of traffic patterns, the use of device statistics, and a “process of elimination”. If you suspect that the use of an ACL is degrading system performance and “Disabling Access Control” improves performance, you are ready to begin the process.

In order to perform the diagnosis, you need to generally identify routed Flows through the switch. You will need this information to match Flows against the F-chip statistics to locate the problem area.

The next step (and this will be repeated as needed) is to note the usage statistics with and without the ACL enabled.

Slow Path versus Fast Path

The goal is to maximize Fast Path traffic (L3 Frame Cache Hits) and minimize Slow Path traffic (L3 Slow Path Frames). The easiest parameters to monitor are found on the Forwarding Statistics Web page found on the switch Web Agent by selecting the following:

Routing > **L3 Forwarding Cache** > **Forwarding Statistics**

The Forwarding Statistics Web page shows the breakdown of traffic by FIRE Fabric port/Chip Index (see Figure 13-6).

Figure 13-6. Frame Forwarding Statistics Web Page

Frame Forwarding Statistics

Select	FIRE Fabric Port/Chip Index	L3 Total	L3 Frame Cache Hits	Percent Cache Hits	L3 Slow Path Frames	Percent Slow Path	L3 Drop Frames	Percent Drops	RX Frame Count (T2)	L2 Frame Count (T2)
☐	1/1	5,325,825	5,010,778	94%	315,047	5%	0	0%	n/a	n/a
☐	4/1	0	0	0%	0	0%	0	0%	n/a	n/a
☐	5/1	0	0	0%	0	0%	0	0%	n/a	n/a
☐	32/1	43,798	14,438	32%	29,360	67%	0	0%	n/a	n/a

CLEAR REFRESH

How Many Cache Entries?

Each F-chip/PRE (see "Identifying the Ports" later in this chapter for PRE column explanation) reserves space for IP unicast, IP multicast, and IPX. By default, the number of entries are limited to 15,000 for each (45,000 total). If the CPU attempts to add an entry in excess of 15,000, the L3FE will not be stored and that Flow will only be forwarded via the Slow Path. When looking at the Active FE Cache (Cache Contents) page, the **Current Entries** column indicates how many entries are currently in the L3F cache. Since the L3 forwarding cache is recalculated periodically, you need to monitor the FE Cache Web page on the switch Web Agent found by selecting the following:

Routing > **L3 Forwarding Cache** > **Cache Contents**

Layer 3 forwarding entries live for a time that is the maximum of two timers:

- The *Age Interval* for entries in the layer 3 forwarding cache (The default setting is 120 seconds. You can set this timer to a value from 120 to 360 seconds.)

  The Age interval timer starts when a new forwarding entry is added to the layer 3 forwarding cache. Each forwarding entry has its own timer.

- The *chassis timer* for layer 3 forwarding entries. This timer is an internal timer that you cannot change. The setting for this timer is dependent on the number and type of modules (number of F-Chips) in the chassis.

  The chassis timer specifies the number of seconds that the CPU takes to rebuild all forwarding caches on all F-chips in the chassis. Each second, the CPU removes expired forwarding entries from one protocol branch (IP-UNI, IP-Multi, or IPX) of the forwarding cache on one F-Chip and then rebuilds the branch.

  On a fully populated P882, rebuilding all forwarding caches on all F-chips takes 400 seconds. On a fully populated P580, this task takes 160 seconds.

The FE Cache page displays as shown below.

Figure 13-7. Active FE Cache Web page

Active FE Cache

Select	Fabric/Chip Index	Type	Mode	Memory Use in Bytes	Total Entries	Current Entries	Aged Entries	Duplicate Add Attempts	Failed Add Attempts	Entries Removed due to Route Deletes	Entries Removed due to Route Changes	Lookup Hits	Lookup Misses	Lookup Levels
☐	1/1	IP_UNI	DA-Only	134	389	1	387	0	0	0	0	1	0	1
☐	1/1	IP_MUL	SA-DA	0	0	0	0	0	0	0	0	0	0	0
☐	1/1	IPX	DA-Only	0	0	0	0	0	0	0	0	0	0	0
☐	4/1	IP_UNI	DA-Only	0	1	0	0	0	0	0	0	1	0	0
☐	4/1	IP_MUL	SA-DA	0	0	0	0	0	0	0	0	0	0	0
☐	4/1	IPX	DA-Only	0	0	0	0	0	0	0	0	0	0	0
☐	5/1	IP_UNI	DA-Only	0	1	0	0	0	0	0	0	1	0	0
☐	5/1	IP_MUL	SA-DA	0	0	0	0	0	0	0	0	0	0	0
☐	5/1	IPX	DA-Only	0	0	0	0	0	0	0	0	0	0	0
☐	32/1	IP_UNI	DA-Only	0	3	0	2	0	0	0	0	1	0	0
☐	32/1	IP_MUL	SA-DA	0	0	0	0	0	0	0	0	0	0	0
☐	32/1	IPX	DA-Only	0	0	0	0	0	0	0	0	0	0	0

Clear | Refresh | Flush Entry | Flush Table

Enabling Routing at the Module

The P580 and P882 have the option of centralized (unlicensed) or distributed (licensed) routing function. When a module is licensed, the on-board F-Chips (FIRE) provide the fast-path processing. When a module is unlicensed, the fast path is handled by the single F-Chip (FORE) on the Supervisor Module. Since the presence of an ACL will consume more F-Chip memory, it is recommended to license all modules in a chassis. This distributes the work load among more resources.

Designing Safe, Efficient ACLs

The entry of ACL rules via the CLI, Web or Avaya Policy Manager does not encourage or enforce any checking beyond correct syntax. The general guideline is that you are configuring a Layer-3 switch, not a firewall. The following are some guidelines for designing safe, efficient ACLs and how they affect performance:

- **Specify both source *and* destination address whenever possible.**

  The wildcard feature is convenient but can dramatically increase the number of flows that the switch identifies. Since the standard ACL implies “any” for the destination, use standard ACLs with care. The wildcard should match a specific set of addresses.

- **Use Protocols/Ports Carefully.**

  Pushing the ACL-to-packet matching up one or two levels of the IP stack refines the granularity of the flows to be very specific in what is matched. A source-port range can cause a large number of “micro” flows to be created. For more information on using protocol and port identifiers in access rules, see Configuring Hash Mode.

- **Do not use ACLs to block protocol or port routing through the switch.**

  You can, however, use ACLs to block protocol or port access to specific interfaces on the switch. For more information, see Configuring Hash Mode.

- **Minimize Rules.**

  The number of rules has a direct impact on the CPU effort to match rules to Flows. This is especially true when there is a high frequency of packets that are “walked down” the entire list and don’t match any rules.

- **Minimize Searching.**

  The goal is to place the most frequently matched rules toward the beginning of the ACL. This requires a good knowledge of traffic patterns. This can be noticeable as ACLs get longer.

- **Permit Management Traffic with High Priority.**

  This include routing updates (unicast for RIP 1, multicast for RIP 2), SNMP (MSNM, HPOV), LDAP (for Avaya Policy Manager). Not doing this can cause loss of management connectivity.

Identifying the Ports

The chassis is organized by slots, fabric ports, PREs or F-chips, and physical ports. The number of F-Chips and physical ports varies according to the module type. This information helps you distribute the workload evenly among resources and identify possible choke points:

- Every Fabric port can manage up to 4 F-Chips
- Slot 1 has 1 Fabric port only
- Slots 2-7 (P580) and Slots 2-17 (P882) have 2 Fabric ports per slot
- The Supervisor has 1 F-chip (FORE)
- The 8-port GigE has 8 F-Chips (4 per Fabric port)
- The 4-port GigE has 4 F-Chips (2 per Fabric port)
- The 24-port Ethernet modules (copper or fiber) have 2 F-Chips (1 per Fabric port). Physical Ports 1-12 correspond with F-Chip 1, and Physical Ports 13-24 correspond with F-Chip 2.
- The 48-port Ethernet has 4 F-Chips (2 per Fabric port) with the following Physical Port to F-Chip correspondence: ports 1-12: F-Chip 1, ports 13-24: F-Chip 2, ports 25-36: F-Chip 3, ports 37-48: F-Chip 4.

Fabric ports are numbered regardless of whether other slots are empty or full.

F-Chips numbers are associated with their respective Fabric ports. To locate the Fabric port and F-Chip for a physical port, you must know the media type and slot.

For example, Physical port 20 on a 24-port Ethernet module that is in slot 4 of the chassis is identified by Fabric Port 7 and F-Chip 2. This will be displayed as Chip 7/2. These absolute addresses are not affected by the placement or type of other modules.

Table 13-3 shows the slots, fabric ports, PREs or F-chips, and physical ports of the switch.

Table 13-3. Fabric Port and F-Chip Assignments

Slot	Fabric Port	F-Chips		Fabric Port	F-Chips	Comments
1	1	1				Supervisor has 1 F-Chip
2	2	2/1-2/4		3	3/1-3/4	
3	4	4/1-4/4		5	5/1-5/4	
4	6	6/1-6/4		7	7/1-7/4	
5	8	8/1-8/4		9	9/1-9/4	
6	10	10/1-10/4		11	11/1-11/4	
7	12	12/1-12/4		13	13/1-13/4	Last slot of a P580
8	14	14/1-14/4		15	15/1-15/4	
9	16	16/1-16/4		17	17/1-17/4	
10	18	18/1-18/4		19	19/1-19/4	
11	20	20/1-20/4		21	21/1-21/4	
12	22	22/1-22/4		23	23/1-23/4	
13	24	24/1-24/4		25	25/1-25/4	
14	26	26/1-26/4		27	27/1-27/4	
15	28	28/1-28/4		29	29/1-29/4	
16	30	30/1-30/4		31	31/1-31/4	
17	32	32/1-32/4		33	33/1-33/4	Last slot of a P882

Configuring Hash Mode

If your diagnosis shows a performance degradation with an ACL enabled, IP Unicast Hash mode may need to be changed from **DA-only** to **SA-DA**.

When you enable an ACL, the switch automatically sets hash mode for IP unicast traffic to **SA-DA**. Avaya recommends that you not change this setting unless changing it to **DA-only** greatly reduces the number of forwarding entries. For more information on the interrelationship of hash mode and specific ACL configurations, see Using Protocol or Port IDs in Access Rules.

SA-DA mode can cause a noticeable increase in the total flows identified and result in an increased usage of F-chip memory.

See "Managing F-chip Memory" in this case.

***Important:** Changing the **Hash Mode** setting affects every L3-enabled F-chip on the P580 or P882.

Web Agent Procedure

To use the Web Agent to manually change the IP unicast hash mode from **DA-only** to **SA-DA**:

1. In the navigation pane, expand the **Routing** > **L3 Forwarding Cache** > **Cache Configuration**. The Layer-3 Forwarding Cache Configuration Web page is displayed in the content pane.
2. In the **Hash Mode** field for **IP Unicast** traffic, select **SA-DA**.
3. Click **APPLY**.

CLI Command

To use the CLI to manually change the IP unicast hash mode from **DA-only** to **SA-DA**, enter the following command:

```
(configure)# ip unicast route-cache hash-mode sa-da
```

Using Protocol or Port IDs in Access Rules

Potential Hashing Issues

Using a port or protocol identifier in an access rule can cause the switch to add many entries to the forwarding cache when traffic between two endpoints includes many flows. The extra entries resulting from the port or protocol identifier hash to the same locations in the forwarding cache because they have the same source and destination address.

For safe, efficient ACLs, do *not* use:

- SA and DA wildcards with any protocol or port identifiers.
- DA wildcard with any protocol or port identifiers.

Be very careful if you use a source wildcard and single destination with protocol or port identifiers. This configuration works for local interface addresses, since all packets destined to local interfaces are forwarded to the slow-path anyway. The interface simply compares the packets to the ACL before processing them and forwarding them to the supervisor.

However, if the destination specified in the access rule is a network host and many simultaneous flows exist, switch performance can degrade. This performance degradation occurs because the switch must generate a large number of forwarding cache entries for the simultaneous flows to further

differentiate packets by protocol and port. These entries all hash to the same value because they have the same source and destination address, and you may observe a degradation of the switch performance.

Example

To block SNMP access to the supervisor from the network, on IP interface 10.10.0.240/255.255.255.0, use the following ACL entry:

ip access-list SNMP 10 deny udp any host 10.10.0.240 eq 161

If you were to use the following command, the switch would block all inter-subnet SNMP traffic, but would also create a forwarding cache entry for every flow that had a different SA, DA, source port, destination port, or protocol.

ip access-list SNMP 10 deny udp any any eq 161

Interrelation with Hash Mode Setting

Using DA-only hashing generally reduces the overall number of forwarding entries, but it can cause performance issues if used when an ACL is enabled. These performance issues are magnified when the ACL uses protocol and port identifiers.

An ACL that specifies a source address, protocol ID, or port ID requires closer analysis of packets than just the destination address. Every flow to the destination needs its own forwarding cache entry based on the ACL criteria, and all of the entries hash to the same value. In this scenario, the switch must sequentially search every entry in the forwarding cache that has the same DA (thus hash-location).

When you set the hash mode to SA-DA, each different source-destination combination hashes to a different value. Thus the number of entries hashed to a single value significantly decreases. However, SA-DA can also cause performance issues in some situations. If many entries that do not match the ACL have similar hash values to those that do, DA-only hashing provides more efficient usage of the forwarding memory.

Managing F-chip Memory

The reconfiguration of Hash Mode can cause a secondary effect: increased cache usage. By default, the IP Unicast Cache size is 15,000 entries per F-chip. Although this can be used up simply due to a high number of flows (for example, a proxy server for the internet), the SA-DA Hash Mode setting always causes more flows to be identified than in the DA-only mode.

The F-chip memory can accommodate approximately 70,000 total entries for routed (L3) flows. This number comprises IP Unicast, IP Multicast, and IPX entries for that F-chip.

To view the current total number of entries for the CPU, expand the **Routing** > **L3 Forwarding Cache** folders, and then click **Cache Configuration**. The **Total Current Entries** field displays the current total entries for the CPU.

To view the current total number of entries for each F-chip, expand the **Routing** > **L3 Forwarding Cache** folders, and then click **Cache Contents**. The Active FE Web page is displayed in the content pane. See Figure 13-7.

As long as these totals stay under 70K, it is safe to increase the IP Unicast Maximum Entries to prevent overflow. If the switch is not routing IPX or has minimal IP Multicast traffic, it is generally safe to double the IP Unicast maximum to 30,000.

When you enable an ACL, the switch may change the maximum number of forwarding entries for IP unicast traffic to improve the performance of the switch. If the maximum number of entries is set to:

- Less than 15,000, the switch automatically doubles the setting.
- Between 15,000 and 30,000, the switch changes the setting to 30,000.
- Greater than 30,000, the switch does not change the setting.

***Important:** Changing the Maximum Cache Entries parameter affects every L3-enabled F-chip on the P580/P882.

Web Agent Procedure

To use the Web Agent to manually double the IP Unicast maximum to 30,000 entries:

1. In the navigation pane, expand the **Routing** > **L3 Forwarding Cache**, and then click **Cache Configuration**. The Layer-3 Forwarding Cache Configuration Web page is displayed in the content pane.
2. In the **Maximum Entries** field for **IP Unicast** traffic, enter **30000**.
3. Click **APPLY**.

CLI Command

To use the CLI to manually double the IP Unicast maximum to 30,000 entries, enter the following CLI command:

```
(configure)# ip unicast route-cache max-size 30000
```

14 Configuring RIP Routing

Overview

*** Note:** To globally enable the Routing Information Protocol (RIP), see Chapter 12, "Configuring IP Routing."

The information and procedures provided in this chapter pertain to 50-series layer 3, and licensed 80-series module configurations only:

This section describes:

- Configuring RIP on the Switch
- Modifying RIP Interfaces
- Creating Trusted RIP Neighbors
- Key Chains
- Viewing RIP Statistics
- NBMA IP Interfaces

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Configuring RIP on the Switch

You can configure RIP from either the Web Agent or the CLI.

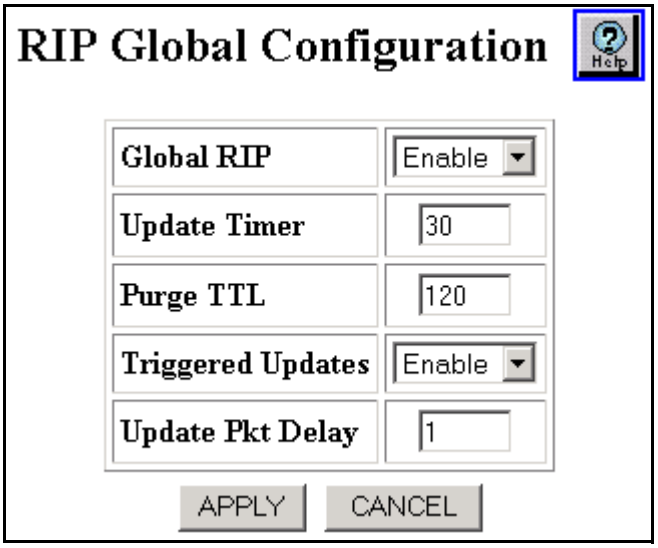
Web Agent Procedure

To configure the RIP globally on your switch using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **RIP** folders, and then click **Global Configuration**.

   The RIP Global Configuration Web page is displayed in the content pane. See Figure 14-1.

Figure 14-1. RIP Global Configuration Web Page



2. See Table 14-1 to configure the RIP Global Configuration Web page parameters

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 14-1. RIP Global Configuration Web Page Parameters

Parameter	Allows you to...
Global RIP	Enable or disable the RIP protocol. The default is **Enable.** **Note:** You must disable the IP Interface before Disabling or Enabling RIP.
Update Timer	Enter the value (in seconds) that represents the time between RIP updates on all interfaces. The parameter range setting is 10-50 seconds. The default value is 30 seconds.
Purge TTL (Garbage Collection Timer)	Enter the “garbage-collection” time. Upon expiration of the time-out, the route is no longer valid, however, it is retained in the routing table for a short time so that neighbors can be notified that the route has been dropped.Global Time To Live (TTL) in seconds that the RIP update persists. The default value is 120 seconds.The parameter range setting is 1-65,535 seconds.
	1 of 2

Table 14-1. RIP Global Configuration Web Page Parameters

Parameter	Allows you to...
Triggered Updates	**Disable** or **Enable** route updates that were triggered by a topology change in the network to be added to the routing table. The default is **Enable**.
Update Pkt Delay	Enter the value (in seconds) that represents the time delay between successive RIP update packets to the neighbor, when the update requires multiple packets. The default value is 1 second. The parameter range setting is 0-50 seconds.
	2 of 2

CLI Commands

To configure or modify the RIP parameters globally on your switch using the CLI, enter the following commands in router:rip mode:

- `(configure router:rip)#` **timers basic <*update*> <*invalid*>**
- `(configure router:rip)#` **output-delay <*delay*>**
- `(configure router:rip)#` **triggered updates**

Modifying RIP Interfaces

You can modify RIP interfaces from either the Web Agent or the CLI.

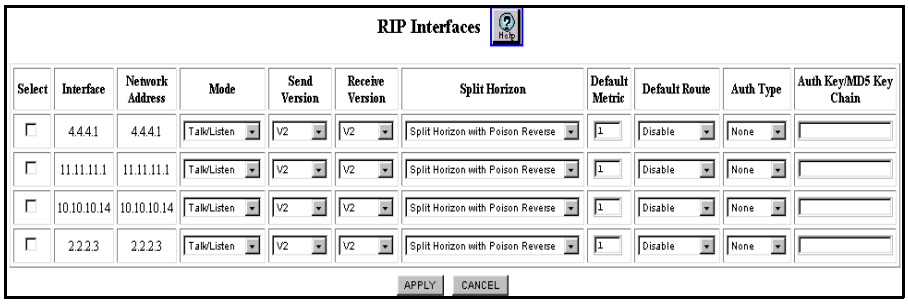
Web Agent Procedure

To modify RIP interfaces using the CLI from the Web Agent:

1. In the navigation pane, expand the **Routing > IP > RIP** folders, and then click **Interfaces**.

   The RIP Interfaces Web page is displayed in the content pane. See Figure 14-2.

Figure 14-2. RIP Interfaces Web Page



2. See Table 14-2 to modify the RIP interfaces Web page parameters.

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 14-2. RIP Interface Web Page Parameters

Parameter	Allows you to...
Select	Select the RIP interface to be modified.
Interface	Displays an interface from a list of interfaces that you previously configured.
Network Address	Displays an IP address to be associated with the displayed interface.
Mode	Specify the RIP State. Different states include: • **talk only** (only send RIP information on this interface) • **listen only** (only receive RIP information on this interface) • **talk/listen** (both send and receive RIP information on this interface) **Talk/Listen** is the default setting.
Send Version	Specify the version of RIP you want to use to send packets across this interface. Selections include: • V1 • V2 • V1/V2 **V1** is the default setting
Receive Version	Specify the version of RIP you want to use to receive packets. Selections include: • V1 • V2 • V1/V2 **V1** is the default setting
Split Horizon	Specify that IP routes learned from an immediate neighbor are not advertised back to the neighbor from which the routes were learned.Selections include: • **Split Horizon** - Routes that leave this interface, will not be advertised back on the same interface. • **Split Horizon with Poison Reverse** - Routes that leave this interface, will be advertised back on the same interface with an infinite metric (16). **Split Horizon with Poison Reverse** is the default setting.
	1 of 2

Table 14-2. RIP Interface Web Page Parameters

Parameter	Allows you to...
Default Metric	Specify the RIP route metric value. Costing metric used when advertising the RIP route on this interface. **1** is the default setting. The parameter range is 0-15.
Default Route	Specify the mode for the default route learning on this interface. Different states include: • **Disable** - Disables the default route. • **Talk Only** - Send RIP information to the network. • **Listen Only** - Receive RIP information from the network. • **Talk/Listen** - Both send and receive RIP information from the network. **Disable** is the default setting.
Auth Type	Specify the type of authentication available for use on a given RIP interface. Authentication types include: • **None** - No authentication required. • **Simple** - Uses a clear-text password for validation. This password is sent unencrypted across the network to neighboring RIP routers.Available for RIP V2 only. • **MD5** - Uses an encrypted **Key** to validate RIP V2 routing updates from neighboring RIP V2 routers. **None** is the default setting
Auth Key	If **Auth Type** is: • **None** - leave this field blank, or this field is ignored. • **Simple** - Enter the clear-text password for validating RIP packets from neighbor RIP V2 routers.(1-16 characters) • **MD5** - Enter the defined **Key Chain** values under **Routing** > **IP** > **RIP** > **Key Chains**. **None** is the default setting.
	2 of 2

CLI Command

To modify RIP interfaces parameters using the CLI enter the following command from Configure mode:

`(configure)#` **interface <*intf-name*> | [type {mbna | ethernet}]**

Creating Trusted RIP Neighbors

Trusted RIP neighbors enable you to enhance the switch's security by enabling one or more neighbors to relay RIP information.

*** Note:** Adding one or more trusted RIP neighbors ensures that your router accepts only information from these neighbors. Consequently, all other information is filtered. Do not create trusted RIP neighbor(s) if you do not wish to filter RIP information from the network.

You can create a trusted neighbor using either the Web Agent or the CLI.

Web Agent Procedure

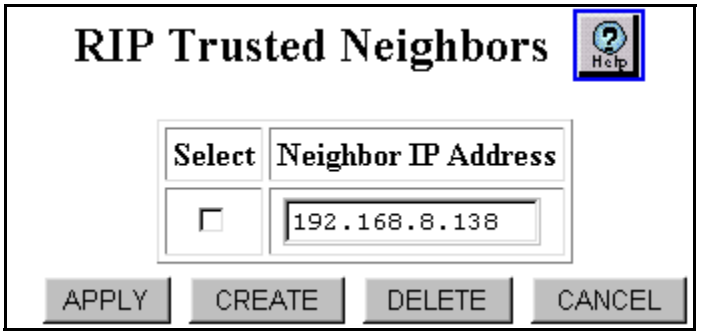
To create a trusted RIP neighbor using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **RIP** folders, and then click **Trusted Neighbors**.

   The RIP Trusted Neighbors Web page is displayed in the content pane. See Figure 14-3.

2. If a RIP Trusted Neighbor hasn't been configured, only the **CREATE** button appears in Trusted Neighbor Web page.

Figure 14-3. RIP Trusted Neighbors Web Page



3. Click **CREATE**. The Add Trusted RIP Neighbor Web page is displayed. See Figure 14-4.

4. Click on the Select column checkbox for the Trusted Neighbor and click on **APPLY** to save your changes

5. Click on the Select column checkbox for the Trusted Neighbor and click on **DELETE** to delete the entry or **CANCEL** to restore previous settings.

Figure 14-4. Add Trusted RIP Neighbors Web Page



6. Enter the IP address for the node that acts as the trusted RIP neighbor in Network Address field.

7. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To create a trusted RIP neighbor, use the following command:

`(configure router:rip)#` **neighbor** <***ip-addr***>

Key Chains

Overview

In a full implementation of Key Chains, a **Key Chain** a set of Keys each with its own set of parameters used for MD5 encryption. The parameters are an encryption Key and the period of time that the key is valid. This key must be identical on each device that will participate in an exchange of information. In the case of RIP (version 2 only), routers can use MD5 encryption to ensure that only routing information packets (RIP packets) from authorized routers are accepted.

In a simple implementation, the user creates one Key in a Key Chain and this key is used until it expires.In an advanced implementation, a Key Chain contains multiple keys each valid for a limited amount of time. This provides a periodic change of the encryption key giving the highest level of security. The drawback to this approach is that every device participating in an exchange, such as RIP V2 routers on the same subnet, must have their clocks synchronized. Failure to synchronize would cause devices to use different keys and hence reject RIP updates from neighboring RIP routers.

Key Chains on the Avaya Multiservice Switch

The Avaya Multiservice switch supports the creation of a single Key Chain with a single Key. Hence, Key and Key Chain are interchangeable terms. The Key has a setting for a start time and duration. Regardless of duration, the Key's start time must be set, ***even if*** the duration is set to infinity.

If the duration is less than 4294967295 seconds, the Key will expire. Ironically, this number is approximately 136 years. A more realistic number would be 18000 seconds (5 Hours). For this time interval, the RIP V2 interface would exchange routes with other RIP V2 routers in the same subnet. Once expired, the RIP V2 interface would no longer accept RIP packets from RIP V2 routers on that subnet.

If the Key's duration is 4294967295 (alternately -1), the key never expires. Again, if you use this value, you must still specify a Key Accept Time. Failure to do so will result in the Avaya switch not creating the Key.

Configuring Key Chains

Web Agent Procedure

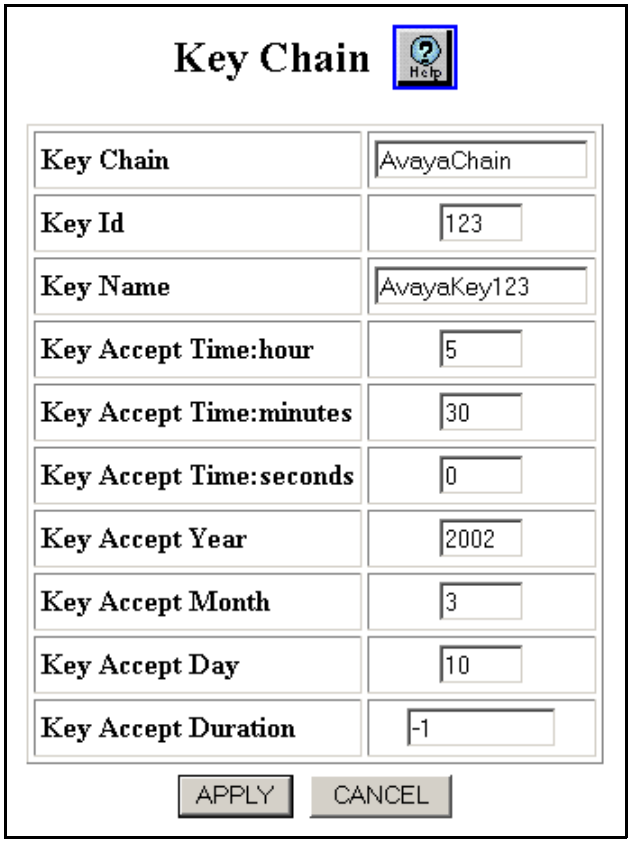
You can configure Key Chains using the Web Agent or the CLI.

To create a Key Chain/Key using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > RIP** folders, and then click **Key Chain**.

   The Key Chain Web page is displayed in the content pane. See Figure 14-5.

Figure 14-5. Key Chain Web Page



2. Configure the key chain as appropriate. See Table 14-3 for an explanation of each field.

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 14-3. Key Chain Parameters

Parameter	Defines the...
Key Chain	1-16 character ASCII name of the Key Chain
Key ID	Three digit number of the Key. Valid range 1 to 255.
Key Name	16 character alpha-numeric key. This is the actual key used by MD5 encryption. This setting must be identical to the Key on other RIP V2 routers on the same subnet. To be compatible with Cisco devices, do not use a digit (0 - 9) as the first character.
Key Accept Time:hour	The hour(1-23) this key becomes valid
Key Accept Time:minutes	The minute(0 - 59) this key becomes valid
	1 of 2

Table 14-3. Key Chain Parameters

Parameter	Defines the...
Key Accept Time:seconds	The second(0 - 59) this key becomes valid
Key Accept Year	The year (2000 - 2009) this key becomes valid
Key Accept Month	The month (1 - 12) this key becomes valid
Key Accept Day	The day (1 - 31) this key becomes valid. If you enter a day that does not exist in a month such as the 31st of February, the Key will not be created.
Key Accept Duration	The duration in seconds that this key is valid. Enter -1 if the key is to last indefinitely.
	2 of 2

CLI Commands

To assign or modify Key Chain parameters using the CLI, enter the following commands from router:rip mode:

- To create the Key Chain:

  `(configure router:rip)#` **key chain <*Key Chain*>**

- To set the Key Chain ID:

  `(configure router:rip)#` **key <*Key ID*>**

- To set the Key Name (MD5 Key):

  `(configure router:rip)#` **key-string <*Key Name*>**

- To set the Key Chain's start time and duration:

  `(configure router:rip)#` **accept-lifetime <*hh:mm:ss*> <*day*> <*month*> <*year*> duration {<*time*> | infinite}**

CLI Example

To create a **Key Chain** called *AvayaChain* with **Key ID** *123*, **Key Name (**MD5 Key) *AvayaKey123* that begins on March 10, 2002 at 5:30 a.m. and lasts indefinitely, enter the commands in the following order:

1. **key chain AvayaChain**
2. **key 123**
3. **key-string AvayaKey123**
4. **accept-lifetime 05:30:00 10 march 2002 duration infinite**

Viewing RIP Statistics

You can view RIP statistics using either the Web Agent or the CLI.

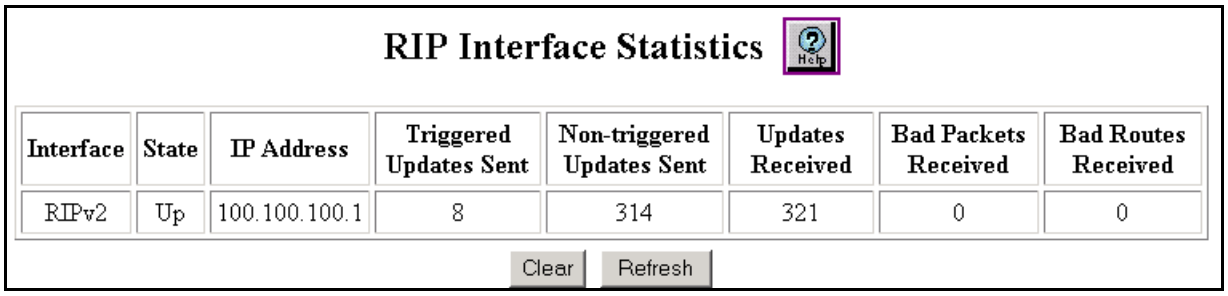
Web Agent Procedure

To view RIP statistics using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **RIP** folders, and then click **Interface Statistics**.

   The RIP Interface Statistics Web page is displayed in the content pane. See Figure 14-6.

Figure 14-6. RIP Interface Statistics Web Page



2. To modify your RIP statistics:

   — Click CLEAR *to reset all the entries*.

   — Click REFRESH to receive the most up-to-date information on the entries.

3. Use Table 14-4 to interpret the RIP statistics:

Table 14-4. RIP Statistical Parameters

Parameter	Defines the...
Interface	Interface associated with the IP address specified.
State	Current status of the RIP interface. UP indicates that the interface is up and RIP can transmit and receive updates.
IP Address	IP address associated with the interface.
Triggered Updates Sent	Number of RIP triggered updates sent.
Non-Triggered Updates Sent	Number of RIP non-triggered updates sent.
	1 of 2

Table 14-4. RIP Statistical Parameters

Parameter	Defines the...
Updates Received	Number of RIP updates received based on route changes in the IP routing table.
Bad Packets Received	Number of bad packets received on this interface.
Bad Routes Received	Number of bad routes received on this interface.
	2 of 2

CLI Command

To view RIP statistics using the CLI, enter the following command from Configure mode:

```
(configure)# show ip rip statistics
```

NBMA IP Interfaces

Non-broadcast multi-access (NBMA) functionality was added to RIP and OSPF routing protocols on the Avaya Multiservice switch software. Support for this functionality was added to enable route exchange over nonbridged connections (routed PVCs).For more information about NBMA, see "NBMA IP Interfaces" in Chapter 15, "Configuring the OSPF Routing Protocol".



CAUTION:

The NBMA feature is not backward compatible. If you have NBMA interfaces set up and need to downgrade from v6.1 to a 5.2 version of software, contact technical support.

15 Configuring the OSPF Routing Protocol

Overview

The following information and procedures provided in this chapter pertain to layer 3 module configuration only:

- What is OSPF?
- Configuring OSPF
- Monitoring Switch Performance Using OSPF Statistics
- NBMA IP Interfaces

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

What is OSPF?

OSPF (Open Shortest Path First) is a link state networking protocol. Each router broadcasts a packet that describes it own local links. Routers collect information from these broadcast packets to build their own network routing tables. These packets that describe the local links are short and cause less traffic congestion than Routing Information Protocol (RIP), which broadcasts large routing tables.

OSPF bases its routing decision on the least-cost path. The cost is administered value, usually based on line speed. If there are multiple areas in an OSPF domain, there must be a backbone area, identified as **area 0**. When areas are configure in the OSPF domain, there are four basic router classifications:

- Internal Router — router with all directly connected networks belonging to the same area. Routers with only interfaces in the backbone area also belong to this category and have a single link-state database.
- Area Border Router (ABR) — router that has directly connected networks belonging to multiple areas. ABRs have multiple link-state databases, one for each area, including the backbone. ABRs

summarize the networks in their areas and advertise them onto the backbone area. The backbone, in turn, distributes the information to the other areas.

- **Backbone Router** — a router that has at lease one directly connected network in the backbone area. This includes all ABRs. However, not all backbone routers must be ABRs. A backbone router that has all its interfaces in the backbone would be an internal backbone router.networks in their areas and advertise them onto the backbone
- **Autonomous System Backbone Router (ASBR)** — router that has directly connected interfaces in non-OSPF networks. These networks are then brought into an OSPF domain. An ASBR is independent of the other router classifications.

Configuring OSPF

This section provides the following information about OSPF:

- Configuring the OSPF Global Configuration
- Creating OSPF Areas
- Deleting OSPF Areas
- Modifying OSPF Areas
- Enabling OSPF on an IP Interface
- Modifying OSPF Interfaces
- Configuring an OSPF Passive-Interface
- Creating OSPF Virtual Links
- Deleting OSPF Virtual Links
- Modifying OSPF Virtual Links
- Creating OSPF Summaries
- Deleting OSPF Summaries
- Modifying OSPF Summaries

Configuring the OSPF Global Configuration

The OSPF global configuration allows you to globally configure OSPF on your switch. It also allows you to specify your router ID and whether or not you want the switch to be the Autonomous System (AS) border router.

You can globally configure OSPF using either the Web Agent or the CLI.

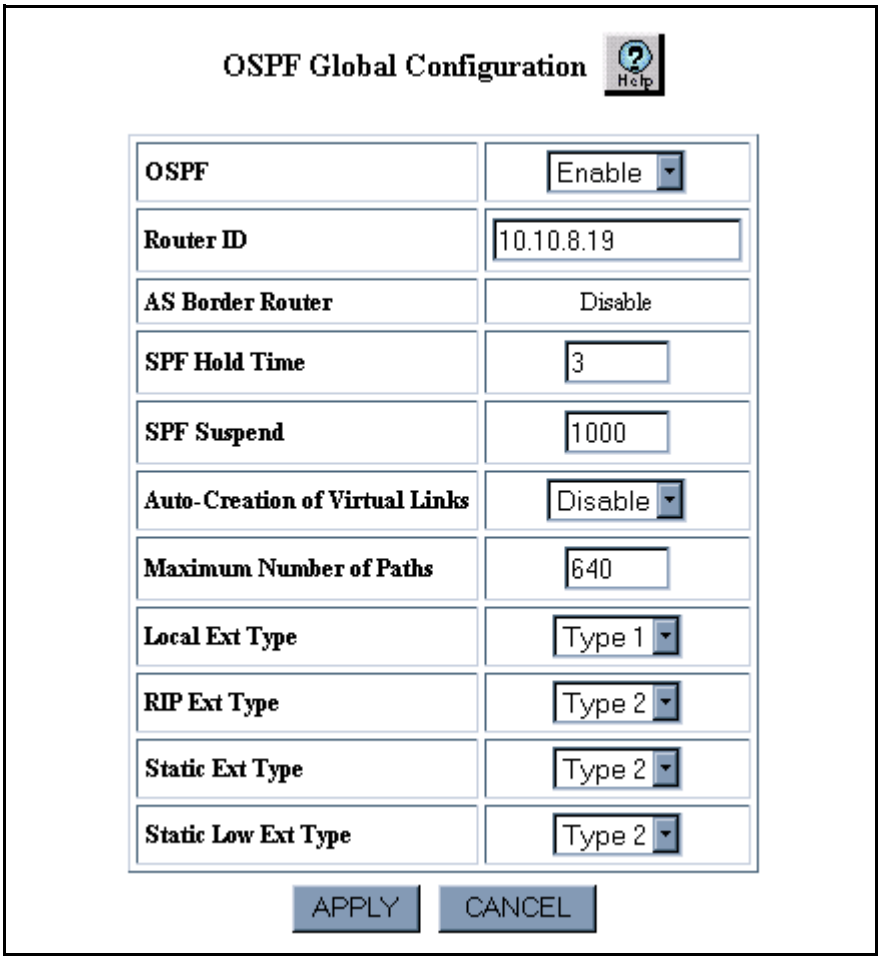
Web Agent Procedure

To globally configure OSPF using the Web Agent:

1. In the navigation pane, expand the **System** > **IP** > **OSPF** folders, and then click **Global Configuration**.

   The OSPF Global Configuration Web page is displayed in the content pane. See Figure 15-1.

Figure 15-1. OSPF Global Configuration Web Page



2. See Table 15-1 to configure the OSPF Global Configuration Web page parameters.

3. Click **APPLY** to save changes, or **CANCEL** to restore previous settings.

Table 15-1. OSPF Global Configuration Web Page Parameters

Parameter	Explanation
OSPF	Select to enable or disable OSPF globally on your switch.
Router ID	Specify the Router ID on the switch. The router ID is a 32-bit number assigned to each router running OSPF. This number uniquely identifies the router within an Autonomous System. If 0.0.0.0 is used, the router uses the IP address of an interface.
AS Border Router	Displays the Autonomous System Border Router (ASBR) status. The switch automatically detects the ASBR status: • If route redistribution filters are configured for OSPF, the ASBR status is **Enable**. • If all interfaces on the switch are in an OSPF stub area, the ASBR status is **Disable**, regardless of whether route redistribution filters are configured. For more information on route redistribution filters, see "Configuring Route Redistribution" in Chapter 12, "Configuring IP Routing."
SPF Hold Time	Specify the minimum number of seconds between SPF (shortest path first) runs.
SPF Suspend	Specify the number of nodes to process SPF runs before suspending.
Auto-Creation of Virtual Links	Enable or disable the function of automating the creation of virtual circuits based on network topology. This feature is Avaya proprietary and only auto creates virtual links on the config. notes.
Maximum Number of Paths	Configure the maximum number of paths used when running OSPF.
Local Ext Type	Specify whether imported local routes are advertised in OSPF with type 1 (internal) or type 2 (external) metrics.
RIP Ext Type	Specify whether imported RIP routes are advertised in OSPF with type 1 (internal) or type 2 (external) metrics.
	1 of 2

Table 15-1. OSPF Global Configuration Web Page Parameters

Parameter	Explanation
Static Ext Type	Specify whether imported high preference static routes are advertised in OSPF with type 1 (internal) or type 2 (external) metrics.
Static Low Ext Type	Specify whether imported low preference static routes are advertised in OSPF with type 1 (internal) or type 2 (external) metrics.
	2 of 2

CLI Command

To globally configure OSPF, use the following command:

```
(configure)# router ospf
```

Creating OSPF Areas

You can create OSPF areas using either the Web Agent or the CLI.

Web Agent Procedure

To create OSPF areas using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **OSPF** folders, and then click **Areas**.

   The OSPF Areas Web page is displayed in the content frame. See Figure 15-2.

*** Note:** The Area ID 0.0.0.0 is a backbone area and always exists in an OSPF configuration.

Figure 15-2. OSPF Areas Web Page



2. Click **CREATE**. The Add OSPF Area Web page is displayed. See Figure 15-3.

Figure 15-3. ADD OSPF Area Web Page



3. See Table 15-2 to configure the Add OSPF Area Web page parameters.

4. Click **CREATE** to save your changes, or **CANCEL** to restore previous settings.

Table 15-2. Add OSPF Area Web Page Parameters

Parameter	Allows you to...
Area ID	Specify the Area ID (32-bit character) for the new area. This must be a unique ID within AS. **Note:** Do not use 0.0.0.0 as an area ID.
Area Type	Select the type of area. Types include: • **Non-Stub** - Non-edge device/router. • **Stub** - An edge device/router that does not leak external advertisements. • **Not-so-stubby** - This is still a stub area, however, this device/router can leak some external advertisements.
Translate 7 into 5	**Enable** or **Disable** the translation of the NSSA ASE Type 7 into an AS External LSA Type 5.
Stub Metric	Specify the stub area default summary cost metric. (Default is 1).
Type 3 ASE Filter	Select to enable or disable the Type 3 summary LSA filter for Stub and NSSA only.

CLI Command

To create OSPF areas, use the following CLI command:

`(configure router:ospf)#` **area** <***area-id***> <***option***>

*** Note:** If you use the **network area** CLI command to add an interface to an OSPF area, be careful not to enter the IP mask for the <*wildcard-mask*> variable. The <*wildcard-mask*> variable is the inverse of a network mask. Enter a 32-bit number in four-part, dotted decimal format. Place ones in the bit positions that you want to mask.

For example, to enable OSPF on interface 10.10.10.1 (where the network mask is 255.255.255.0) and assign it to area 2.2.2.2, enter: **network 10.10.10.1 0.0.0.255 area 2.2.2.2**.

If you enter the IP network mask (255.255.255.0) instead of the wildcard mask (0.0.0.255), all OSPF interfaces that have 1 for the last octet of their IP address will be added to area 2.2.2.2.

Deleting OSPF Areas

You can delete an OSPF area using either the Web Agent or the CLI.

Web Agent Procedure

To delete an OSPF area using the Web Agent:

*** Note:** Before deleting an OSPF area, assign any interfaces that are associated with the area to a different area or delete the interfaces.

1. In the navigation pane, expand the **Routing** > **IP** > **OSPF** folders, and then click **Areas**.

   The OSPF Area Web page is displayed in the content pane. See Figure 15-2.

2. Select the OSPF Area to be removed from the **Select** field.

3. Click **DELETE**. The OSPF area you selected is deleted.

CLI Command

To delete an OSPF area, use the following command:

*** Note:** Before deleting an OSPF area, assign any interfaces that are associated with the area to a different area or delete the interfaces.

`(configure router:ospf)#` **no area** <***area-id***>

Modifying OSPF Areas

You can modify the parameters of an OSPF area using either the Web Agent or the CLI.

Web Agent Procedure

To modify an OSPF area using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **OSPF** folders, and then click **Areas**.

   The OSPF Areas Web page is displayed in the content pane. See Figure 15-2.

2. Select an OSPF Area to be modified from the **Select** field.

3. See Table 15-2 and modify the OSPF Area Web page parameters.

4. Click **Apply** to save your changes, or **Cancel** to ignore your changes.

CLI Command

To modify an OSPF area using the CLI, use the following command:

`(configure router:ospf)#` **area** <***area-id***> <***option***>

Enabling OSPF on an IP Interface

You can create an OSPF interface using either the Web Agent or the CLI after you setup a VLAN and create an IP interface.

Web Agent Procedure

After you have set up a VLAN, you must create the IP interface that your VLAN and your subnet will use to communicate. While creating the IP interface, you must assign it to the VLAN. The last step is to enable IP forwarding (Routing) global parameters.

To enable OSPF on the IP interface:

1. Administratively bring down the IP interface. Select the down arrow from the **Admin. State** field and select **DOWN**.

   ***Note:** You must bring the IP interface Down to successfully enable OSPF on the IP interface. If you do not, you will get an error message and OSPF will not be enabled.

2. Select the down arrow from the **OSPF** field.

3. Select **Enable**.

4. Administratively bring the IP interface **UP**.

Modifying OSPF Interfaces

You can modify the parameters on an OSPF interface using either the Web Agent or the CLI.

Web Agent Procedure

To modify an OSPF interface using the Web Agent:

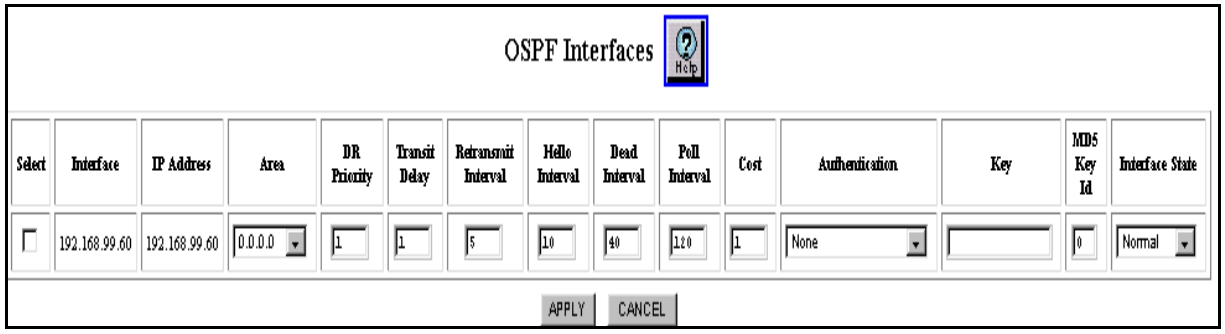
1. In the navigation pane, expand **Routing** > **IP** > **Configuration** folders.
2. Click **Interfaces**.

   The IP Interfaces Web page is displayed in the content pane.
3. Select the **Select** check box for the interface that you want to modify.
4. In the **Admin. State** column for the interface that you want to modify, select **DOWN**.

   This step disables the interface and must be performed before you perform any OSPF modifications.
5. Click **Apply**.
6. In the navigation pane, expand the **Routing** > **IP** > **OSPF** folders.
7. Click **Interfaces**.

   The OSPF Interfaces Web page is displayed in the content pane (Figure 15-4).

Figure 15-4. OSPF Interfaces Web Page



8. Select the **Select** check box for the interface that you want to modify.
9. Modify the OSPF interface as necessary. See Table 15-3 for an explanation of each field on this Web page.
10. Click **Apply**.

11. Reenable the interface:

a. In the navigation pane, expand **Routing** > **IP** > **Configuration** folders.

b. Click **Interfaces**.

The IP Interfaces Web page is displayed in the content pane.

c. Select the **Select** check box for the interface that you want to reenable.

d. In the **Admin. State** column for the interface that you want modified, select **UP**.

e. Click **Apply**.

Table 15-3. OSPF Interface Web Page Parameters

Parameter	Definition
Select	Select the OSPF interface to be modified.
Interface	Displays IP interface (VLAN) that has OSPF enabled. **Note:** This is a read-only field.
IP Address	Displays IP address associated with the OSPF interface. **Note:** This is a read-only field.
Area	Enter the area ID configured for this interface. The default is 0.0.0.0.
DR Priority	Enter the decimal value for this interface for DR priority functionality. The value range is 0-255. The default is 1.
Transmit Delay	Enter the estimated time (seconds) it takes to transmit a link state update packet over this interface.The value range is 1-3600. The default is 1.
Retransmit Interval	Enter the time (seconds) between link-state advertisement retransmissions, for adjacencies belonging to this interface. This value is also used when retransmitting database description and link-state request packets. The value range is 1-3600. The default is 5.
Hello Interval	Enter the time (seconds) between the Hello packets that the router sends on the interface. This value must be the same for all routers attached to a common network. The value range is 1-65535. The default is 10.
	1 of 2

Table 15-3. OSPF Interface Web Page Parameters

Parameter	Definition
Dead Interval	Enter the time (seconds) that a router's Hello packets have not been seen before it's neighbors declare the router down. This should be some multiple of the Hello interval. This value must be the same for all routers attached to a common network. The value range is 1-65535. The default is 40.
Poll Interval	Enter the larger time interval (seconds) between the Hello packets sent to an inactive non-broadcast multiaccess neighbor. he value range is 1-65535. The default is 120.
Cost	Enter the cost metric associated with this interface. The value range is 1-65535. The default is 1.
Authentication	Select the type of authentication available for use on a given OSPF interface. Authentication types are: • None (default) • Simple Password • MD5 Authentication
Key	Enter the authorization key value for the interface.
MD5 Key ID	Enter the MD5 authentication key ID as a decimal value. The value range is 1- 255.
Interface State	Select the interface state. The available sites are: • Normal (default) • Passive
	2 of 2

CLI Command

To modify an OSPF interface using the CLI, enter the following command:

`(configure router:ospf)#` **network <*ip-address*> <*wildcard-mask*> <*area-id*>**

If you use the **network area** CLI command to add an interface to an OSPF area, be careful not to enter the IP mask for the <*wildcard-mask*> variable.

The <*wildcard-mask*> variable is the inverse of a network mask. Enter a 32-bit number in four-part, dotted decimal format. Place ones in the bit positions that you want to mask.

For example, to enable OSPF on interface 10.10.10.1 (where the network mask is 255.255.255.0) and assign it to area 2.2.2.2, enter: **network 10.10.10.1 0.0.0.255 area 2.2.2.2**

If you enter the IP network mask (255.255.255.0) instead of the wildcard mask (0.0.0.255), all OSPF interfaces that have 1 for the last octet of their IP address will be added to area 2.2.2.2.

Configuring an OSPF Passive-Interface

You can prevent OSPF from sending routing updates across the network by using the **OSPF passive-interface** command. Enabling this command prevents OSPF from sending hello packets across the network. Any OSPF hello packets received from other OSPF routers on the network are also ignored. OSPF advertises the passive-interface as a stub network in the router updates it sends to other OSPF interfaces.

You can set OSPF passive-interface using either the Web Agent or the CLI.

Web Agent Procedure

You can set OSPF passive-interface from the OSPF Interface Web page. The **Interface-State** field on the OSPF Interface Web page can be set to **Passive**, which disables the sending of OSPF routing updates. **Normal**, which enables the sending of OSPF routing updates, is the default.

CLI Commands

You can set OSPF Passive Interface from the CLI using the following commands:

- To define an OSPF interface as passive-interface:

  `(configure router:ospf)#` **passive-interface {<*interface-name*> | <*ip-addr*>}**

- To set the state of an OSPF interface that is configured as a passive-interface to normal:

  `(configure router:ospf)#` **no passive-interface {<*interface-name*> | <*ip-addr*>}**

Creating OSPF Virtual Links

You can create OSPF Virtual links using either the Web Agent or the CLI.

Web Agent Procedure

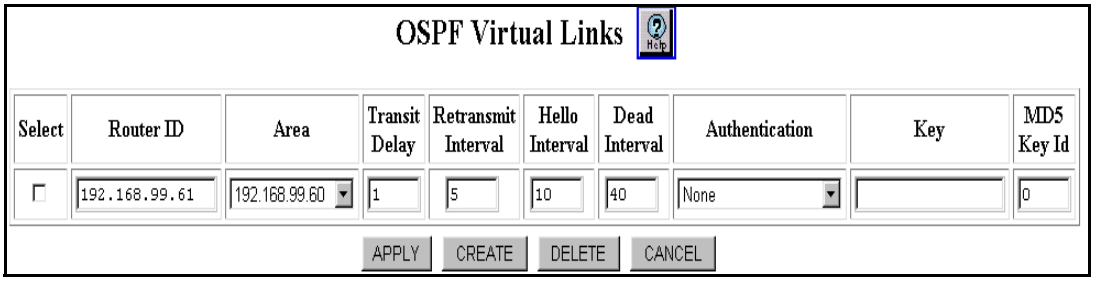
To create a OSPF virtual link using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **OSPF** folders, and then click **Virtual Links**.

   The OSPF Virtual Links Web page is displayed in the content pane. See Figure 15-5.

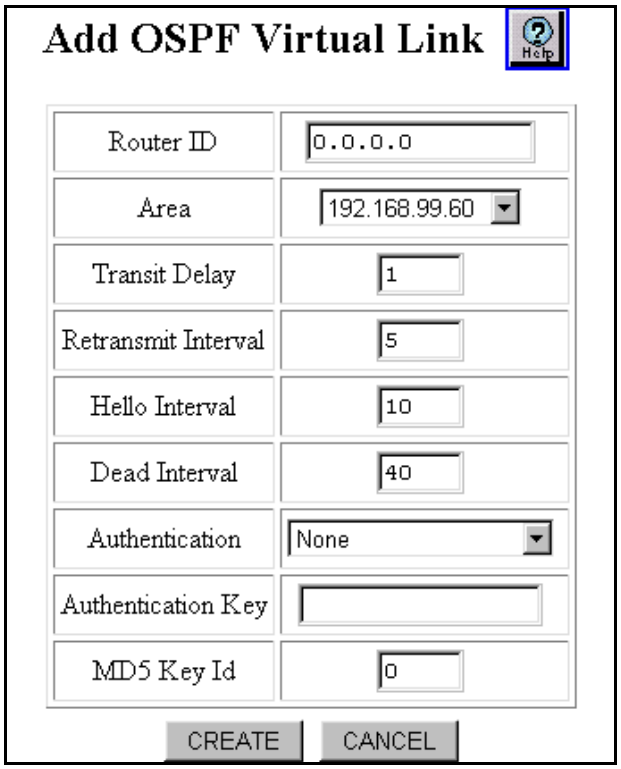
   ***Note:** The OSPF Web page displays only if OSPF virtual links were previously configured.

Figure 15-5. OSPF Virtual Links Web Page



2. Select **CREATE**. The Add OSPF Virtual Link Web page is displayed. See Figure 15-6.

Figure 15-6. Add OSPF Virtual Link Web Page



3. Enter the new OSPF Link information. See Table 15-4 for an explanation of the Add OSPF Virtual Link Web page parameters.

4. Click **CREATE** to create the virtual link, or **CANCEL** if you do not want to create the OSPF virtual link.

Table 15-4. Add OSPF Virtual Link Web Page Parameters

Parameter	Defines the...
Router ID	Router ID for the far end of the virtual link.
Area	Area ID through which the virtual link travels.
Transmit Delay	Estimated number of seconds it takes to transmit a link state update packet over this virtual link. The value range is 1-3600. The default is 1.
Retransmit Interval	Number of seconds between link-state advertisement retransmissions for adjacencies belonging to this virtual link. This value is also used when retransmitting database description and link-state request packets. The value range is 1-3600. The default is 5.
Hello Interval	Time in seconds between the Hello packets that the router sends on the virtual link. This value must be the same for all routers attached to a common network. The value range is 1-65535. The default is 10.
Dead Interval	Time in seconds that a router's Hello packets have not been seen before it's neighbors declare the router down. This should be some multiple of the Hello interval. This value must be the same for all routers attached to a common network.The value range is 1-65535. The default is 40.
Authentication	Type of authentication available for use on a given OSPF interface. Authentication types include: • None (default) • Simple Password • MD5 Authentication
Authentication Key	Authentication key for the Area.
MD5 Key ID	MD5 authentication key ID as a decimal value. The values range is 1-255.

CLI Command

To create a OSPF virtual link, use the following command:

`(configure router:ospf)#` **area** <***area-id***> **virtual-link** <***router-id***>

Deleting OSPF Virtual Links

You can delete OSPF Virtual links using either the Web Agent or the CLI.

Web Agent Procedure

To delete an OSPF virtual link using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > OSPF** folders, and then click **Virtual Links**.

   The OSPF Virtual Links Web page is displayed in the content pane. See Figure 15-5.

2. Select the virtual link that you want to remove.

3. Click **DELETE**. The virtual link is removed.

CLI Command

To delete an OSPF virtual link, use the following command:

`(configure router:ospf)#` **[no] area <*area-id*> virtual-link <*router-id*>**

Modifying OSPF Virtual Links

You can modify OSPF virtual links using either the Web Agent or the CLI.

Web Agent Procedure

To modify an OSPF virtual link using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > OSPF** folders, and then click **Virtual Links**.

   The OSPF Virtual Links Web page is displayed in the content pane. See Figure 15-5.

2. Select the virtual link that you want to modify.

3. See Table 15-4 to modify the OSPF Virtual Links Web page parameters.

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To modify an OSPF virtual link, use the following command:

`(configure router:ospf)#` **area <*area-id*> virtual-link <*router-id*> [hello-interval <*seconds*>] [retransmit-interval <*seconds*>] [dead-interval <*seconds*>] [{authentication-key <*key*> | message-digest-key <*key-id*> md5 <*key*>}]**

Creating OSPF Summaries

The primary purpose of the OSPF summary is route aggregation. Route aggregation is a group range of IP addresses that are linked to a single address.

You can create a new OSPF summary from either the Web Agent or the CLI.

Web Agent Procedure

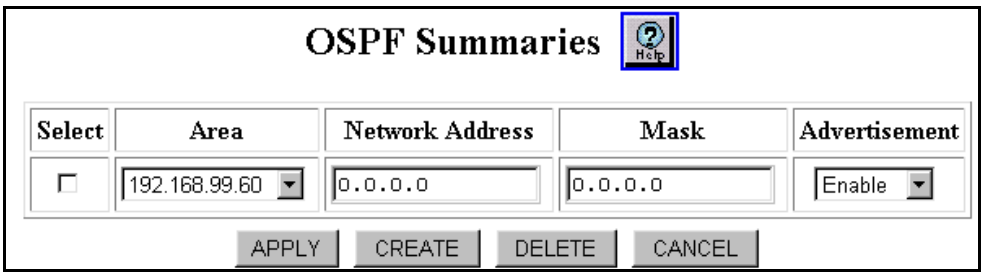
To create a new OSPF summary using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > OSPF** folders, and then click **Summaries**.

   The OSPF Summaries Web page is displayed in the content pane. See Figure 15-7.

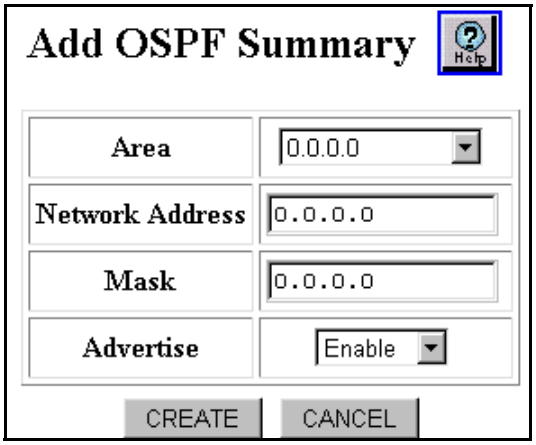
   ***Note:** This Web page will display OSPF summaries only if you have previously configured one or more summaries.

Figure 15-7. OSPF Summaries Web Page



2. Select **CREATE**. The Add OSPF Summary Web page is displayed. See Figure 15-8.

Figure 15-8. Add OSPF Summary Web Page



3. See Table 15-5 for details about the Add OSPF Summary Web page parameters:

Table 15-5. Add OSPF Summary Web Page Parameters

Parameter	Definition
Area	Select the area ID of the area from which the routes are aggregated (summary IP address).
Network Address	Enter the IP address of the network to be advertised.
Mask	Enter the subnet mask of the network to be advertised.
Advertise	Select the ability to suppress (disable) or enable advertisements of this summary. When suppressing, advertisements of IP routes in this range are also suppressed.

4. Click **CREATE** to save the new OSPF summary, or **CANCEL** to ignore the new summary.

CLI Command

To create a new OSPF summary, use the following command:

`(configure router:ospf)#` **area <*area-id*> range <*ip address*> <*mask*> [no-advertise]**

Deleting OSPF Summaries

You can delete OSPF summaries using either the Web Agent or the CLI.

Web Agent Procedure

To delete an OSPF summary using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **OSPF** folders, and then click **Summaries**.

   The OSPF Summaries Web page is displayed in the content pane. See Figure 15-7.

2. Select the OSPF summary that you want to remove.

3. Click **DELETE**. The OSPF summary you selected is removed.

CLI Command

To delete an OSPF summary, use the following command:

`(configure router:ospf)#` **[no] area <*area-id*> range <*ip-address*> <*mask*>**

Modifying OSPF Summaries

You can modify an OSPF summary using wither the Web Agent or the CLI.

Web Agent Procedure

To modify an OSPF summary using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **OSPF** folders, and then click **Summaries**.

   The OSPF Summaries Web page is displayed in the content pane. See Figure 15-7.

2. Select the OSPF summary to be modified.

3. See Table 15-5 to configure the OSPF Summaries Web page parameters.

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To modify an OSPF summary, use the following command:

`(configure router:ospf)#` **area** <***area-id***> **range** <***ip-address***> <***mask***>

Monitoring Switch Performance Using OSPF Statistics

You can monitor switch performance using the following OSPF statistics:

- OSPF Links
- OSPF Neighbors
- OSPF Link Database

This section contains the following procedures:

- Displaying OSPF Statistics
- Displaying OSPF Links
- Displaying OSPF Neighbors
- Searching the OSPF Link State Database

Displaying OSPF Statistics

You can display OSPF statistics using either the Web Agent or the CLI.

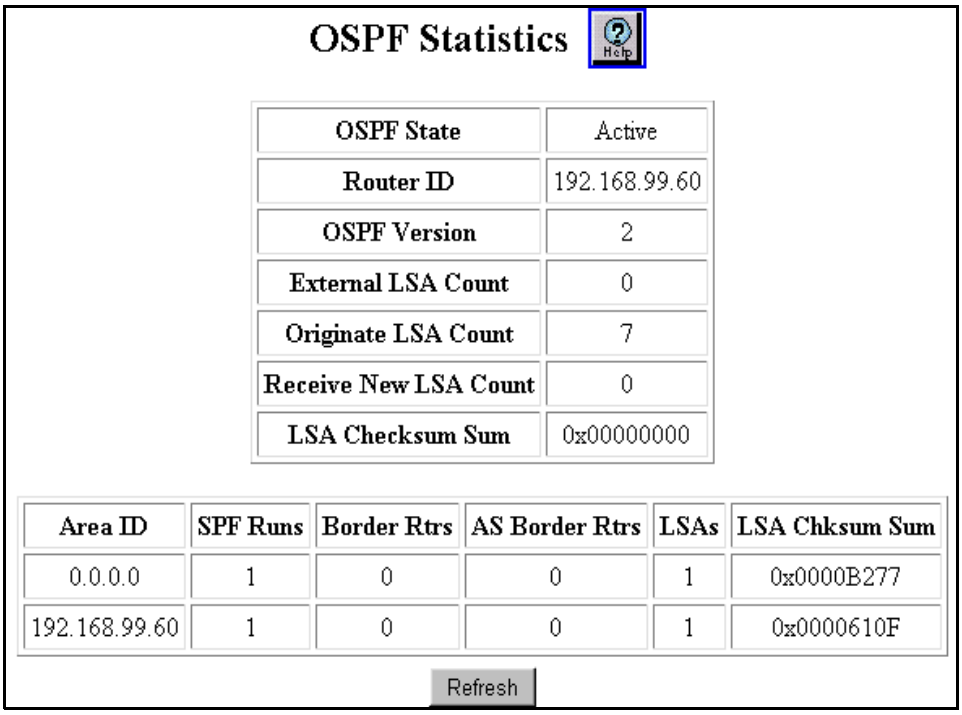
Web Agent Procedure

To display the OSPF global statistics using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > OSPF** folders, and then click **Statistics**.

   The OSPF Statistics Web page is displayed in the content pane. See Figure 15-9.

Figure 15-9. OSPF Statistics Web Page



2. See Table 15-6 for an explanation of the OSPF Statistics Web page parameters.

3. Click **Refresh** to reset the counters with the latest information.

Table 15-6. OSPF Statistical Parameters

Parameter	Displays the...
OSPF State	Current state of OSPF.
Router ID	Router ID for OSPF.
	1 of 2

Table 15-6. OSPF Statistical Parameters

Parameter	Displays the...
OSPF Version	Current version of OSPF. The P580/P882 with Integrated Routing supports OSPFv2.
External LSA Count	Number of external (LS type 5) link state advertisements (LSAs) in the link-state database.
Originate LSA Count	Number of LSAs originated by this router.
Receive New LSA Count	Number of new LSAs received by this router.
LSA Checksum Sum (global OSPF system)	32-bit unsigned sum of the LS checksums of the external link-state advertisements contained in the link-state database. This sum can be used to determine if there has been a change in a router's link state database, and to compare the link-state database of two routers.
Area ID	Area ID of the area in question. It has the function of defining a summarization point for Link State Advertisements.
SPF Runs	Number of times that the intra-area route table has been calculated using this area's link-state database.
Border Rtrs	Total number of area border routers reachable within this area. This is initially zero, and is calculated in each SPF Pass.
AS Border Rtrs	Total number of Autonomous System border routers reachable within this area. This is initially zero, and is calculated in each SPF Pass.
LSAs	Total number of link-state advertisements in this area's link-state database, excluding AS External LSAs.
LSA Chksum Sum (per area, not globally)	32-bit unsigned sum of the LS checksums of the external link-state advertisements contained in the link-state database. This sum can be used to determine if there has been a change in a router's link state database, and to compare the link-state database of two routers.
	2 of 2

CLI Commands

To display the OSPF global statistics using the CLI, use the following commands:

- **> show ip ospf stats**
- **> show ip ospf virtual-links**

Displaying OSPF Links

The OSPF link display provides information on the currently active OSPF links. You can display the OSPF using either the Web Agent or the CLI.

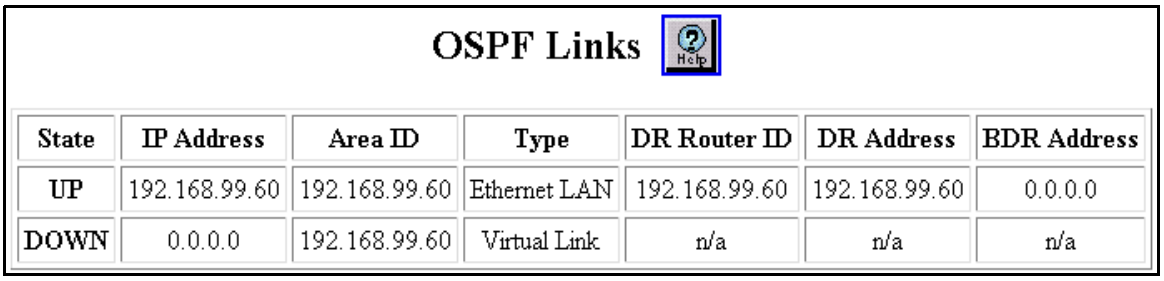
Web Agent Procedure

To display OSPF links using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IP** > **OSPF** folders, and then click **Links**.

   The OSPF Links Web page is displayed in the content pane. See Figure 15-10.

Figure 15-10. OSPF Links Web Page



2. See Table 15-7 for an explanation of the OSPF Links Web page parameters:

Table 15-7. OSPF Link Parameters

Parameter	Displays the...
State	State of the OSPF link.
IP Address	IP address of the OSPF link.
Area ID	Area ID (IP address) associated with the OSPF link.
Type	Type of OSPF link.
DR Router ID	Router ID of the designated router.
DR Address	IP address of the designated router.
BDR Address	IP address of the designated border router.

CLI Commands

To display OSPF links using the CLI, enter the following command:

- **> show ip ospf interface**
- **> show ip ospf virtual-links**

Displaying OSPF Neighbors

The OSPF neighbors table summarizes the list of OSPF interfaces and their associated neighbors. You can display OSPF neighbors using either the Web Agent or the CLI.

Web Agent Procedure

To display the OSPF neighbors using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > OSPF** folders, and then click **Neighbors**.

   The OSPF Neighbors Web page is displayed in the content pane. See Figure 15-11.

Figure 15-11. OSPF Neighbors Web Page

OSPF Neighbors

IP Address	State	Router ID	Master	DD Number	DR Priority	E-Option	T-Option	MC-Option	N-Option	OPQ-Option	DR Choice	BDR Choice
15.0.7.1	FULL	15.0.7.1	Slave	0x00001107	1	Yes	No	No	No	No	15.0.7.1	15.0.7.2
15.0.21.2	FULL	15.0.21.2	Master	0x0000161A	1	Yes	No	No	No	No	15.0.21.2	15.0.21.1

2. See Table 15-8 to review the OSPF Neighbors Web page parameters:

Table 15-8. OSPF Neighbors Web Page Parameters

Parameter	**Definition**
IP Address	The IP address associated with the OSPF neighbor.
State	The functional level of an interface. States include: • **Down** - This is the initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. • **Attempt** - Indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor • **INIT** - Indicates that the Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor. • **2-Way** - Communication between the two routers is bidirectional. • **ExStart** - This is the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial Database Description sequence number. Neighbor conversations in this state or greater are called adjacencies. • **Exchange** - Router is describing its entire link state database by sending DD (Database Description) packets to the neighbor. • **Loading** - Link State Request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the Exchange state. • **Full** - The neighboring routers are fully adjacent. These adjacencies appears in router-LSAs and network-LSAs.
Router ID	The router ID of the neighbor.
Master	The state of the neighbor: master or slave.
DD Number	The hexadecimal number used to sequence the collection of Database Description Packets. The initial value (indicated by the Init bit being set) should be unique. The DD sequence number then increments until the complete database description has been sent.
DR Priority	Displays the DR Priority of the neighboring router. This is used to determine whether this neighbor is eligible to become the *Backup* Designated Router. If 0, the router is ineligible to become the Backup Designated Router.
E-Option	The method used to flood AS-external-LSAs.
T-Option	Specifies whether this neighbor is an ASBR.
MC-Option	Specifies whether this neighbor supports MOSPF.
N-Option	Specifies whether this neighbor supports the handling of Type-7 LSAs.
OPQ-Option	Specifies whether this neighbor supports opaque LSAs (LSA type 9-11).
DR Choice	Interface IP address of the designated router.
BDR Choice	Interface IP address of the backup designated router.

CLI Command

To display the OSPF neighbors using the CLI, enter the following command:

> show ip ospf neighbor

Searching the OSPF Link State Database

You can search the OSPF kink state database using either the Web Agent or the CLI.

Web Agent Procedure

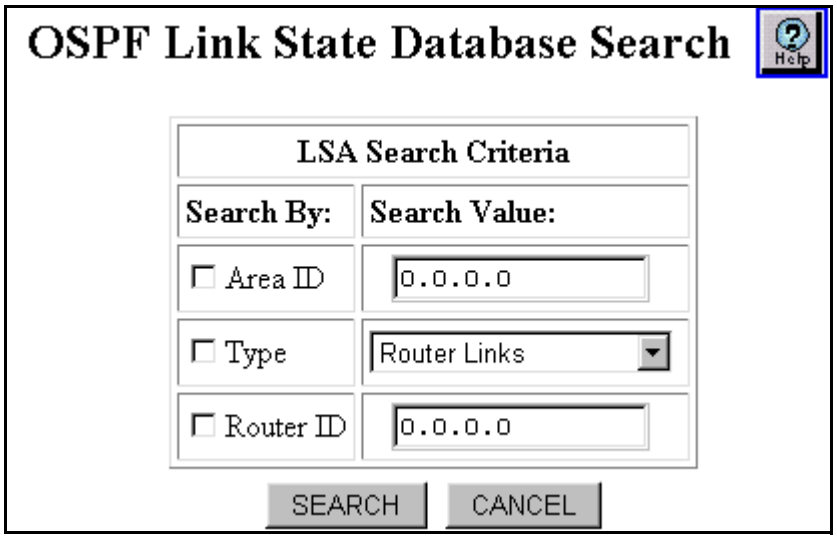
To search the OSPF link state database using the Web Agent:

1. In the navigation pane, expand the **Routing > IP > OSPF** folders, and then click **Link State Database Search**.

   The Link State Database Search Web page is displayed in the content pane. See Figure 15-12.

   ***Note:** You can select more than one item in the **Search By** column to help narrow your search results.

Figure 15-12. Link State Database Search Web Page



2. To search by:

   — **Area ID** — click the Area ID checkbox and specify the IP address of the Area ID that you want to find in the database, and click SEARCH.

   — **Type** — click the Type checkbox and from the Type pull-down menu, select the type of search you want to perform, and click SEARCH.

   — **Router ID** — click the Router ID checkbox and specify the router IP address, and click SEARCH.

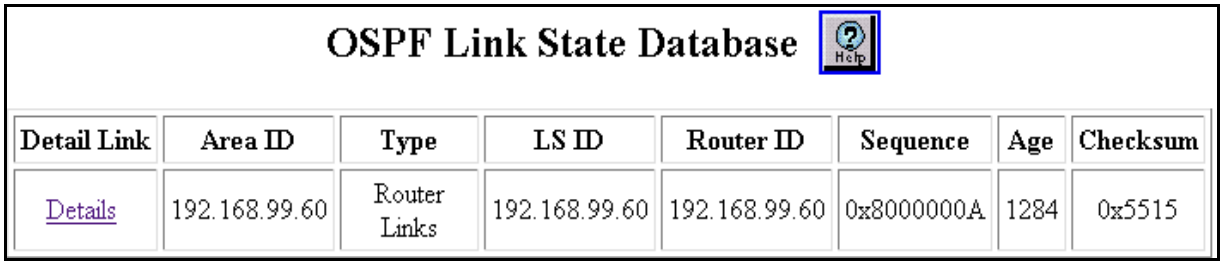
3. See Table 15-9 for an explanation of the OSPF Link State Database Search Web page parameters:

Table 15-9. OSPF Link State Database Search Web Page Parameters

Parameter	Allows you to...
Area ID	Searches the database for the 32-bit identifier of the area from which a LSA was received.
Type	Searches the database for all entries that match one of the following types: • **Router Links** - These packets describe the states of the router's links to the area and are only flooded within a particular area. • **Network Links** - These packets are generated by Designated Routers and describe the set of routers attached to a particular network. • **Summary Network** - These summaries are generated by Area Border Routers and describe inter-area routes to various networks. They can also be used for aggregating routes. • **Summary AS border** - This describes links to Autonomous System Border Routers and are generated by Area Border Routers. • **AS external** - These packets are generated by Autonomous System Border Routers and describe routes to destination external to the Autonomous system. They are flooded everywhere except stub areas. • **Multicast group** - These packets are generated by multicast groups. • **NSSA external** - These packets are generated by Area Border Routers and describe routes within the NSSA (Not-So-Stubby-Area).
Router ID	Searches the database for all entries that this router originated.

4. If your search produces results, the detailed information displays in the OSPF Link State Database Web page (Figure 15-13).

Figure 15-13. OSPF Link State Database Web Page



5. See Table 15-10 for an explanation of the OSPF Link State Database Web page parameters.

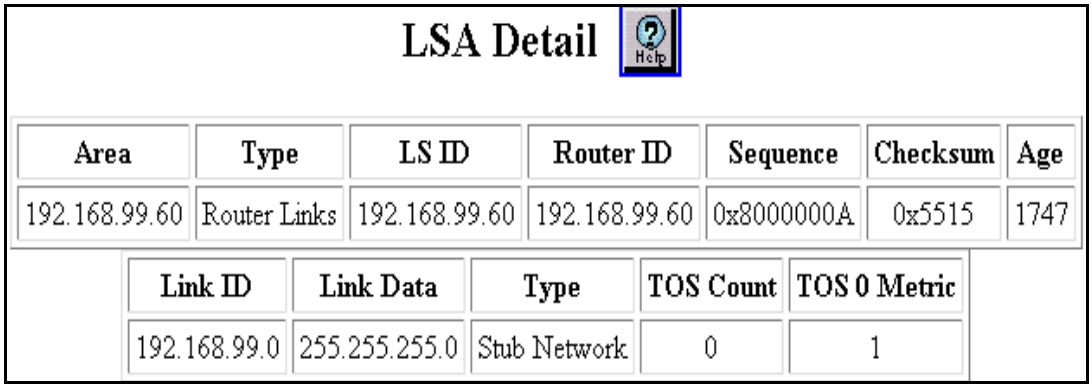
Table 15-10. OSPF Link State Database Parameters

Parameter	Displays...
Detail Link	A link to the LSA Detail Web page.
Area ID	The 32-bit identifier of the area from which the LSA was received.
Type	The LSA format and function. Types include: • **Router Links** - These packets describe the states of the router's links to the area and are only flooded within a particular area. • **Network Links** - These packets are generated by Designated Routers and describe the set of routers attached to a particular network. • **Summary Network** - These summaries are generated by Area Border Routers and describe inter-area routes to various networks. They can also be used for aggregating routes. • **Summary AS Border** - This describes links to Autonomous System Border Routers and are generated by Area Border Routers. • **AS External** - These packets are generated by Autonomous System Border Routers and describe routes to destinations external to the Autonomous system. They are flooded everywhere except stub areas. • **Multicast Group** - These packets are generated by multicast groups. • **NSSA External** - These packets are generated by Area Border Routers and describe routes within the NSSA (Not-So-Stubby-Area).
LS ID	The piece of routing domain that is being described by the advertisement. Depending on the advertisements LS type, the LS ID displays different values.
Router ID	The 32-bit number that uniquely identifies the originating router in the Autonomous System.
Sequence	The sequence number.
Age	Displays the age (in seconds) of the LSA.
Checksum	Displays the checksum of the complete contents of the advertisement, except the age field.

6. To update all LSA entries in the database, select **Refresh Table**.

7. To display the details of a specific link state advertisement (LSA), select **Details** from the **Detail Link** column The LSA Detail Web page displays with additional search information (Figure 15-14).

Figure 15-14. LSA Detail Web Page



8. See Table 15-11 for an explanation of the LSA Detail Web page parameters.

Table 15-11. LSA Detail Web Page Parameters

Parameter	Definition
Area	Displays the 32-bit identifier of the area from which the LSA was received.
Type	Displays the link state type. Types include: • Router Links • Network Links • Summary Network • Summary AS Border • AS External • Multicast Group • **NSSA External**
LS ID	Displays the link-state ID. The link-state ID is an LS type specific field containing either a router ID or an IP address that identifies the piece of the routing domain that is being described by the advertisement.
Router ID	Displays the router ID of the originator of the link state advertisement.
Sequence	Displays the link-state sequence number. The sequence number is a 32-bit signed integer.
Checksum	Displays the checksum of the complete contents of the advertisement, except the age field.
Age	Displays the time (in seconds) of the LSA.
	1 of 2

Table 15-11. LSA Detail Web Page Parameters

Parameter	Definition
Network Mask	Displays the network mask for the LSA.
Attached Router ID 1 &2	Displays the router ID for the attached router(s).
	2 of 2

9. To update the LSA entry, select **Refresh Entry**.

CLI Command

To search the OSPF link state database using the CLI, enter the following command:

> show ip ospf database <*options*>

NBMA IP Interfaces

Overview

Non-broadcast multi-access (NBMA) functionality was added to RIP and OSPF routing protocols on the Avaya Multiservice switch software. Support for this functionality was added to enable route exchange over nonbridged connections (routed PVCs).

This feature makes it possible for the switch to exchange routing information over nonbridged connections (routed virtual switch ports). NBMA functionality has been added to the RIP and OSPF routing protocols in the switch software and has been tested with other routers.

Operation of OSPF over NBMA is almost identical to operation of OSPF over broadcast LANs. Flooding uses the designated router, and both subnets are represented identically within the OSPF link-state database by network-LSAs. The only real difference between broadcast subnets and NBMA subnets is in the discovery of neighboring routers. On broadcast networks, a router can discover its neighbors dynamically by sending multicast Hello packets; on NBMA networks, a router's neighbors may have to be configured.This feature makes it possible for you to set up an IP interface as NBMA. Each IP routing protocol may then have IP addresses that are designated as NBMA neighbors.

Each IP routing protocol may then have IP addresses that are designated as NBMA neighbors. This allows the switch to work with other routers that are routed over PVCs.

CAUTION: The NBMA feature is not backward compatible. If you have NBMA interfaces set up and need to downgrade from v6.1 to a 5.2 version of software, contact technical support.

This section contains procedures for the following tasks:

- Setting Up an NBMA IP Interface
- Setting Up NBMA Neighbors
- Removing NBMA Neighbors

Setting Up an NBMA IP Interface

You can set up an NBMA IP interface by using either the Web Agent or the Command Line Interface (CLI).

Web Agent Procedure

To use the Web Agent to set up an NBMA IP interface:

*** Note:** RIP or OSPF must be enabled to set up NBMA neighbors.

*** Note:** See "Creating and Assigning IP Interfaces to the VLAN" in Chapter 12, "Configuring IP Routing," for option settings.

1. In the navigation pane, expand the **Routing** > **IP** > **Configuration** folders, and then click **Interfaces**.

   The IP Interfaces Web page is displayed.

2. Click **Create**.

   The Add IP Interface Web page is displayed.

3. Enter a name for the interface in the **Name** field.

4. From the **Interface Type** list, select **NBMA**.

   ***Note:** If you enter an IP address to create an NBMA neighbor entry to an interface that is not an NBMA type, the switch will return you to the NBMA Neighbor display page without displaying an error message.

5. Enter the appropriate information in the remaining fields.

6. In either the **RIP** or **OSPF** list or both lists, select **Enable**.

7. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

   ***Note:** You must bring the IP interface Down to successfully enable OSPF on the IP interface. If you do not, you will get an error message and OSPF will not be enabled.

CLI Commands

To use the CLI to set up an NBMA IP interface:

1. Enter Configuration mode. The CLI displays the following prompt:

   ```
   (configure)#
   ```

2. Enter the following command at the prompt:

   `(configure)#` **interface <*interface-name*>**

   The following prompt displays:

   ```
   (configure-if:<interface-name>)#
   ```

3. Enter the following command at the prompt:

   ip address <*ip-addr*> <*mask*>

4. Enter the following command at the prompt:

   type nbma

Setting Up NBMA Neighbors

Web Agent Procedure

To use the Web Agent to set up NBMA neighbors:

1. Open the **Routing** > **IP** > **RIP** or **OSPF** folders, whichever one was created with an NBMA type interface.
2. Click **NBMA Neighbors**. The RIP or OSPF NBMA Neighbors Web page is displayed.
3. Click **Create**. The Add NBMA RIP or OSPF Neighbor Web page is displayed.
4. Enter a neighbor IP interface address in the **Network Address** field.
5. Click **CREATE**, or **CANCEL** to restore previous setting.

CLI Commands

To set up NBMA neighbors on OSPF, you must first create an OSPF interface.

After you create an OSPF interface, use the following command to set up NBMA neighbors:

`(configure router:ospf)#` **nbma-neighbor** <***ip address***>

Removing NBMA Neighbors

You can remove NBMA neighbors using either the Web Agent or the CLI.

CLI Command

To remove the NBMA neighbors using the CLI, enter the following command at the `(configure router:ospf)#` prompt:

`(configure router:ospf)#` **no nbma-neighbor** <***ip address***>

*** Note:** Do not change the type of interface after you have created it. If you need to change the interface type, delete the old one and create a new one.

*** Note:** Poll-interval of NBMA interface decreases in value to 0, during the operation. It does not affect the operation of OSPF on an NBMA interface. A panic message will display. Ignore this message.

*** Note:** NBMA RIP2 Poison Reverse in MD5 mode does not advertise all the routes in the Route Table.

RIP2 MD5 Authentication sequence numbers do not persist after you reset the switch.

CLI Commands for RIP

To set up an NBMA neighbor on RIP, you must first create a RIP interface.

1. After you create a RIP interface, use the following command to set up NBMA neighbors:

   `(configure router:rip)#` **nbma-neighbor** <***ip address***>

16 Configuring IPX Routing

Overview

This chapter provides the following information and procedures for configuring layer 2 modules or licensed layer 3 modules. See Chapter 1 "Introduction" for more details about routing IP and Internetwork Packet Exchange (IPX) Protocol through the Avaya Multiservice Switch.

- IPX Overview
- Configuring the Avaya Switch as an IPX Router
- Monitoring Your Switch Using IPX

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

IPX Overview

The IPX protocol is connectionless and performs datagram delivery and routing in Novell NetWare networks. Each IPX address consists of:

- **Network Number** — A 32-bit (8 characters) number that is normally assigned by the network administrator.
- **Node Number** — A 48-bit (12 characters) number that is normally the MAC layer address of the physical interface.
- **Socket Number** — A number used to route packets to different processes within the same node.

The syntax for entering an IPX address is: *network node socket*

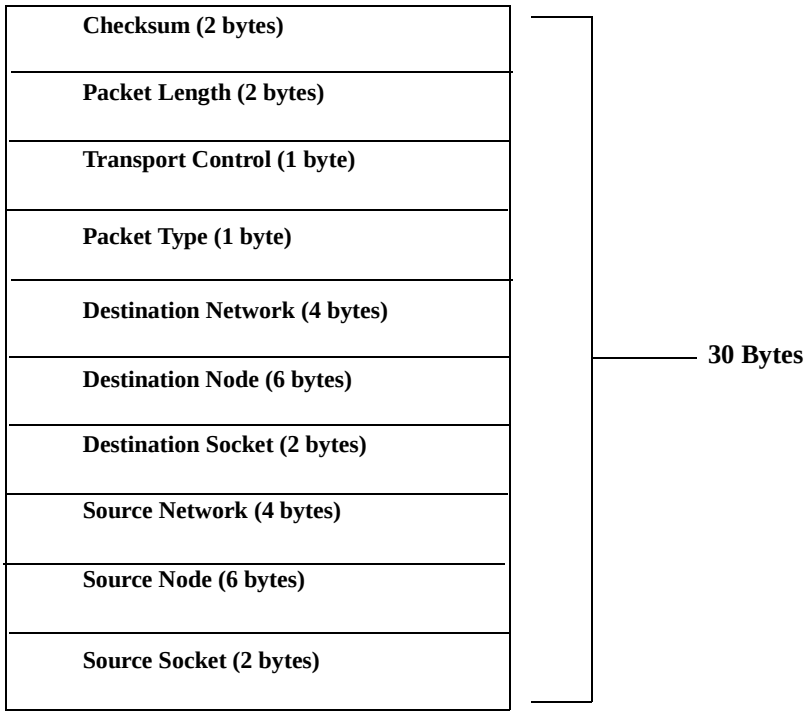
For example: 000000AAh 00e03b124213h 4003h

where `000000AAh` is the network number, `00e03b124213h` is the node number, and `4003h` is the socket number associated with a running process on the end node (for example, RIP, NetWare Link State Protocol (NLSP)).

IPX Datagram Structure

The IPX datagram contains an IPX header and any data to be transferred on the network. The IPX header is a 30 byte header that contains 10 fields. Figure 16-1 illustrates a conceptual view of an IPX datagram:

Figure 16-1. IPX Datagram in Detail



See Table 16-1 for a description of the IPX datagram Structure fields:

Table 16-1. IPX Datagram Structure Fields

Field	Definition
Checksum	Provides integrity checking. **Note:** Checksum is normally *not enabled* in IPX networks and is usually set to 0xFFFF.
Packet Length	Length (in bytes) of the packet.
Transport Control	Number of routers a packet has traversed. This is used to discard a packet if the packet traverses a maximum number of routers (16).

1 of 2

Table 16-1. IPX Datagram Structure Fields

Field	Definition
Packet Type	Indicates the type of service required or offered by the packet. Types include: • Sequenced Packet Exchange (SPX packet) • NetWare Core Protocol (NCP packet) • NetBIOS (propagated packet)
Destination Network	The IPX network address of the destination network.
Destination Node	The MAC address of the destination node.
Destination Socket	Address of the process running in the destination node. Sockets route packets to different processes within the same node.
Source Network	The network address of the source network.
Source Node	The MAC address of the source node.
Source Socket	Address of the process running in the source node.
	2 of 2

Configuring the Avaya Switch as an IPX Router

You can configure the Avaya P580 or P882 Multiservice switch as an IPX router. Specific hardware is required to accomplish this configuration.

This section contains the following information and procedures:

- Hardware Requirements
- Configuring IPX Routing Globally
- Configuring IPX Interfaces
- Creating IPX Static Routes
- Deleting IPX Static Routes
- Modifying IPX Static Routes
- Creating IPX Static Services
- Deleting IPX Static Service
- Modifying IPX Static Service

Hardware Requirements

To configure your switch as an IPX router, your Avaya P580 or P882 Multiservice switch must be configured with the following hardware:

- Layer 3 50-series media modules.
- 80-Series Supervisor Modules
- 80-Series Licensed Layer 3 media modules

Configuring IPX Routing Globally

You can configure IPX routing globally using the Web Agent or the CLI.

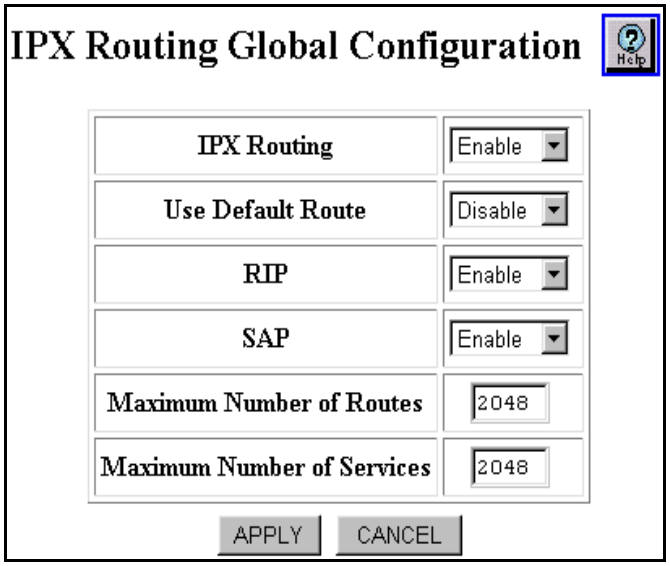
Web Agent Procedure

To Enable or Disable IPX routing globally using the Web Agent:

1. In the navigation pane, expand the **Routing > IPX > Configuration** folders and then click **Global Configuration**.

   The IPX Global Configuration Web page is displayed in the content pane. See Figure 16-2.

Figure 16-2. IPX Global Configuration Web Page



2. Configure IPX Routing Global Configuration Web page parameters to make your switch an IPX router. See Table 16-2 for an explanation of the Web page parameters.

3. Click APPLY to save your changes, or CANCEL to restore previous settings.

Table 16-2. IPX Routing Global Configuration Web Page Parameters

Parameter	Allows you to...
IPX Routing	Select Disable to prevent IPX from routing globally. The default value is Enable.
Use Default Route	Select Enable if you know the default route. The default value is Disable.
RIP	Select Disable to prevent IPX RIP from routing globally. This affects all IPX interfaces set up to use the IPX RIP routing protocol. The default value is Enable.
SAP	Select Disable to prevent IPX SAP from routing globally. This affects all IPX interfaces set up to use the IPX SAP routing protocol. The default value is Enable.
Maximum Number of Routes	Enter the maximum number of routes that can be added to the routing table. The switch rounds your entry to the nearest power of 2. For example, if you enter 1000, the system rounds this number up to 1024 routes. The default value is 2048.
Maximum Number of Services	Enter the maximum number of services that can be added. The switch rounds your entry to the nearest power of 2. For example, if you enter 1000, the system rounds this number up to 1024 services. The default value is 2048.

CLI Commands

To enable IPX routing globally using the CLI, enter the following command from Configure mode:

- To enable IPX routing globally, `(configure)#` **ipx routing**
- To disable IPX routing globally, `(configure)#` **no ipx routing**

Configuring IPX Interfaces

You can create new IPX interfaces and configure up to four interfaces of different encapsulation types to a VLAN (multinetting). You can configure an IPX interface using either the Web Agent or the CLI.

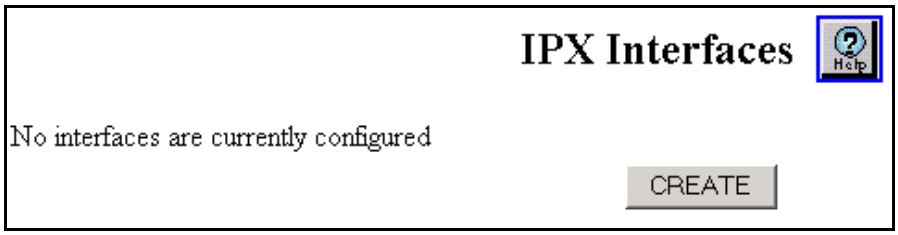
Web Agent Procedure

To create a new IPX interface using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **Configuration** folders, and then click **Interfaces**.

   The IPX Interfaces Web page is displayed in the content pane. See Figure 16-5.

Figure 16-3. IPX Interfaces Web Page

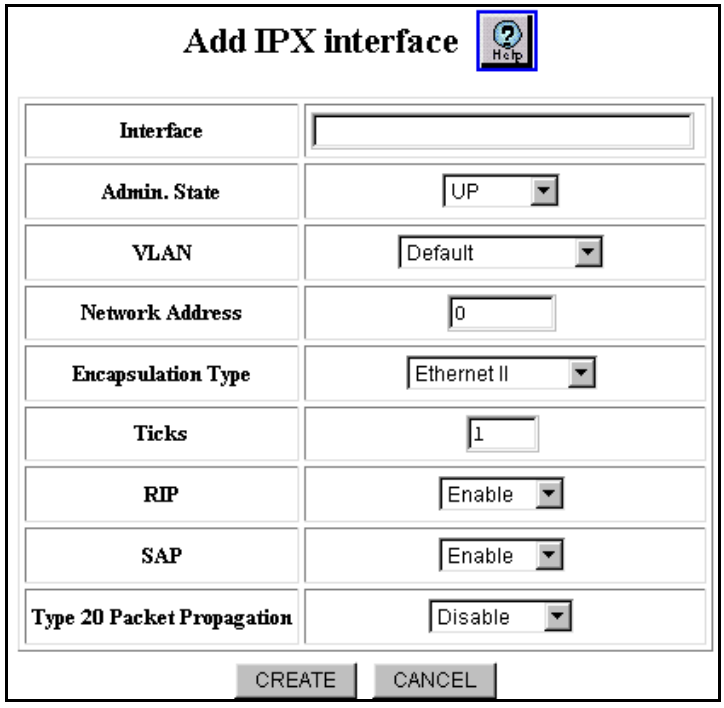


2. To create a new IPX interface Click on **CREATE**.

*** Note:** Only the create button is shown if no IPX Interfaces are configured.

The Add IPX interface Web page is displayed. See Figure 16-4

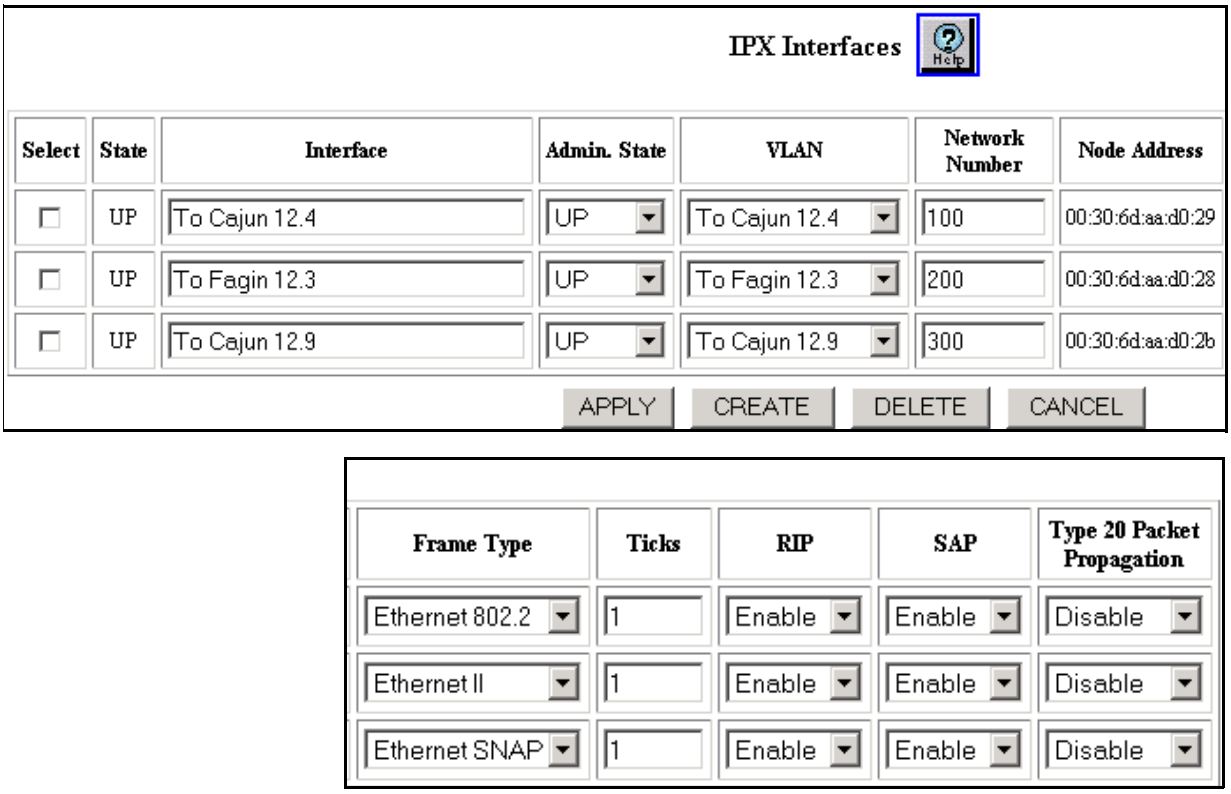
Figure 16-4. ADD IPX Interface Web Page



3. See Table 16-3 to configure an IPX Interface.

4. Select **CREATE**. The IPX Interfaces Web page is displayed. See Figure 16-5.

Figure 16-5. IPX Interfaces Web Page



***Note:** Due to its length, the IPX Interfaces Web page is split into two views.

5. To view or modify IPX Interfaces, expand the **Routing** > **IPX** > **Configuration** folders in the navigation pane, and then click **Interfaces**.

   The IPX Interfaces Web page is displayed in the content pane. See Figure 16-5.

6. Click **APPLY** to save your changes or **CANCEL** to restore previous settings.

Table 16-3. IPX Interface Web Page Parameters

Parameter	Allows you to...
Select	Select the Interface to be configured. This field id displayed in the IPX Interfaces Web page not in the Add Interfaces Web page.
Admin. State	Select the administration state of the interface. Options include: **Up** **Down**
Interface	Enter the name of the IPX interface.Up to 31 Alpha-numeric characters
VLAN	Select the VLAN that corresponds to the IPX interface you selected. **Note:** If you need to create a new VLAN, see "Creating and Implementing VLANs" in Chapter 6, "Using VLANs, Hunt Groups, and VTP Snooping.":
Network Number	Enter the number of the IPX network you want to assign to the IPX interface. This number is a hexadecimal 32 bit (8 characters) number.
Node Address	Displays the node address on which the IPX interface resides. A 48-bit (12 characters) number. This parameter is displayed in the IPX Interface Web page, not in the Add IPX Interface Web page.
Encapsulation Frame Type	Select the format of the MAC header on the IPX packets sent by the router on the interface. Formats include: **Ethernet II** (Maximum Transmission Unit (MTU) = 1500) **Ethernet 802.2** (MTU = 1497) **Ethernet SNAP** (MTU = 1492) **Ethernet 802.3** (MTU = 1500) The default is **Ethernet II**
Ticks	Specify the amount of time (in ticks) that the packet takes to reach the network number you specified. A tick is approximately 1/18th of a second. The default value is **1**
RIP	Select Disable to prevent IPX RIP from routing globally. This affects all IPX interfaces set up to use the IPX RIP routing protocol. The default value is **Enable.**
	1 of 2

Table 16-3. IPX Interface Web Page Parameters

Parameter	Allows you to...
SAP	Select **Disable** to prevent IPX SAP from routing globally. This affects all IPX interfaces set up to use the IPX SAP routing protocol. The default value is **Enable**.
Type 20 Packet Propagation	Specify whether or not Type 20 packets can be sent, received, disabled, or sent and received on a given interface. Options include: **Disable** - Type 20 packets are neither sent nor received. **Inbound** - Type 20 packets are received. **Outbound** - Type 20 packets are sent. **Both** - Type 20 packets are sent and received. The default value is **Disable**
	2 of 2

CLI Command

To begin the creation of a new IPX interface using the CLI, enter the following command in Configure mode:

`(configure)#` **interface** <***interface name***>

Creating IPX Static Routes

You can create IPX static routes using either the Web Agent or the CLI.

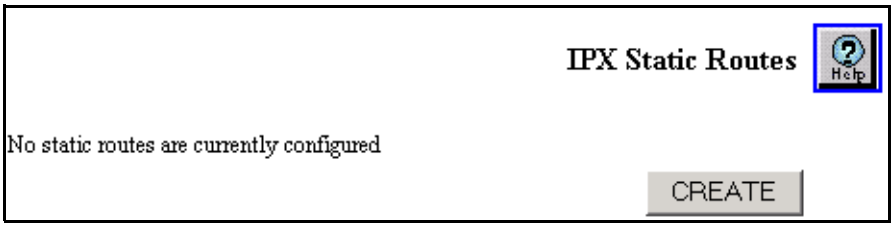
Web Agent Procedure

To create IPX static routes using the Web Agent:

1. In the navigation pane, expand the **Routing > IPX > Configuration** folders, and then click **Static Routes**.

   The IPX Static Routes Web page is displayed in the content pane. See Figure 16-6.

Figure 16-6. IPX Static Routes Web Page

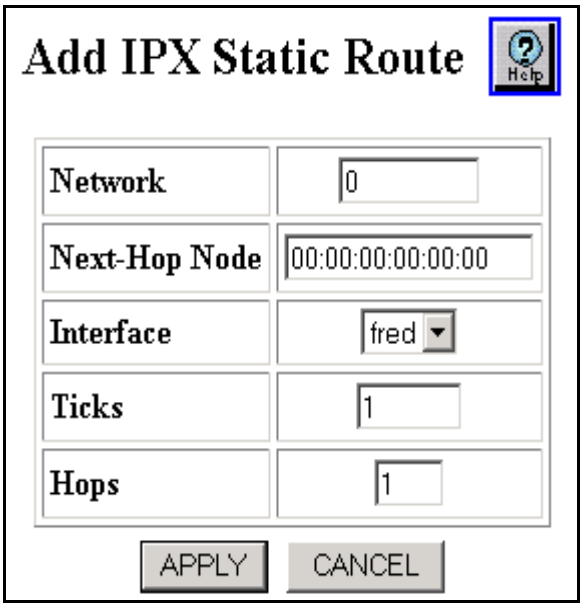


2. To create a new IPX Static Route, click **CREATE**.

*** Note:** Only the create button is shown if no IPX Static Routes are configured.

The Add IPX Static Route Web page is displayed. See Figure 16-7.

Figure 16-7. Add IPX Static Route Web Page



3. See Table 16-4 to configure the IPX Static Route parameters.

Table 16-4. IPX Static Route Parameters

Parameter	Allows you to...
Network	Enter the IPX network number that you want to assign to the IPX static route.
Next-Hop Node	Enter the MAC address for the next destination to which the packet is routed. Format of the value to enter is **aa:bb:cc:dd:ee:ff**.
Interface	View the IPX Interface associated with the next hop.
Ticks	Enter the amount of time (in ticks) that the packet takes to reach the network number you specified. A tick is approximately 1/18th of a second.
Hops	Enter the number of routers (hops) that the packet must pass through before reaching the network number associated with the IPX network.

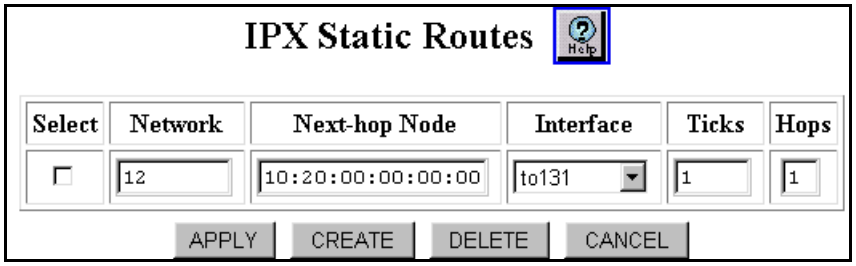
4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

5. To view or modify IPX Static Routes, expand the **Routing** > **IPX** > **Configuration** folders in the navigation pane, and then click **Static Routes**.

   The IPX Static Routes Web page is displayed in the content pane. See Figure 16-8.

   ***Note:** It is also possible to create a static route by, making a RIP source route from the IPX Static Route Table.

Figure 16-8. IPX Static Routes Web Page



CLI Command

To create IPX static routes using the CLI, enter the following command in Interface mode:

`(config-if:<if-name>)#` **ipx route** <***network address*** | ***range***> <***options***>

Deleting IPX Static Routes

You can delete an IPX static route using either the Web Agent or the CLI.

Web Agent Procedure

To delete an IPX static route using the Web Agent:

1. In the navigation pane, expand the **Routing > IPX > Configuration** folders, and then click **Static Routes**.

   The IPX Static Routes Web page is displayed in the content pane. See Figure 16-8.

2. Select the checkbox for the IPX static route that you want to delete from the Select column.

3. Click **DELETE** to remove the IPX static route, or CANCEL to restore the IPX static route.

CLI Command

To delete an IPX static route using the CLI, enter the following command from Configure mode:

`(configure)#` **clear ipx route {<*network*> default | *}**

Modifying IPX Static Routes

You can modify IPX static routes form either the Web Agent or the CLI.

Web Agent Procedure

To modify an IPX static route using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **Configuration** folders, and then click **Static Routes**.

   The IPX Static Routes Web page is displayed in the content pane. See Figure 16-7.

2. Select the checkbox for the IPX static route that you want to modify from the Select column.

3. See Table 16-4 to modify the IPX Static Route configuration.

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To modify an IPX static route using the CLI, enter the following command Interface mode:

`(config-if:<`*`if-name`*`>)#` **ipx route** <***network address*** | ***range***> <***options***>

Creating IPX Static Services

You can create IPX static services using either the Web Agent or the CLI.

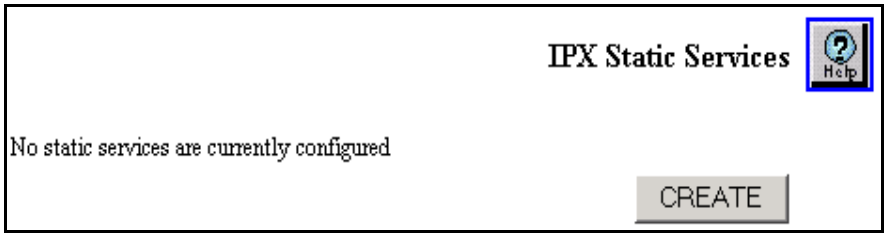
Web Agent Procedure

To create IPX static services using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **Configuration** folders, and then click **Static Services**.

   The IPX Static Services Web page is displayed in the content pane. See Figure 16-9.

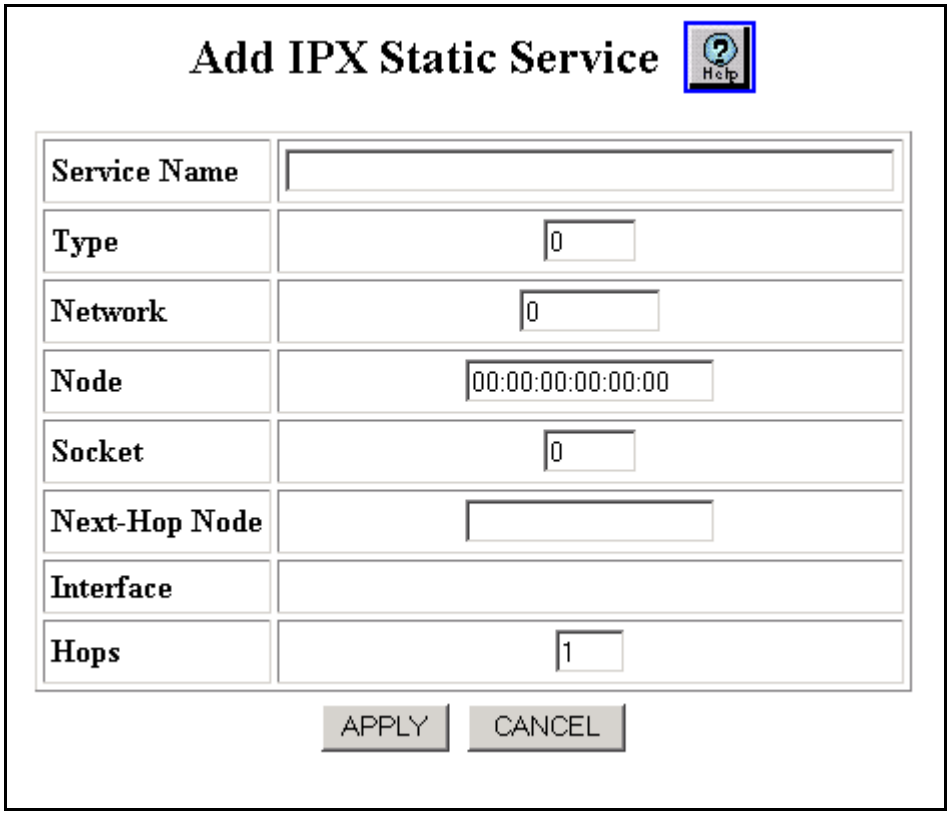
Figure 16-9. IPX static Services Web Page



2. Select **CREATE** to create a new static service. The Add IPX Static Service Web page is displayed. See Figure 16-10.

*** Note:** Only the Create button is shown if no IPX Static Services are configured.

Figure 16-10. Add IPX Static Service Web Page



3. See Table 16-5 to configure the Add IPX Static Service Web page parameters.

Table 16-5. IPX Static Service Web Page Parameters

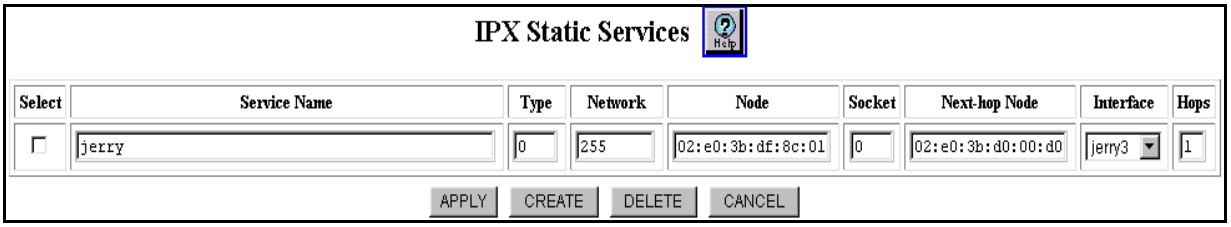
Parameter	Allows you to...
Service Name	Enter the IPX Static Service name. For example, FS_ENG01. Use SLIST (NetWare 3.x) or NLIST (NetWare 4.x) to list your current server names and types from your NetWare server. You may want to explicitly add a service so that clients always use the services of a particular server.
Type	Enter the service type (in hex) that identifies the type of IPX static service the server provides. Well-known service types include: • Unknown (0) • Print Queue (3) • File Server (4) • Job Server (5) • Print Server (7) • Archive Server (9) • Remote Bridge Server (24) • Advertising Print Server (47)
Network	Enter the IPX Static Service network number (in Hex).
Node	Enter the IPX Static Service node address (in hex). The format of the node value is **aa:bb:cc:dd:ee:ff**.
Socket	Enter the number (in hex) associated with a running process on the end node (for example, RIP, NLSP).
Next-Hop Node	Enter the MAC address of the next destination through which the service is reached. The format of the node value is **aa:bb:cc:dd:ee:ff**.
Interface	Enter the interface that corresponds to the next-hop node.
Hops	Enter the number of routers (hops) that the packet must pass through before reaching the network number associated with the service.

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

5. To view or modify IPX Static Services, expand the **Routing > IPX > Configuration** folders in the navigation pane, and then click **Static Services**.

   The IPX Static Services Web page is displayed in the content pane. See Figure 16-11.

Figure 16-11. IPX Static Services Web Page



CLI Command

To create IPX static services using the CLI, enter the following command in Configure mode:

`(configure)#` **ipx service** <***service-type type*** | ***range***>

Deleting IPX Static Service

You can delete an IPX static service using either the Web Agent or the CLI.

Web Agent Procedure

To delete an IPX static service using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **Configuration** folders, and then click **Static Services**.

   The IPX Static Services Web page is displayed in the content pane. See Figure 16-9.

2. Select the checkbox for the IPX Static Service that you want to remove.

3. Click **DELETE** to remove the IPX static service, or **CANCEL** to keep the IPX Static Service.

CLI Command

To delete an IPX static service using the CLI, enter the following command in Configuration mode:

`(configure)#` **clear ipx service {**<***service-type***> <***service-name***> | ***}**

Modifying IPX Static Service

You can modify an IPX static service using either the Web Agent or the CLI.

Web Agent Procedure

To modify an IPX static service using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **Configuration** folders, and then click **Static Services**.

   The IPX Static Services Web page is displayed in the content pane. See Figure 16-9.

2. Select the checkbox for the IPX Static Service that you want to modify from the Select column.

3. See Table 16-5 to modify the IPX Static Services Web page parameters.

CLI Command

To modify an IPX static service using the CLI, enter the following command in Configure mode:

`(configure)#` **ipx service** <***service-type***> <***service-name***> <***network***> <***node***> <***socket***> <***network.next-hop-node***> <***hops***>

Monitoring Your Switch Using IPX

This section explains how to use IPX to do the following on your switch:

- Displaying IPX Global Statistics
- Searching the IPX Route Table
- Displaying the IPX Route Table
- Displaying IPX Route Table Statistics
- Searching the IPX Service Table
- Displaying the IPX Service Table
- Displaying IPX Service Table Statistics

Displaying IPX Global Statistics

IPX Global Statistics count only the IPX packets that are received by or sent from the Supervisor module, not those packets routed in hardware.

You can display the IPX Global statistics using either the Web Agent or the CLI.

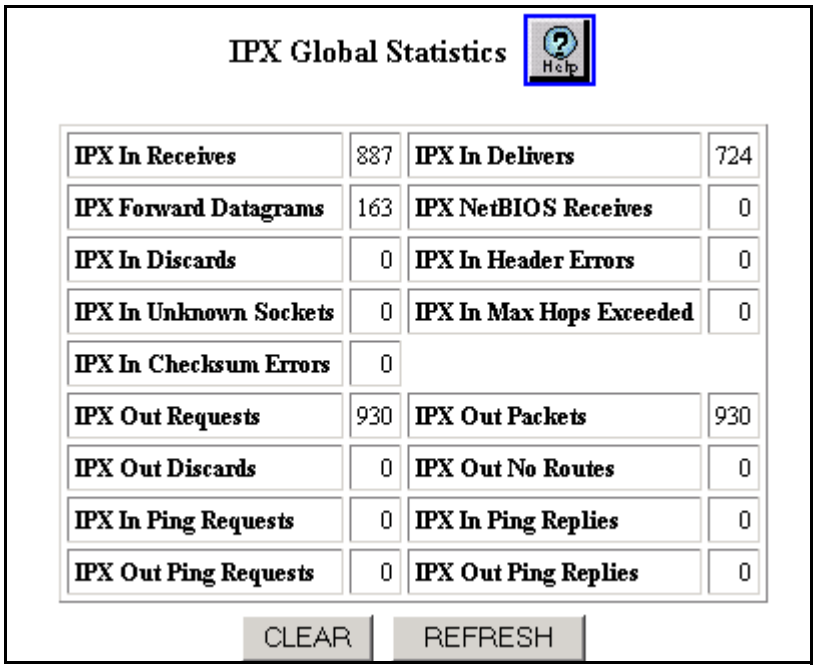
Web Agent Procedure

To display the global IPX routing statistics using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **Display** folders, and then click **Global Statistics**.

   The IPX Routing Global Statistics Web page is displayed in the content pane. See Figure 16-12.

Figure 16-12. IPX Global Statistics Dialogue Box



2. To view the latest statistics available, click **REFRESH** to update all statistics.

3. To reset all of the statistics that are currently displayed, click **CLEAR.** The statistics all reset to zero.

4. See Table 16-6 for a explanation of the IPX Global Statistics Web page parameters:

Table 16-6. IPX Global Statistics Dialogue Box Parameters

Parameter	Allows you to...
IPX In Receives	View the total number of IPX packets received (including errors).
IPX In Delivers	View the total number of IPX packets delivered locally. This includes packets from local applications.
IPX Forward Datagrams	View the number of IPX packets forwarded.
IPX Netbios Receives	View the number of NetBIOS packets received.
IPX In Discards	View the number of IPX packets received but discarded.
IPX In Header Errors	View the number of IPX packets discarded because of errors in the packet header. This includes packets that are less than the minimum 30 byte length.
	1 of 2

Table 16-6. IPX Global Statistics Dialogue Box Parameters

Parameter	Allows you to...
IPX In Unknown Sockets	View the number of IPX packets discarded because the destination socket was not open.
IPX In Max Hops Exceeded	View the number of IPX packets discarded because the Transport Control field is greater than or equal to 16.
IPX In Checksum Errors	View the number of IPX packets received with bad checksums.
IPX Out Requests	View the number of IPX packets supplied locally for transmission. This does not include any packets counted in IPX Forward Datagrams.
IPX Out Packets	View the total number of IPX packets transmitted.
IPX Out Discards	View the number of outgoing IPX packets discarded.
IPX Out No Routes	View the number of IPX packets that cannot be transmitted because no routes are available.
IPX In Ping Requests	View the number of received ping requests.
IPX In Ping Replies	View the number of received replies made to ping requests.
IPX Out Ping Requests	View the number of transmitted ping requests.
IPX Out Ping Replies	View the number of transmitted replies made to ping requests.
	2 of 2

CLI Command

To display the global IPX routing statistics use the following command:

> show ipx traffic

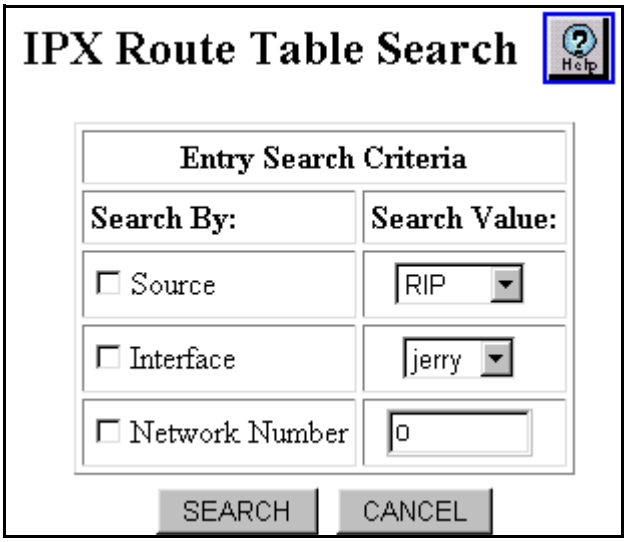
Searching the IPX Route Table

To search the IPX route table:

1. In the navigation pane, expand the **Routing > IPX > Display** folders, and then click **Route Table Search**.

   The IPX Route Table Search Web page is displayed in the content pane. See Figure 16-13.

Figure 16-13. IPX Route Table Search Web Page



2. Select the **SEARCH** Criteria you want to use to find information on specific routes. For example, if you want to find all static routes that are presently configured on your switch, search by source and specify static as your search value.

3. See Table 16-7 to determine the search parameters:

Table 16-7. IPX Route Table Search Parameters

Parameter	Allows you to...
Source	Search the IPX route table using of the following parameters: • **RIP** - RIP routing updates • **Static** - User configuration • **Local** - Directly connected routes Once you select a parameter, you can search to find the routes associated with that parameter.
Interface	Search the IPX Route table using the interface associated with the next-hop to the IPX network.
Network Number	Search the IPX Route table using the IPX Network address (hexadecimal number) you specify.

4. Click on the **Search** button to start:

   — If the available routes are found they are displayed in the IPX Route Table Web page.

   — If no routes are available, a message is displayed in the IPX Route Table Web page.

Displaying the IPX Route Table

To display the IPX route table:

1. Perform an IPX Route Table search as described in "Searching the IPX Route Table." If available routes are found they are displayed in the IPX Route Table Web page.
2. See Table 16-8 to review your configuration:

Table 16-8. IPX Route Table Parameters

Parameter	Defines the...
Select	Parameter selected.
Network	Network number (in hex) of the IPX network.
Interface	Interface associated with the IPX network.
Source	Method by which the network was learned. For example, RIP, local, or static.
TTL	Number of seconds before the route expires.
Ticks	Amount of time (in ticks) that the packet takes to reach the network number you specified. A tick is approximately 1/18th of a second.
Hops	Number of routers (hops) that the packet must pass through before reaching the network number associated with the IPX network.
Next-Hop MAC Address	MAC address of the next destination through which the network is reached.

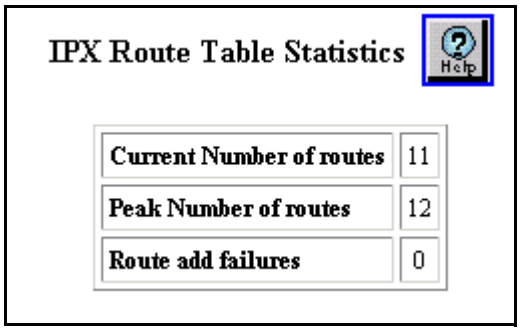
Displaying IPX Route Table Statistics

To display your IPX route table statistics:

1. In the navigation pane, expand the **Routing > IPX > Display** folders, and then click **Route Table Statistics**.

   The IPX Route Table Statistics Web page is displayed in the content pane. See Figure 16-14.

Figure 16-14. IPX Route Table Statistics Web Page



2. See Table 16-9 and review each statistic:

Table 16-9. IPX Route Table Statistics

Statistic	Definition
Current Number of Routes	Indicates the current number of IPX routes.
Peak Number of Routes	Lists the peak number of routes.
Route Add Failures	Indicates the number of failed attempts to add a route to the routing table.

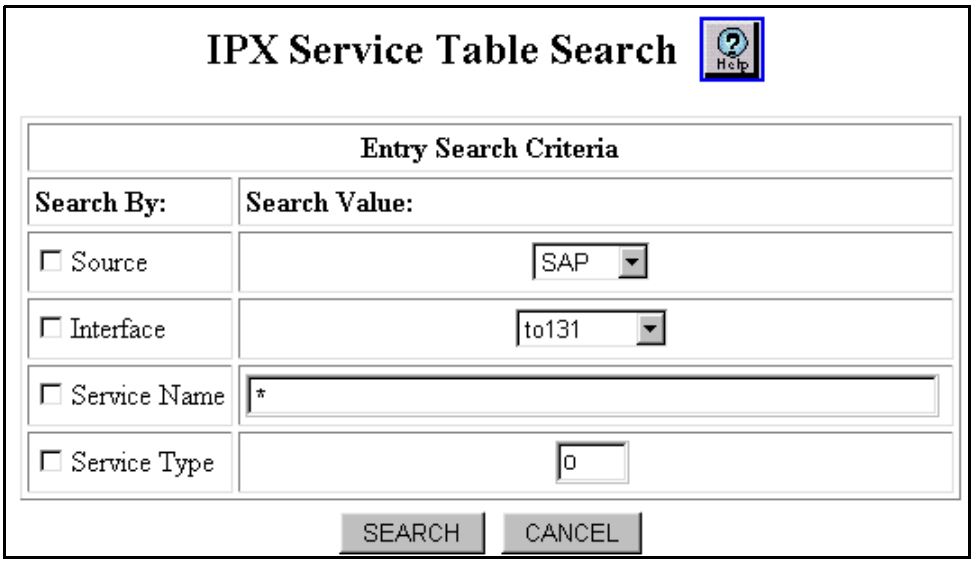
Searching the IPX Service Table

To search the IPX service table:

1. In the navigation pane, expand the **Routing > IPX > Display** folders, and then click **Service Table Search**.

The IPX Service Table Search Web page is displayed in the content pane. See Figure 16-15.

Figure 16-15. IPX Service Table Search Web Page



2. Select the search criteria you want to use to find information on specific services in the Search By column. For example, if you want to find all static services that are presently configured on your switch, select Source from the Search By column and specify Static as your search value.

3. See Table 16-10 to determine the search parameters:

Table 16-10. IPX Service Table Search Parameters

Parameter	Allows you to...
Source	Search the IPX Service Table using one of the following parameters: • **SAP** - Services learned via the SAP protocol • **Static** - User configuration • **Local** - Local services Once you select a parameter, you can search to find the routes associated with that parameter.
Interface	Search the IPX Service Table using the interface associated with the next-hop to the service.
	1 of 2

Table 16-10. IPX Service Table Search Parameters

Parameter	Allows you to...
Service Name	Search the IPX Service Table using a service name you specify. Note that you can specify a single asterisk (*) to indicate a wildcard character that will match all characters entered before the wildcard. For example, entering FS_ENG* will find all service names that start with FS_ENG.
Service Type	Search the IPX Service Table using the service type you specify. For example, to specify a print queue type, you would enter 3 (0003h) in the space provided.
	2 of 2

4. Click on the **SEARCH** to start the search.

   — If available services are found they are displayed in IPX Service Table Web page.

   — If services are not available, a message displays in the IPX Service Table Web page.

Displaying the IPX Service Table

To display the IPX service table:

1. Perform an IPX Service table search as described in "Searching the IPX Service Table" earlier in this chapter. If available services are found they are displayed in the IPX Service Table Web page.

2. See Table 16-11 to review each field in the IPX Services table:

Table 16-11. IPX Service Table Parameters

Parameter	Defines the...
Select	Parameter selected.
Name	Name of the service in question. For example, FS_ENG01. (Use SLIST (NetWare 3.x™) or NLIST (NetWare 4.x™) to list your current server names and types from your NetWare server.)
	1 of 2

Table 16-11. IPX Service Table Parameters

Parameter	Defines the...
Type	Service type that identifies the type of service the server provides. Well known service types include: • Unknown (0) • Print Queue (3) • File Server (4) • Job Server (5) • Print Server (7) • Archive Server (9) • Remote Bridge Server (24) • Advertising Print Server (47)
Network	Network number of the IPX service.
Node	Node address of the service.
Socket	Number associated with a running process on the end node (for example, RIP, NLSP).
Interface	Interface associated with the next hop to the service.
Source	Method by which the service was learned. For example, SAP or static.
TTL	Number of seconds before the service expires.
Hops	Number of routers (hops) that the packet must pass through before reaching the network number associated with the service.
Next-Hop MAC Address	MAC address of the next destination through which the service is reached.
	2 of 2

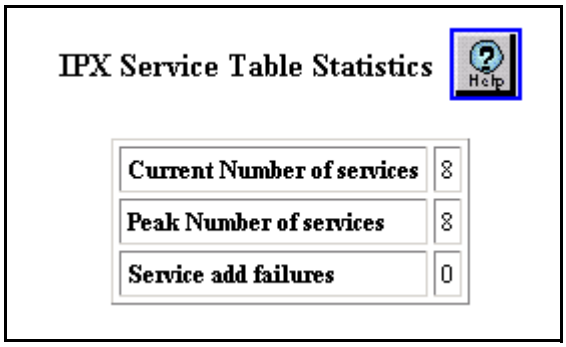
Displaying IPX Service Table Statistics

To display the IPX service table statistics:

1. In the navigation pane, expand the **Routing** > **IPX** > **Display** folders, and then click **Service Table Statistics**.

   The IPX Service Table Statistics Web page is displayed in the content pane. See Figure 16-16.

Figure 16-16. IPX Service Table Statistics



2. See Table 16-12 to review each statistic:

Table 16-12. IPX Service Table Statistics Web Page

Statistic	Defines the...
Current Number of Services	The current number of IPX services.
Peak Number of Services	The peak number of services.
Service Add Failures	The number of failed attempts to add a service to the routing table.

17 Configuring the IPX RIP Protocol

Overview

The information and procedures provided in this chapter pertain to layer 3 module configuration only. This chapter provides information about the Internetwork Packet Exchange (IPX) Routing Interface Protocol (RIP) and procedures for the following tasks:

- Configuring IPX RIP Interfaces
- Creating and Modifying IPX RIP Filters
- Viewing RIP Interface Statistics

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Configuring IPX RIP Interfaces

You can configure IPX RIP interfaces using either the Web Agent or the CLI.

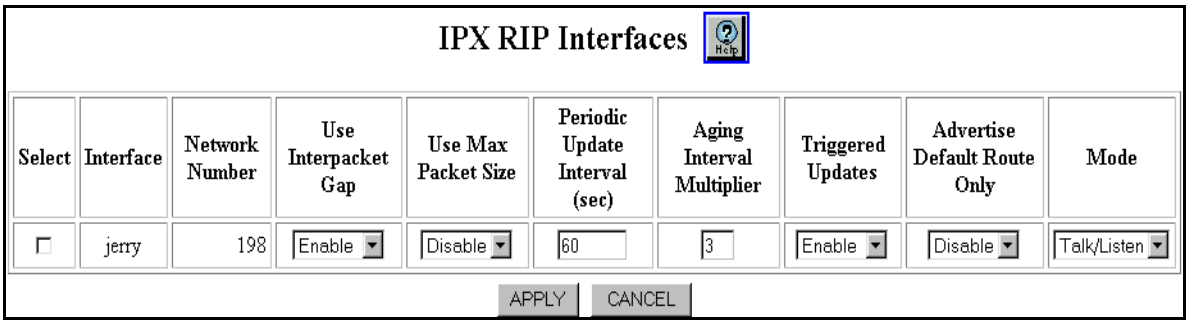
Web Agent Procedure

To configure IPX RIP interfaces using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **RIP** folders, and then click **Interfaces**.

   The IPX RIP Interfaces Web page is displayed in the content pane. See Figure 17-1. If no IPX interfaces are configured, then the `No Interfaces are currently configured` message is displayed.

Figure 17-1. IPX RIP Interfaces Web Page



2. See Table 17-1 to modify the IPX RIP Interfaces Web page parameters.

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 17-1. IPX RIP Interfaces Web Page Parameters

Parameter	Allows you to...
Select	Select a RIP interface to modify. **Note:** This field is displayed in the IPX RIP Interface Web page.
Interface	Identify the IPX interface associated with the RIP interface.
Network Number	Identify the number of the IPX network associated with the interface.
Use Interpacket Gap	Select **Disable** if you do not want RIP updates sent out over an interface to have an interpacket transmission delay. The default is **Enable**.
Use Max Packet Size	Enable or disable the transmission size of RIP packets sent out an interface: • **If** *enabled*, RIP packets can contain the maximum allowed by the MTU of the RIP interface. • **If** *disabled*, RIP packets are limited to 50 network entries. • The default is **Disabled**
Periodic Update Interval (sec)	Specify the length of time for the periodic update interval. The default is 60 sec.
Aging Interval Multiplier	Specify the length of time that information from received RIP updates are kept as a multiplier of the Periodic Update Interval. The default is **3**.
	1 of 2

Table 17-1. IPX RIP Interfaces Web Page Parameters

Parameter	Allows you to...
Triggered Updates	Select **Disable** to disallow RIP updates to be immediately transmitted to the network in response to changes in the network topology. The default is **Enable**.
Advertise Default Route Only	Select **Enable** to allow the advertising of the default network exclusively (subject to a route to the default network being known to the switch). The default is **Disable**.
Mode	Select the one of the following modes for the RIP interface: • **Talk/Listen** - Send and receive advertisements. • **Talk Only** - Send advertisements. • **Listen Only** - Receive advertisements. The default is **Talk/Listen**
	2 of 2

CL Command

To configure IPX RIP interfaces using the CLI, enter the following command in Interface mode:

`(config-if: <`*`if name`*`>)#` **ipx rip**

Creating and Modifying IPX RIP Filters

You can create and modify IPX RIP filters from either the Web Agent or the CLI.

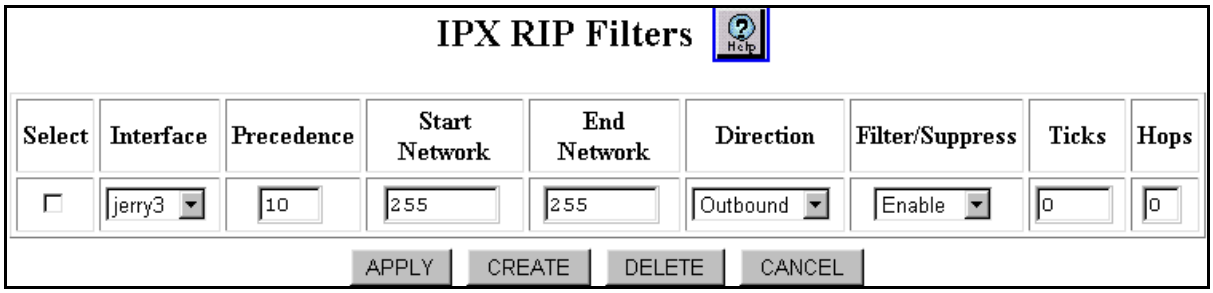
Web Agent Procedure

To create and modify IPX RIP filters using the Web Agent:

1. In the navigation pane, expand the **Routing > IPX > RIP** folders, and then click **Filters**.

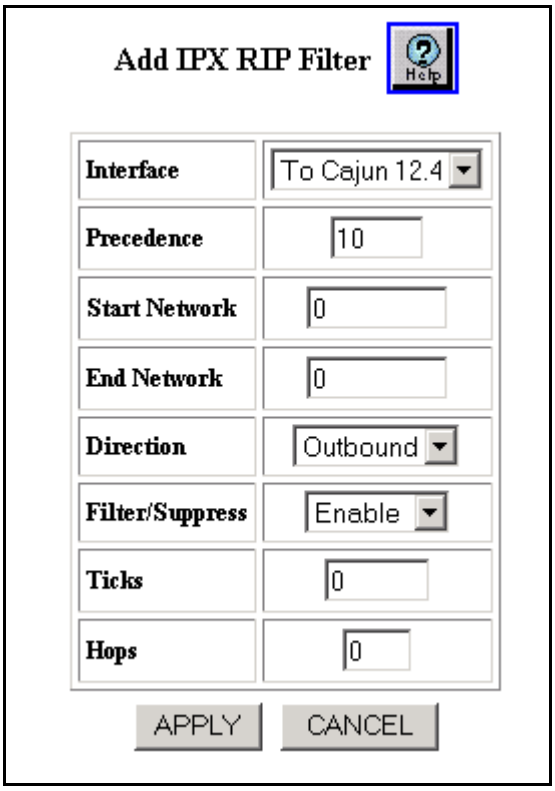
   The IPX RIP Filters Web page is displayed in the content pane. See Figure 17-2. Only the **CREATE** button is displayed when no filters are configured.

Figure 17-2. IPX RIP Filters Web Page



2. Select **CREATE** to create a new filter. The Add IPX RIP Filter Web page is displayed. See Figure 17-3.

Figure 17-3. Add IPX RIP Filter Web Page



3. See Table 17-2 to modify the Add IPX RIP Filter Web page parameters.

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 17-2. Add IPX RIP Filters Web Page Parameters

Parameter	Allows you to...
Select	The Select chechbox selects the IPX Interface. This box only appears with the IPX RIP Filters Web page.
Interface	Select the interface to which this filter will be applied to RIP packets sent and/or received on the interface.
Precedence	Specify the filter precedence (in order of importance) with 0 equal to most important. **Note:** All filters on the same interface, must be assigned a unique filter precedence.
Start Network	Specify the first IPX network number in the range.The default is **0**
End Network	Specify the last IPX network number in the range.The default is **0**
Direction	Select the direction of the filter from a pull-down menu. The pull-down filter menu options are: • **Inbound** - Apply the filter only to RIP packets received on the interface. • **Outbound** - Apply the filter only to RIP packets sent on the interface. • **Both** - Apply the filter to RIP packets both sent and received on the interface. • The default is **Outbound**
Filter/Suppress	Disable the IPX networks (within the specified range in the filter) to be filtered (inbound) or suppressed (outbound). • **Filtered** - Apply the filter only to RIP packets received (inbound direction) on the interface. • **Suppressed** - Apply the filter only to RIP packets sent (outbound direction) on the interface. **Enable** allows you to filter and suppress traffic and is the default. Select **Disable** to disable the filtering and suppression of traffic.The default is **Enable**
	1 of 2

Table 17-2. Add IPX RIP Filters Web Page Parameters

Parameter	Allows you to...
Ticks	Specify the time (in ticks) that the packet takes to reach the network number you specified. A tick is approximately 1/18th of a second. This entry is used to override the value in the RIP packet. If you enter **0**, there will be no override in the transmission of data on the network.The default is **o**
Hops	Specify the number of routers (hops) that the packet must pass through before reaching the network number associated with the IPX interface. This entry is used to override the value in the RIP packet. If you enter **0**, there will be no override in the transmission of data on the network. The default is **0**
	2 of 2

Example #1: Configuring your switch to suppress advertising of network 10

To configure your switch to suppress the advertising of network 10 on the IPX interface named Backbone, complete the following steps:

1. In the navigation pane, expand the **Routing > IPX > RIP** folders, and then click **Filters**.

   The IPX RIP Filters Web page is displayed in the content pane. See Figure 17-2.

2. Select **CREATE** to create a new filter. The Add IPX RIP Filter Web page is displayed. See Figure 17-3.

3. Configure the following parameters on the Add IPX RIP Filter Web page:

   a. Select **Backbone** from the Interface pull-down menu.

   b. Enter **0** in the **Precedence** field.

   c. Enter **10** in the **Start Network** field.

   d. Enter **10** in the **End Network** field.

   e. Open the **Direction** pull-down menu and select **Outbound**.

   f. Open the **Filter/Suppress** pull-down menu and select **Enable**.

   g. Enter **0** in the **Ticks and Hops** fields.

4. Click **Apply** to save your changes, or **CANCEL** to restore previous settings.

The filter created in this example ensures that all advertising of network 10 from the interface labeled Backbone is suppressed.

Example #2: Suppress Advertising of all Networks Except 10

To configure your switch to suppress the advertising of all networks except network 10 on the IPX interface named Backbone, you must first create two filters.

To create filter 1 from the Web Agent:

1. In the navigation pane, expand the **Routing > IPX > RIP** folders, and then click **Filters**.

   The IPX RIP Filter Web page is displayed in the content pane. See Figure 17-2.

2. Select **Create** to create a new filter. The Add IPX RIP Filter Web page is displayed.

3. Configure the following parameters:

   a. Select **Backbone from the Interface pull-down menu.**

   b. Enter **0** in the **Precedence** field.

   c. Enter **10** in the **Start Network** field.

   d. Enter **10** in the **End Network** field.

   e. Select **Outbound** from the **Direction** pull-down menu.

   f. Select **Disable** from the **Filter/Suppress** pull-down menu.

   g. Enter **0** in the **Ticks and Hops** fields.

Filter 1 ensures that the advertising of network 10 on the interface Backbone will not be suppressed. To suppress all other networks, proceed with the creation of Filter 2.

To create Filter 2 from the Web Agent:

1. Repeat steps 1- 2 in the creating Filter #1 procedure.

2. Configure the following parameters:

   a. Select **Backbone** from the **Interface** pull-down menu.

   b. Enter **1** in the **Precedence** field.

c. Enter **0** in the **Start Network** field.

d. Enter **ffffffff** in the **End Network** field.

e. Select **Outbound** from the **Direction** pull-down menu.

f. Select **Enable** from the **Filter/Suppress** pull-down menu.

g. Enter **0** in the **Ticks** and **Hops** fields.

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Filter 2 ensures that the advertising of all networks on the interface Backbone will be suppressed.

Together, Filter 1 and Filter 2 will act to suppress all network advertising with the *exception* of *network 10*. It is important to note that Filter 2 had the Precedence field set to **1** and Filter 1 had the Precedence field set to **0**. Any filter with a Precedence of 0 will always override a filter with a precedence of 1 or higher.

CLI Command

To create and modify IPX RIP filters using the CLI, enter the following command from **Configure** mode:

`(configure)#` **ipx rip-filter** <***precedence***> <***start-network***> <***end-network***> **{outbound | inbound | both} {filter | allow} [**<***filter-ticks***> **[**<***filter-hops***>**]]**

Viewing RIP Interface Statistics

You can view RIP Interface Statistics from either the Web Agent or the CLI.

Web Agent Procedure

To view **RIP Interface Statistics** using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **RIP** folders, and then click **Interface Statistics**.

   The RIP Interface Statistics Web page is displayed in the content pane.

2. See Table 17-3 for an explanation of the IPX RIP interface Statistics.

3. Click **CLEAR** to remove the statistics, or **REFRESH** to access current interfaces statistics.

Table 17-3. IPX RIP Interface Statistical Parameters

Parameter	Definition
Interface	The IPX interface associated with this RIP interface.
State	The current state of the RIP interface.
Network Number	The network number of the IPX network associated with the interface.
Triggered Updates Sent	The number of triggered updates sent from the RIP interface.
Non-triggered Updates Sent	The number of non-triggered updates sent from the RIP interface.
Updates Received	The number of updates received. Updates may be received periodically even if no changes have occurred.
Requests Received	The number of requests received on the RIP interface.
Bad Packets Received	The number of bad packets received on the RIP interface.

CLI Command

To interpret IPX RIP interface statistics using the CLI, enter the following command from Interface mode:

`(config-if <interface name>)#` **show ipx interface [<*intf-name*>]**

18 Configuring the IPX SAP Protocol

Overview

The information and procedures provided in this chapter pertain to layer 3 module configuration only. This chapter provides information about the Internetwork Packet Exchange (IPX) Service Advertising Protocol (SAP) and procedures for the following tasks:

- Configuring IPX SAP Interfaces
- Creating IPX SAP Filters
- Interpreting IPX SAP Interface Statistics

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Configuring IPX SAP Interfaces

You can configure IPX SAP using either the Web Agent or the CLI.

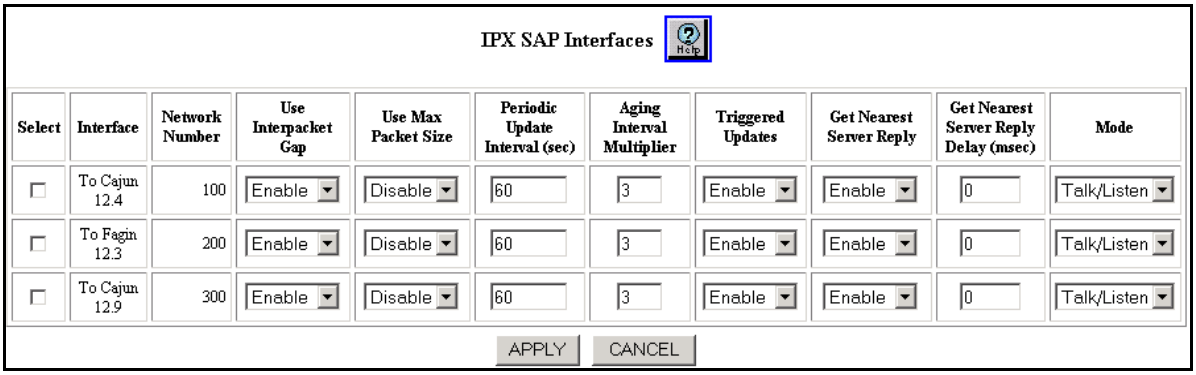
Web Agent Procedure

To configure the IPX SAP interfaces using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **SAP** folders, and then click **Interfaces**.

   The IPX SAP Interfaces Web page is displayed in the content pane. See Figure 18-1. If no IPX interfaces are configured, then the `No Interfaces are currently configured` message is displayed.

Figure 18-1. IPX SAP Interfaces Web Page



2. See Table 18-1 to configure the IPX SAP Interface Web page parameters.

3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 18-1. IPX SAP Interface Web Page Parameters

Parameter	Allows you to...
Select	Select a SAP interface to modify.
Interface	Identifies the IPX interface associated with the SAP interface.
Network Number	Identifies the number of the IPX network associated with the interface.
Use Interpacket Gap	Select one of the following from the pull-down menu: • **Disable –** if you do not want the IPX SAP periodic update packets to be sent out over an interface with no interpacket transmission delay. • **Enable** – If you want to send the IPX SAP periodic update packets to be sent out over an interface with an interpacket transmission delay (default).
Use Max Packet Size	Select one of the following from the pull-down menu: • **Disable –** To limit SAP packets to 7 service entries(default). • **Enable** – To allow SAP) packets to contain the maximum number of services entries that are allowed by the SAP Interface MTU (default).
Periodic Update Interval (sec)	Enter (in seconds) the interval at which periodic SAP updates are sent out an interface. Default is 60 seconds.
	1 of 2

Table 18-1. IPX SAP Interface Web Page Parameters

Parameter	Allows you to...
Aging Interval Multiplier	Enter the length of time that information from received SAP updates are kept as a multiplier of the Periodic Update Interval.Default is 3.
Triggered Updates	Select one of the following from the pull-down menu: • **Disable –** To prevent SAP updates from being immediately transmitted to the network in response to changes in the network topology • **Enable** – To allow SAP updates to be immediately transmitted to the network in response to changes in the network topology (default).
Get Nearest Server Reply	Select one of the following from the pull-down menu: • **Disable –** To prevent the router from responding to **Get Nearest Server** requests received on the interface • **Enable** – To allow the router to respond to **Get Nearest Server** requests received on the interface (default).
Get Nearest Server Reply Delay(msec.)	Specify the delay (in msecs) to wait before responding to a Get Nearest Service request received on the interface.Default is 0.
Mode	Select one of the following from the pull-down menu: • **Talk/Listen** - Send and receive advertisements (default). • **Talk Only** - Send advertisements. • **Listen Only** - Receive advertisements.
	2 of 2

CLI Command

To configure the IPX SAP interfaces using the CLI, enter the following command from Interface mode:

`(config-if <`*`interface-name`*`>)#` **ipx sap**

Creating IPX SAP Filters

This section provides the following procedures:

- Creating IPX SAP Name Filters
- Creating IPX SAP Network Filters

Creating IPX SAP Name Filters

You can create IPX SAP name filter suing either the Web Agent or the CLI.

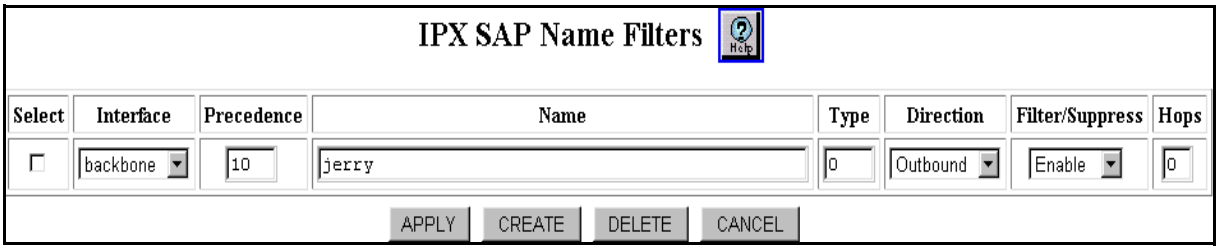
Web Agent Procedure

To create IPX SAP name filters using the Web Agent:

1. In the navigation pane, expand the **Routing > IPX > SAP** folders, and then click **Name Filters.**

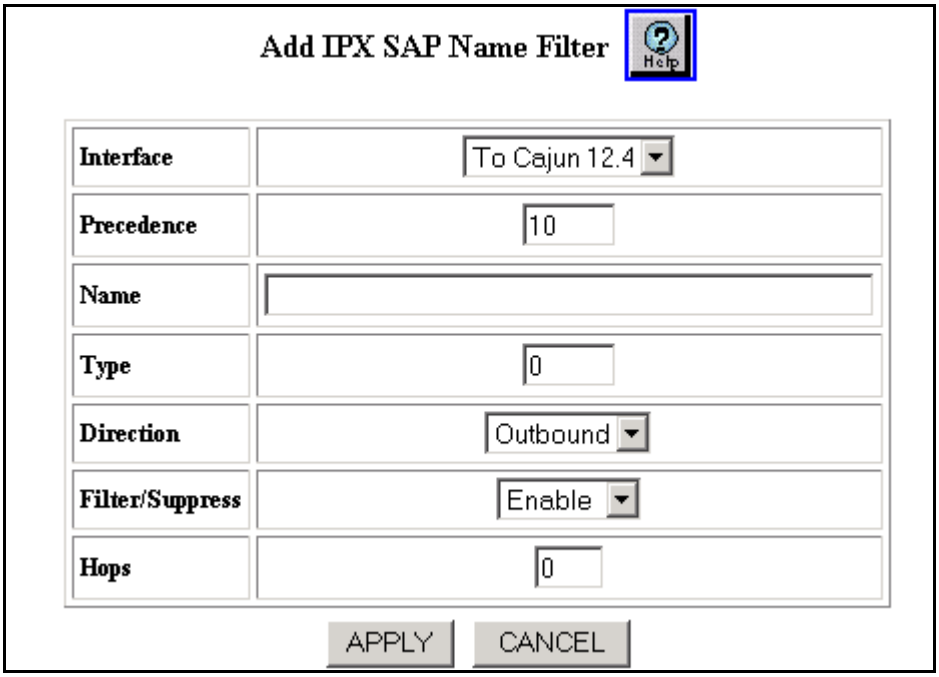
   The IPX SAP Name Filters Web page is displayed in the content pane. See Figure 18-2. Only the **CREATE** button is displayed when no SAP filters are configured.

Figure 18-2. IPX SAP Name Filters Web Page



2. Select **CREATE** to create a new filter. The Add IPX SAP Name Filter Web page is displayed. See Figure 18-3.

Figure 18-3. Add IPX SAP Name Filter Web Page



3. Configure an IPX SAP name filter. See Table 18-2 for an explanation of the Add IPX SAP Name Filter Web page parameters

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 18-2. IPX SAP Name Filter Web Page Parameters

Parameter	Allows you to...
Interface	Select the interface to which this filter will be applied to SAP packets sent and/or received on the interface from the pull-down menu.
Precedence	Enter the filter precedence (in order of importance) with 0 equal to most important. **Note:** All SAP filters on the same interface must be assigned a unique precedence value.
Name	Enter a service name. For example, FS_ENG001. A single asterisk may be present as the last character, which will match all remaining characters of a service name.
Type	Enter the service type that identifies the type of service the server provides. Well-known service types include: • Unknown (0) • Print Queue (3) • File Server (4) • Job Server (5) • Print Server (7) • Archive Server (9) • Remote Bridge Server (24) • Advertising Print Server (47) • NetWare Directory Services [NDS] (278) A value of ffff matches all service type values.
Direction	Select the filter direction from the following pull-down menu options: • **Inbound** - Applies the filter only to SAP packets received on the interface. • **Outbound** - Applies the filter only to SAP packets sent on the interface (Default) • **Both** - Applies the filter to SAP packets both sent and received on the interface.

1 of 2

Table 18-2. IPX SAP Name Filter Web Page Parameters

Parameter	Allows you to...
Filter/Suppress	Select to enable or disable whether the services matching Name and Type are filtered (inbound) or suppressed (outbound). • **Filtered** - Applies the filter only to SAP packets received (inbound direction) on the interface. • **Suppressed** - Applies the filter only to SAP packets sent (outbound direction) on the interface. Select **Enable** to filter/suppress traffic. Select **Disable** to disable the filtering and suppression of traffic.Default value is Enable.
Hops	Specify the number of routers (hops) that the packet must pass through before reaching the service(s) matched by the filter. This entry is used to override the value in the SAP packet. Entering 0 ensures that there is no override in the transmission of data on the network.
	2 of 2

Example: Suppress Advertising of all Print Servers on Interface Remote

To configure your switch to suppress the advertising of all Print Servers (for example, type 7) on Interface Remote:

1. In the navigation pane, expand the **Routing > IPX > SAP** folders, and then click **Name Filters**.

   The IPX SAP Name Filters Web page is displayed in the content pane. See Figure 18-2.

2. Click **CREATE** to create a new filter. The Add IPX SAP Name Filter Web page is displayed. See Figure 18-3.

3. Configure the following parameters:

   a. Select **Remote** from the **Interface** pull-down menu.

   b. Enter 0 in the Precedence field.

   c. Enter * (asterisk) in the Name field. The asterisk represents a wildcard that applies to all server names.

   d. Enter 7 in the **Type** field.

   e. Select **Outbound** from the **Direction** pull-down menu.

f. Select **Enable** from the **Filter/Suppress** pull-down menu.

g. Enter **0** in the **Hops** field. This ensures that there is no override in the transmission of data on the network.

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

*** Note:** This filter ensures that all advertising of any known print server on Interface Remote will be suppressed.

CLI Command

To create IPX SAP name filters using the CLI, enter the following command from Configure mode:

`(configure)#` **ipx sap-name-filter** <***precedence***> <***filter-name***> <***service-type***> **{outbound | inbound | both} {filter | allow} [**<***filter-hops***>**]**

Creating IPX SAP Network Filters

You can create IPX SAP network filters using wither the Web Agent or the CLI.

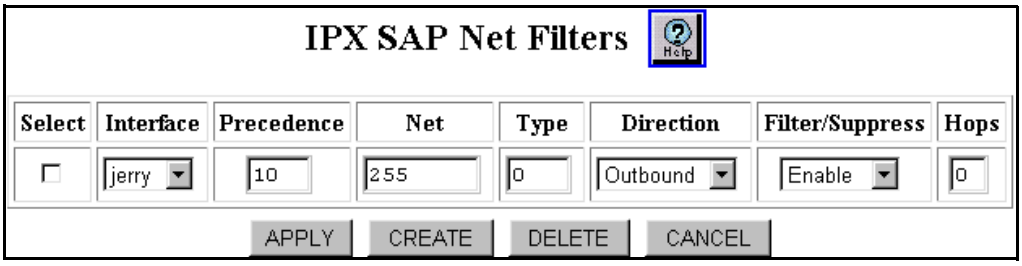
Web Agent Procedure

To create IPX SAP Network filters using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **SAP** folders, and then click **Net Filters**.

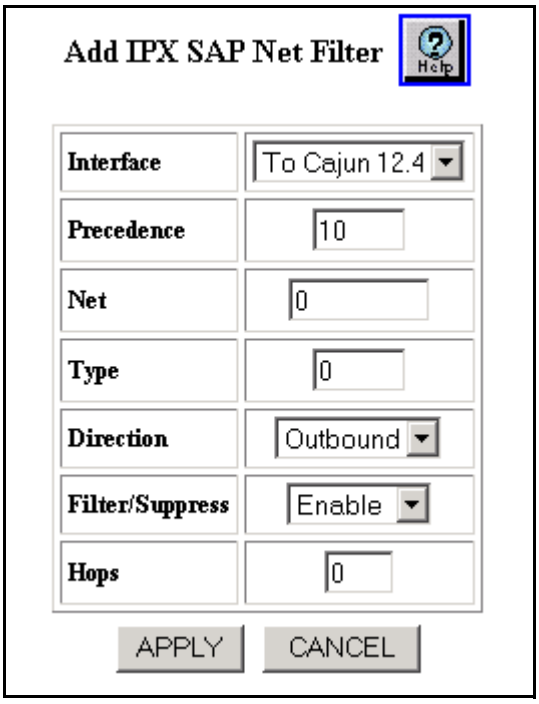
The IPX SAP Net Filters Web page is displayed in the content pane. See Figure 18-4.Only the **CREATE** button is displayed when no IPX SAP Network Filters are configured.

Figure 18-4. IPX SAP Net Filters Web Page



2. Select **CREATE**. The Add IPX SAP Net Filter Web page is displayed. See Figure 18-5.

Figure 18-5. Add IPX SAP Net Filter Web Page



3. See Table 18-3 to configure the Add IPX SAP Net Filters Web page parameters.

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Table 18-3. Add IPX SAP Network Filter Web Page Parameters

Parameter	Allows you to...
Select	Select the IPX interface. **Note:** This field is displayed in the IPX SAP Net Filter Web page, not in the Add IPX SAP Net Filter Web page.
Interface	Select the interface to which this filter will be applied to SAP packets sent and/or received on the interface.
Precedence	Specify the filter precedence (in order of importance) with 0 equal to most important. **Note:** All SAP filters on the same interface must be assigned a unique filter precedence.
Net	Specify the network on which the server resides. A network number of ffffffff represents all networks.
	1 of 2

Table 18-3. Add IPX SAP Network Filter Web Page Parameters

Parameter	Allows you to...
Type	Specify the service type (in hex) that identifies the type of service the server provides. Well-known service types include: • Unknown (0) • Print Queue (3) • File Server (4) • Job Server (5) • Print Server (7) • Archive Server (9) • Remote Bridge Server (24) • Advertising Print Server (47) • NetWare Directory Services (278) A value of ffff matches all service type values.
Direction	Select the direction of the filter in question. Filter choices include: • **Inbound** - Applies the filter only to SAP packets received on the interface. • **Outbound** - Applies the filter only to SAP packets sent on the interface. • **Both** - Applies the filter to SAP packets both sent and received on the interface Default value is outbound.
Filter/Suppress	Disable whether the services matching Net and Type are filtered (inbound) or suppressed (outbound). • **Filtered** - Applies the filter only to SAP packets received (inbound direction) on the interface. • **Suppressed** - Applies the filter only to SAP packets sent (outbound direction) on the interface. Select **Disable** to disable the filtering and suppression of traffic. The default is **Enable**.
Hops	Specify the number of routers (hops) that the packet must pass through before reaching the service(s) matched by the filter. This entry is used to override the value in the SAP packet. Entering 0 ensures that there is no override in the transmission of data on the network.
	2 of 2

Filtering all Services Except Netware Directory Services (NDS)

To configure your switch to filter all services except NDS a specific Interface, you must create two filters. Together, Filter 1 and Filter 2 filter services learned on the indicated interface except for NDS advertisements.

To create these filters using the Web Agent:

Filter 1

This filter ensures that all NDS packets received are not filtered on Interface Remote.

1. In the navigation pane, expand the **Routing > IPX > SAP** folders, and then click **Net Filters**.

   The IPX SAP Net Filters Web page is displayed in the content pane. See Figure 18-4.

2. Select **CREATE**. The Add IPX SAP Net Filter Web page is displayed. See Figure 18-5.

3. Configure the following parameters:

   a. Select an interface name from the **Interface** pull-down menu.

   b. Enter 0 in the Precedence field.

   ***Note:** Any filter with a Precedence of 0 will always override a filter with a precedence of 1 or higher.

   c. Enter **ffffffff** (which represents all networks) in the **Net** field.

   d. Enter **278** (which represents the type for NDS) in the **Type** field.

   e. Select **Inbound** from the **Direction** pull-down menu.

   f. Select **Disable** from the **Filter/Suppress** pull-down menu to ensure that NDS advertisements are not filtered.

   g. Enter 0 in the **Hops** field to ensure that there is no override in the transmission of data on the network.

4. Click **APPLY** to save your changes, or CANCEL to restore previous settings.

Filter 2

Filter 2 ensures that all networks and service types are filtered for the indicated interface.

1. In the navigation pane, expand the **Routing > IPX > SAP** folders, and then click **Net Filters**.

   The IPX SAP Net Filters Web page is displayed in the content pane. See Figure 18-4.

2. Click **CREATE** to create a new filter. The Add IPX SAP Net Filter Web page is displayed. See Figure 18-5.

3. Configure the following parameters:

   a. Select an interface name from the **Interface** pull-down menu.

   b. Enter 1 in the **Precedence** field.

   c. Enter **ffffffff** (which represents all networks) in the Net field.

   d. Enter **ffff** (which represents all services/types) in the Type field.

   e. Select **Inbound** from the **Direction** pull-down menu.

   f. Select **Enable** from the **Filter/Suppress** pull-down menu.

   g. Enter 0 in the **Hops** field. Entering 0 ensures that there is no override in the transmission of data on the network.

4. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To create IPX SAP Network filters using the CLI, enter the following command from Interface mode:

`(config if-<interface name>)#` **ipx sap-network-filter <*precedence*> <*filter-network*> <*service-type*> {outbound | inbound | both} {filter | allow} [<*filter-hops*>]**

Interpreting IPX SAP Interface Statistics

You can interpret IPX SAP interface statistics using either the Web Agent or the CLI.

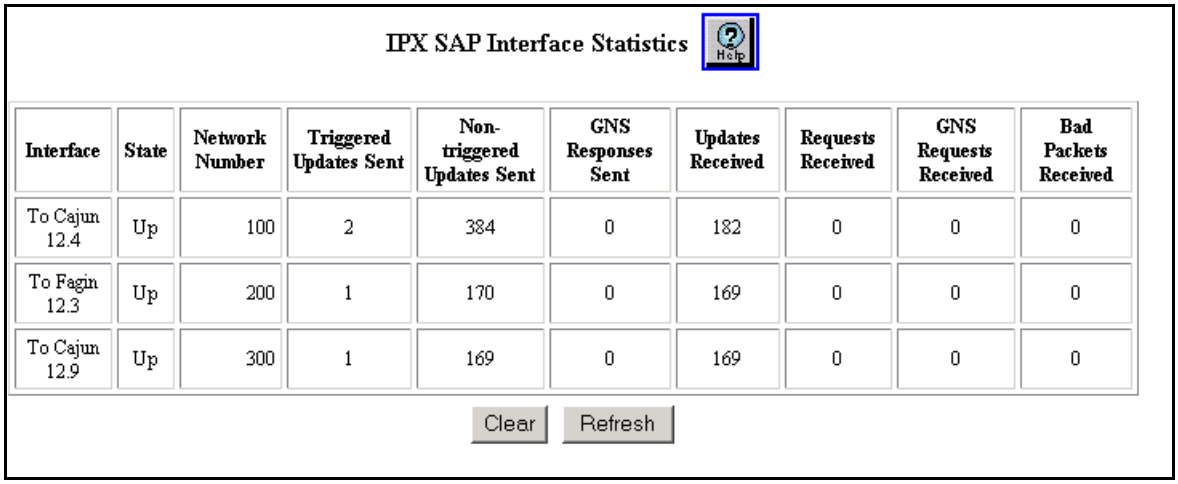
Web Agent Procedure

To interpret IPX SAP interface statistic using the Web Agent:

1. In the navigation pane, expand the **Routing** > **IPX** > **SAP** folders, and then click **Interface Statistics**.

   The IPX SAP Interface Statistics Web page is displayed in the content pane. See Figure 18-6.

Figure 18-6. IPX SAP Interface Statistics Web Page



2. See Table 18-4 for an explanation of the **IPX SAP interface** statistics.

3. Click on **CLEAR** to remove the statistics, or **REFRESH** to access current interface statistics.

Table 18-4. IPX SAP Interface Statistical Parameters

Parameter	Definition
Interface	The IPX interface associated with this SAP interface.
State	The current state of the SAP interface.
Network Number	The network number of the IPX network associated with the interface.
Triggered Updates Sent	The number of triggered updates sent from the SAP interface.
	1 of 2

Table 18-4. IPX SAP Interface Statistical Parameters

Parameter	Definition
Non-triggered Updates Sent	The number of non-triggered updates sent from the SAP interface.
GNS Responses Sent	The number of GNS responses sent from the SAP interface.
Updates Received	The number of updates received. Updates may be received periodically even if no changes have occurred.
Requests Received	The number of requests received on the SAP interface.
GNS Requests Received	The number of GNS requests received on the SAP interface.
Bad Packets Received	The number of bad packets received on the SAP interface.
	2 of 2

CLI Command

To interpret IPX SAP interface statistics using the CLI, enter the following command from Interface mode:

`(config-if <interface name>)#` **show ipx interface [<*intf-name*>]**

19 Configuring AppleTalk Routing

Overview

The following information and procedures provided in this chapter pertain to layer 3 module configuration only:

- Implementation
- Configuring AppleTalk Routing
- Viewing AppleTalk Statistics

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Implementation

This section introduces AppleTalk routing, explains how it operates, and explains why you would want to configure it on the Avaya P580 or P882 Multiservice switch.

What is AppleTalk Routing?

AppleTalk Phase I was originally designed for local work groups. AppleTalk Phase II extends the number of nodes in an internetwork to over 16 million and the number of zones per port to 254. The Avaya switch supports both AppleTalk Phase I and Phase II. However, the translation from AppleTalk Phase I to Phase II is not supported.

*** Note:** The Avaya P580 and P882 Multiservice switches support AppleTalk over Ethernet only.

The supported Ethernet versions are:

- Ethernet SNAP
- Ethernet II

The Avaya P580 and P882 Multiservice switches support the following AppleTalk protocols:

- AppleTalk Address Resolution Protocol (AARP)

  This is an AppleTalk support protocol that maps the hardware address of an AppleTalk node to an Appletalk protocol address.

- Routing Table Maintenance Protocol (RTMP)

  This protocol maintains information about AppleTalk addresses and connections between different networks. Specially, it tells each router to:

  — Learn a new route to other routers

  — Delete a route if the local router has not broadcast that route to the network for a certain period of time.

- Name-Binding Protocol (NBP)

  This protocol translates alphanumeric entity names to AppleTalk addresses. NBP maintains a table of node addresses and entities within each node. Because each node also maintains it own list of named entities, the name directory within an AppleTalk network is not centralized. The names directory database is distributed among all nodes on the intranetwork.

- Datagram Delivery Protocol (DDP)

  This protocol transfers data in packets called datagrams.

- AppleTalk Echo Protocol (AEP)

  This protocol is used to send datagrams to other nodes in the network. The transmitted AEP datagram causes the destination node to return, or echo, that datagram to the sending node. This protocol determines whether a node is accessible before any sessions are started, and it enables users to estimate the round-trip delay time between nodes.

- Zone Information Protocol (ZIP)

  AppleTalk routers use this protocol to map network numbers to zones. Each AppleTalk router maintains a Zone information Table which lists the zone-to-network mapping information.

For more detailed information about these protocols, see the AppleTalk documentation.

Need for AppleTalk Routing

If you configure your Avaya P580 or P882 Multiservice switch into a network running AppleTalk routing, you should be aware of the following:

- Appletalk is built into all Apple devices, thus making them automatically network capable. This makes Apple an extremely easy network system to install and maintain.
- Due to the naming mechanism that AppleTalk uses, users do not have to understand how AppleTalk works.
- AppleTalk supports a peer-to-peer network, thus dedicated servers or centralized network are not required.
- AppleTalk is plug-and-play (auto-configuring). Therefore, users can plug an Appletalk device into an AppleTalk network and use it immediately. No network configuration or assignments of network addresses are needed when you add a device to an AppleTalk network.
- AppleTalk support zones, which makes it easier for network administrators to define work groups that consist of users and services that can span multiple networks segments.

Configuring AppleTalk Routing

This section provides procedures for:

- Enabling AppleTalk Global Routing
- Creating an AppleTalk Routing Interface
- Editing AppleTalk Interfaces
- Deleting an AppleTalk Interface
- Creating an AppleTalk Static Route
- Editing AppleTalk Static Routes
- Deleting an AppleTalk Static Route
- Creating an AppleTalk Name-Binding Protocol (NBP) Filter
- Editing an AppleTalk NBP Filter
- Adding or Deleting Interfaces to an NBP Filter
- Creating an AppleTalk Zone Filter
- Editing an AppleTalk Zone Filter
- Adding or Deleting Interfaces to a Zone Filter

Enabling AppleTalk Global Routing

To enable AppleTalk routing globally:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Global Configuration**.

   The AppleTalk Routing Global Configuration Web page is displayed in the content pane. See Figure 19-1.

*** Note:** AppleTalk Routing is disabled by default.

Figure 19-1. AppleTalk Routing Global Configuration Web Page



2. Select **Enable** from the **AppleTalk Routing** pull-down menu.
3. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

Creating an AppleTalk Routing Interface

You can create an AppleTalk routing interface using the Web Agent or CLI.

Web Agent Procedure

To create an AppleTalk routing interface using the Web Agent:

1. Enable AppleTalk routing globally. See "Enabling AppleTalk Global Routing" earlier in this chapter.
2. In the navigation pane, expand the **Routing > AppleTalk > Configuration** folders, and then click **Interfaces**.

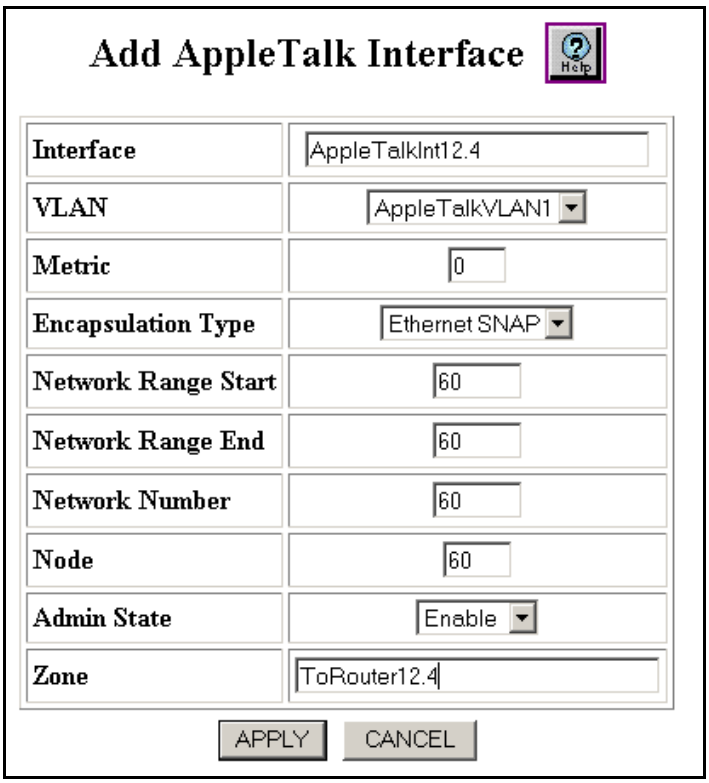
   The AppleTalk Interfaces Web page is displayed in the content pane. See Figure 19-2.

Figure 19-2. AppleTalk Interfaces Web Page



3. Select **CREATE**. The Add AppleTalk Interfaces Web page is displayed. See Figure 19-3.

Figure 19-3. Add AppleTalk Interfaces Web Page



4. See Table 19-1 to configure the Add AppleTalk Interface Web page parameters:

Table 19-1. Add AppleTalk Interface Web Page Parameters

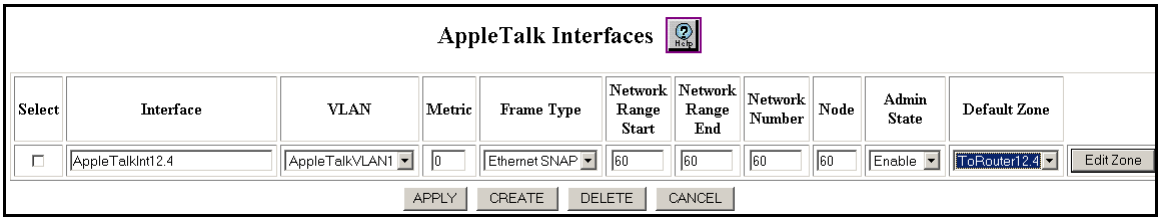
Parameter	Definition
Interface	Enter the name of the AppleTalk interface to be created. 1 to 32 Characters
VLAN	Select the VLAN to be associated with the AppleTalk interface.
Metric	Enter the metric associated with the AppleTalk interface.
Encapsulation Type	Select the encapsulation type to be associated with the AppleTalk interface.The options include: • **Ethernet II** - MTU = 1500 • **Ethernet SNAP** - MTU =1492
	1 of 2

Table 19-1. Add AppleTalk Interface Web Page Parameters

Parameter	Definition
Network Range Start	Enter the starting network number. The network number specifies the range of AppleTalk network numbers for extended networks. Each number in the range must be an integer between 0 and 65279. **Note:** The Network Range Start value must be less than or equal to the Network Range End value.
Network Range End	Enter the ending network number. The network number specifies the range of AppleTalk network numbers for extended networks. Each number in the range must be an integer between 0 and 65279. **Note:** If the Network Range Start value equals 0, then the Network Range End value must also equal 0.
Network Number	Enter the interface network number. This number specifies the network number the interface is using. **Note:** This value must be within the Network Range or be equal to 0.
Node	Enter the interface node identification number. This number must be between 1 and 253. **Note:** Only if the Network Number is equal to 0, can the Node also be equal to 0.
Admin State	Select whether to enable or disable the administrator state. The administrator state determines if the interface is operational from a management point of view.
Zone	Select the default AppleTalk zone to be used by this interface. Zone names may be up to 32 characters in length.
	2 of 2

5. Click **Apply** to add the new AppleTalk interface, or **CANCEL** to restore previous settings. Once you click Apply, you are returned to the AppleTalk Interfaces Web page (Figure 19-4)

Figure 19-4. AppleTalk Interfaces



CLI Commands

Use the following CLI commands in interface mode to create an AppleTalk interface:

- `(config-if:<interface-name>)#` **appletalk cable-range** <***cable-range***>
- `(config-if:<interface-name>)#` **appletalk address** <***network.node***>
- `(config-if:<interface-name>)#` **appletalk zone** <***zone-name***>
- `(config-if:<interface-name>)#` **appletalk vlan** <***vlan-id***>

Editing AppleTalk Interfaces

You can edit AppleTalk interfaces using either the Web Agent or the CLI.

Web Agent Procedure

To edit an AppleTalk interface using the Web Agent:

*** Note:** You must enable AppleTalk globally and create an AppleTalk interface before you can edit an interface. See "Enabling AppleTalk Global Routing" and "Creating an AppleTalk Routing Interface" earlier in this chapter.

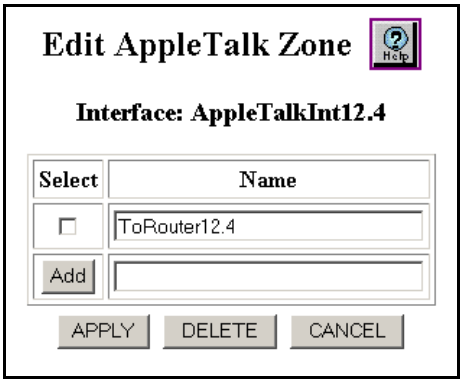
1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Interfaces**.

   The AppleTalk Interfaces Web page is displayed in the content pane. See Figure 19-4.

2. If there is no zone to select or if you want to edit the zones available, complete this step through step x. If not, skip to step x+1. Select **Edit Zone** to edit the AppleTalk zone for this interface.

   The Edit AppleTalk Zone Web page is displayed. See Figure 19-5.

Figure 19-5. Edit AppleTalk Zone Web Page



3. To add a new zone, enter the new AppleTalk zone name in the **Add text** field and click Add. Repeat this step as needed for each new zone on this interface.

4. To change a zone name(s), place a check mark(s) in the **Select** column and make the change(s) to the zone name(s). Click **Apply** to save the new AppleTalk zone, or either **Delete** or **Cancel** to remove the new zone.

5. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Interfaces**.

   The AppleTalk Interfaces Web page is redisplayed in the content pane. See Figure 19-4.

6. Select the new zone that you just created from the **Default Zone** field pull-down menu.

7. Make sure there is a check mark next to the interface(s) you will modify.

8. Select the VLAN to be associated with the interface from the **VLAN** field pull-down menu.

9. Enter the new port metric to be associated with the interface in the **Metric** field.

10. Select the new frame type to be associated with the interface from the **Frame Type** pull-down menu.

11. Enter the new network range start number in the **Network Range Start** field.

12. Enter the new network range end number in the **Network Range End** field.

13. Enter the new network number to be associated with the interface in the **Network Number** field.

14. Enter the new node number to be associated with the interface in the **Node** field.

15. Select **Enable** from the **Admin. State** field pull-down menu.

16. Select the new default zone associated with the interface from the **Default Zone** field pull-down menu.

17. Select the AppleTalk interface that you want to update in the **Select** column.

18. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

***Note:** You may select and change multiple interfaces simultaneously with one APPLY operation.

CLI Command

To edit an AppleTalk interface using the CLI, enter the following command from Interface mode:

`(config-if:<interface-name>)#` **appletalk** <***options***>

Deleting an AppleTalk Interface

You can delete an AppleTalk interface using either the Web Agent or the CLI.

Web Agent Procedure

To delete an AppleTalk interface using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Interfaces**.

   The AppleTalk Interfaces Web page is displayed in the content pane. See Figure 19-4.

2. Select the AppleTalk interface that you want to delete from the **Select** column.

3. Click **DELETE** to remove the AppleTalk interface.

CLI Commands

Use the following CLI commands in interface mode to delete an AppleTalk interface:

- `(config-if:<interface-name>)#` **no appletalk cable-range** <***cable-range***>
- `(config-if:<interface-name>)#` **no appletalk address** <***network.node***>
- `(config-if:<interface-name>)#` **no appletalk zone** <***zone-name***>
- `(config-if:<interface-name>)#` **no appletalk vlan** <***vlan-id***>

Creating an AppleTalk Static Route

You can create an AppleTalk static route using the Web Agent or the CLI.

Web Agent Procedure

To create an AppleTalk static route using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Static Route**.

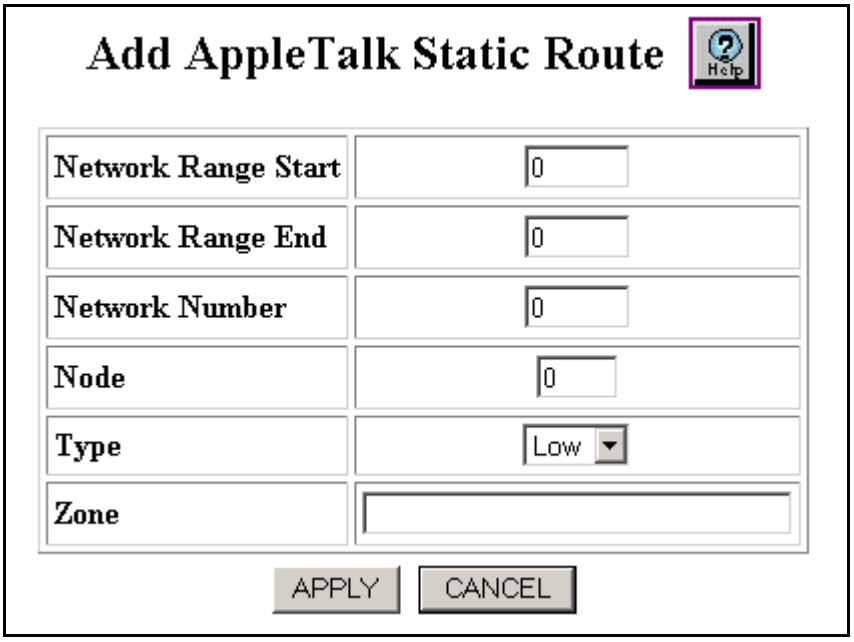
   The AppleTalk Static Route Web page is displayed in the content pane. See Figure 19-6.

Figure 19-6. AppleTalk Static Route Web Page



2. Select **CREATE** to add a new static route. The Add AppleTalk Static Route Web page is displayed. See Figure 19-7.

Figure 19-7. Add AppleTalk Static Route Web Page



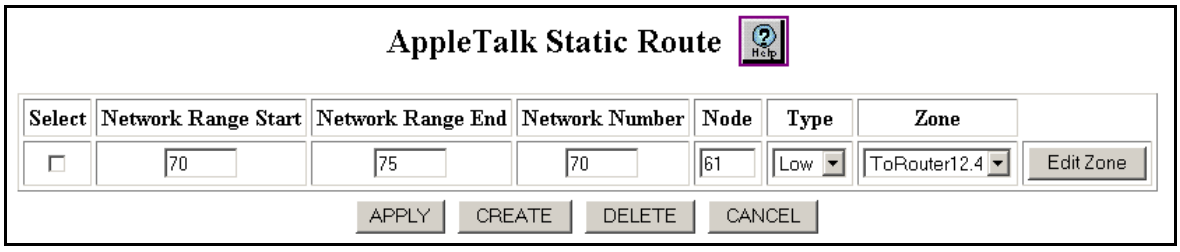
3. See Table 19-2 to configure the Add AppleTalk Static Route Web page parameters.

Table 19-2. Add AppleTalk Static Route Web Page Parameters

Parameter	Definition
Network Range Start	Enter the starting network number. The network number specifies the range of AppleTalk network numbers for extended networks. Each number in the range must be an integer between 1 and 65279. **Note:** Network Range Start must be less than or equal to Network Range End.
Network Range End	Enter the ending network number. The network number specifies the range of AppleTalk network numbers for extended networks. Each number in the range must be an integer between 1 and 65279.
Network Number	Enter the next hop network number. This number specifies the network number of the next hop router for the network range.
Node	Enter the next hop node identification number. This number must be between 1 and 253.
Type	Select the type of static route. The options include: • **High** - The static route is not superseded by a route update. • **Low** - The static route can be superseded by a route update.
Zone	Enter an AppleTalk zone name assigned to this route. Zone names may be up to 32 characters in length.

4. Click **APPLY** to add the static route, or **CANCEL** to restore previous settings. When you click **APPLY**, you are returned to the AppleTalk Static Route Web page. See Figure 19-8.

Figure 19-8. AppleTalk Static Route Web Page.



CLI Command

To create an AppleTalk static route using the CLI, enter the following command in Interface mode:

`(config-if:<interface-name>)#` **appletalk static cable-range** <***cable-range***> **to** <***network.node***> **[floating] zone** <***zone-name***>

Editing AppleTalk Static Routes

You can edit AppleTalk Static routes using either the Web Agent or the CLI.

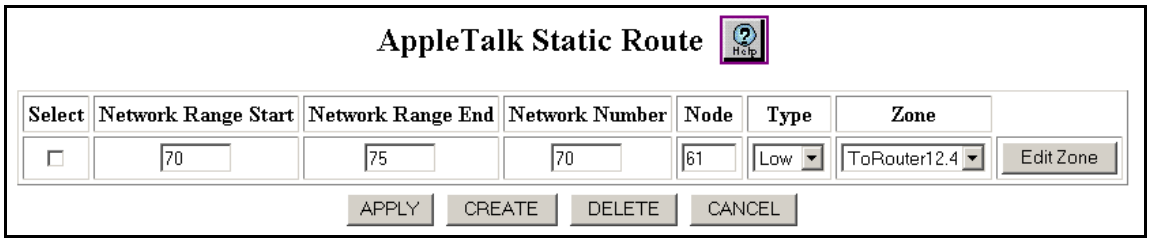
Web Agent Procedure

To edit an AppleTalk static route using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Static Route**.

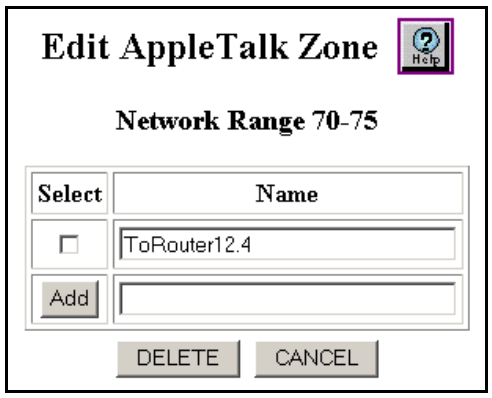
   The AppleTalk Static Route Web page is displayed in the content pane. See Figure 19-6.

Figure 19-9. AppleTalk Static Routes



2. If you need to Add or Delete Zones, select **Edit Zone** otherwise, skip to step 4. The Edit AppleTalk Zone Web page is displayed. See Figure 19-10

Figure 19-10. Edit AppleTalk Zone



3. If you need to add a new zone, enter the new AppleTalk zone name and click **Add** in the **Add text** field. If you need to delete a zone(s), place a check in the **Select** column next to the zone you need to delete then click the **DELETE** button.

4. Navigate back the AppleTalk Static Route Web page by expanding the **Routing** > **AppleTalk** > **Configuration** folders, and then clicking **Static Route**.

   The AppleTalk Static Route Web page is displayed in the content pane. See Figure 19-6.Select the AppleTalk Static Route that you want to edit from the **Select** column.

5. Enter the new network range start number in the **Network Range Start** field.

6. Enter the new network range end number in the **Network Range End** field.

7. Enter the new network number of the next hop for the static route in the **Network Number** field.

8. Enter the new node number of the next hop for the static route in the **Node** field.

9. Enter the type to be associated with the static route in the **Type** field.

10. Click **APPLY** to save the new zone.

CLI Command

To edit an AppleTalk static route using the CLI, enter the following command from Interface mode:

`(config-if:<interface-name>)#` **appletalk static cable-range** <***options***>

Deleting an AppleTalk Static Route

You can delete an AppleTalk static route using either the Web Agent or the CLI.

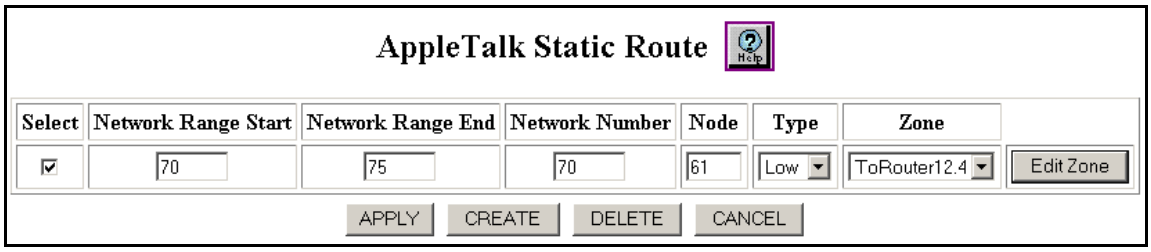
Web Agent Procedure

To delete an AppleTalk static route using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Static Route**.

   The AppleTalk Interfaces Web page is displayed in the content pane. See Figure 19-11.

Figure 19-11. AppleTalk Static Route



2. Select the AppleTalk static route that you want to delete from the **Select** column.

3. Click **DELETE** to remove the AppleTalk static route.

CLI Command

To delete an AppleTalk static route using the CLI, enter the following command in **Configure** mode:

`(config-if:<interface-name>)#` **no appletalk static**

Creating an AppleTalk Name-Binding Protocol (NBP) Filter

The NBP performs a conversion from named AppleTalk entities to their AppleTalk protocol addresses. Multiple names can exist for the same entity (alias). NBP also performs:

- Name registration
- Name deletion
- Name lookup
- Name confirmation

NBP allows you to bind a name to the internal storage address for your entity and register this mapping so that other entities can look it up. You can display NBP names to users and use addresses internally to locate entities. When you register your entity's name and address pair, NBP validates its uniqueness.

An NBP Filter prevents hosts on one or more interfaces from accessing hosts or services beyond that interface.

You can create an NBP Filter using either the Web Agent or the CLI.

Web Agent Procedure

To create an NBP Filter using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **NBP Filter**.

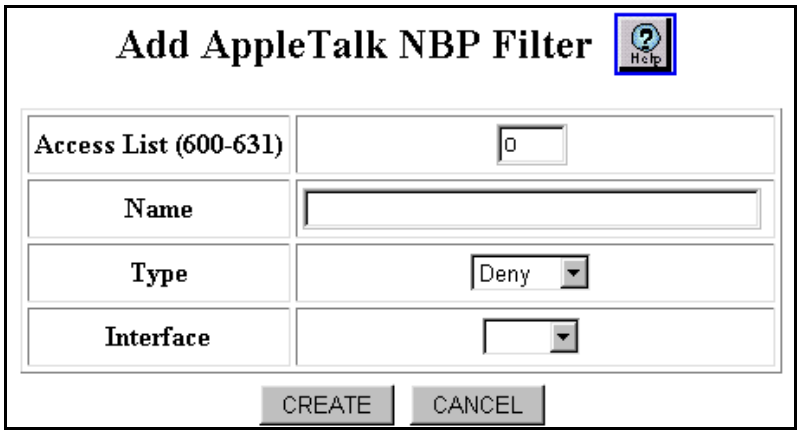
   The AppleTalk NBP Filter Web page is displayed in the content pane. See Figure 19-12.

Figure 19-12. AppleTalk NBP Filter Web Page



2. Select **CREATE** to add a new NBP filter. The Add AppleTalk NBP Filter Web page is displayed. See Figure 19-13.

Figure 19-13. Add AppleTalk NBP Filter Web Page



3. See Table 19-3 to configure the Add AppleTalk NBP filter Web page.

4. Click **CREATE** to add your new static route, or **CANCEL** to restore previous settings.

Table 19-3. AppleTalk NBP Filter Parameters

Parameter	Definition
Access List	Enter the access list number to be associated with the NBP filter. Valid values are 600-631.
Name	Enter the name of the NBP name object to be filtered.
	1 of 2

Table 19-3. AppleTalk NBP Filter Parameters

Parameter	Definition
Type	Select the type of filtering. The options include: • **Deny** • Permit
Interface	Select the interface(s) to apply this filter to.
	2 of 2

CLI Command

To create an NBP filter using the CLI, enter the following command from Interface mode:

`(config-if:<interface-name>)#` **appletalk access-list <*access-list-number*> {permit | deny} nbp <*string*>**

Editing an AppleTalk NBP Filter

You can edit an AppleTalk NBP filter using either the Web Agent or the CLI.

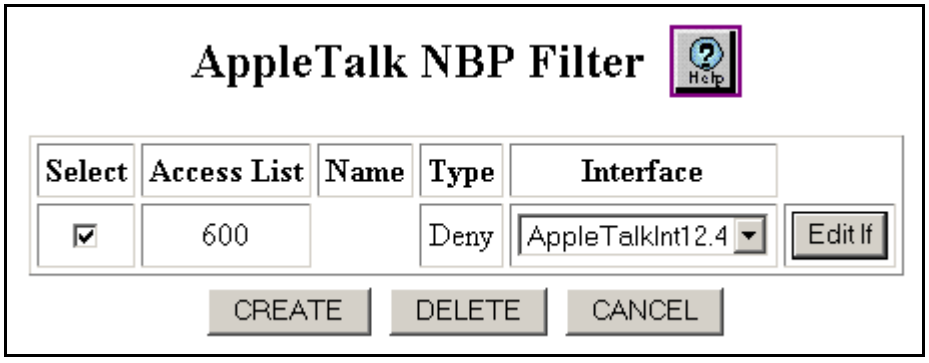
Web Agent Procedure

To edit an AppleTalk NBP filter using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **NBP Filter**.

   The AppleTalk NBP Filter Web page is displayed in the content pane. See Figure 19-14.

Figure 19-14. AppleTalk NBP Filter



2. See Table 19-3 to edit the AppleTalk NBP Filter Web page parameters.

3. Click:

   — **CREATE** to add a new filter. The Add AppleTalk NBP Filter Web page is displayed. See "Creating an AppleTalk Name-Binding Protocol (NBP) Filter" earlier in this chapter for more information.

   — **DELETE** to remove the selected NBP filter.

   — **CANCEL** to restore previous settings.

   — **Edit If** to add or delete this NBP filter to/from an interface. The Add/Delete Interface to NBP Filter Web page is displayed. See "Adding or Deleting Interfaces to an NBP Filter" later in this chapter for more information.

CLI Command

To edit an AppleTalk NBP filter using the CLI, enter the following command in Interface mode:

`(config-if:<interface-name>)#` **appletalk access-list** <***access-list-number***> **{permit | deny} nbp** <***string***>

Adding or Deleting Interfaces to an NBP Filter

You can add or delete interfaces to an NBP filter using either the Web Agent or the CLI.

Web Agent Procedure

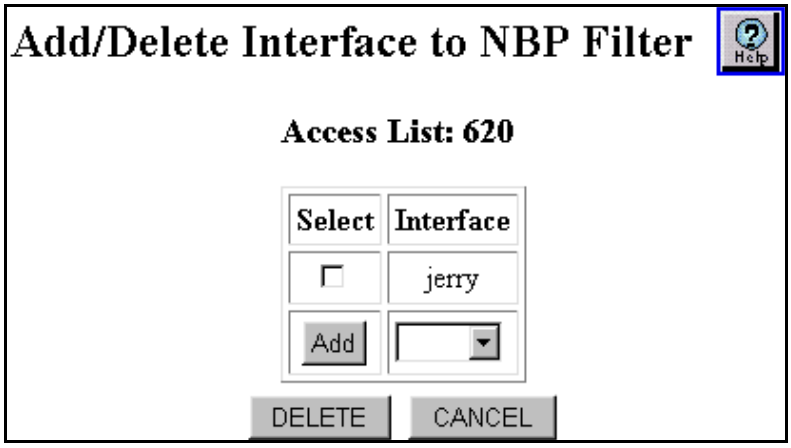
To add or delete interfaces to an NBP filter using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **NBP Filter**.

   The AppleTalk NBP Filter Web page is displayed in the content pane. See Figure 19-14.

2. Select **Edit If**. The Add/Delete Interface to NBP Filter Web page is displayed. See Figure 19-15.

Figure 19-15. Add/Delete Interface to NBP Filter Web Page



3. Select the interface to be added from the **Add** field pull-down menu.

   Or

   Click an interface from the **Select** column for the interface that you want to be remove.

4. Click:

   — **Add** to add this NBP filter to the selected interface.

   — **DELETE** to remove this NBP filter from the selected interface.

   — **CANCEL** to restore previous settings.

CLI Commands

To add interfaces to an NBP filter using the CLI, use the following command:

`(configure)#` **appletalk access-list <*access-list-number*> {permit | deny} nbp <*string*>**

To delete interfaces to an NBP filter using the CLI, enter the following command from **Configure** mode:

`(configure)#` **no appletalk access-list <*access-list-number*> {permit | deny} nbp <*string*>**

Creating an AppleTalk Zone Filter

You can create an AppleTalk Zone Filter using either the Web Agent or the CLI.

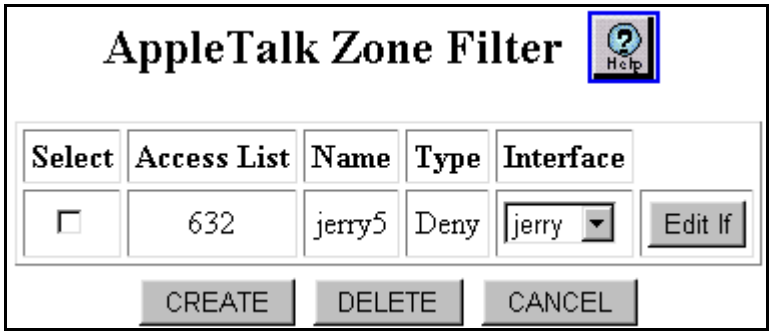
Web Agent Procedure

To create an AppleTalk zone filter using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Zone Filter**.

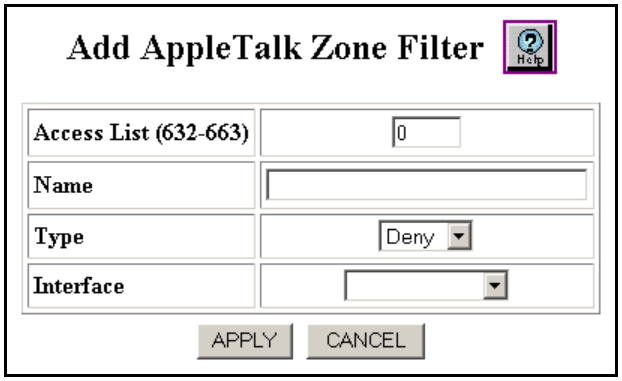
   The AppleTalk Zone Filter Web page is displayed in the content pane. See Figure 19-16.

Figure 19-16. AppleTalk Zone Filter Web Page



2. Select **CREATE**. The Add AppleTalk Zone Filter Web page is displayed. See Figure 19-17.

Figure 19-17. Add AppleTalk Zone Filter Web Page



3. See Table 19-4 to configure the Add AppleTalk Zone Filter Web page parameters.

4. Click **APPLY** to add the new zone filter, or **CANCEL** to restore previous settings.

Table 19-4. Add AppleTalk Zone Filter Web Page Parameters

Parameter	Definition
Access List	Enter the access list number to be associated with the zone filter. Valid values are 632-663.
Name	Enter the name of the zone to be filtered. You can also leave this field blank for wildcarding.
Type	Select the type of filtering. The options include: • **Deny** • **Permit**
Interface	Select the interface to apply this filter to.

CLI Command

To create an AppleTalk zone filter using the CLI, use the following command:

`(configure)#` **appletalk access-list <*access-list-number*> {permit | deny} zone <*string*>**

Editing an AppleTalk Zone Filter

You can edit an AppleTalk zone filter using either the Web Agent or the CLI.

Web Agent Procedure

To edit an AppleTalk zone filter using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Zone Filter**.

   The AppleTalk Zone Filter Web page is displayed in the content pane. See Figure 19-17.

2. Click the checkbox in the **Select** column next to the AppleTalk zone filter that you want to be edit.

3. See Table 19-4 to edit the AppleTalk Zone Filter Web page parameters:.

4. Click:

   — **CREATE** to add a new zone filter. The Add AppleTalk Zone Filter Web page is displayed. See "Creating an AppleTalk Zone Filter" earlier in this chapter for more information.

   — **DELETE** to remove the selected zone filter.

— **CANCEL** to restore previous settings.

— **Edit If** to add or delete this zone filter to/from an interface. The Add/Delete Interface to Zone Filter Web page is displayed. See "Adding or Deleting Interfaces to a Zone Filter" later in this chapter for more information.

CLI Command

To edit an AppleTalk zone filter using the CLI, use the following command:

`(configure)#` **appletalk access-list <*access-list-number*> {permit | deny} zone <*string*>**

Adding or Deleting Interfaces to a Zone Filter

You can add or delete interfaces to a zone filter using either the Web Agent or the CLI.

Web Agent Procedure

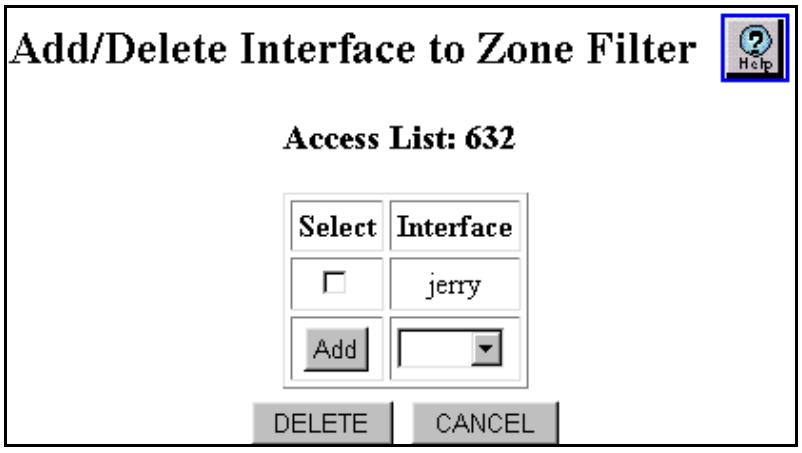
To add or delete interfaces to a zone filter using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Configuration** folders, and then click **Zone Filter**.

   The AppleTalk Zone Filter Web page is displayed in the content pane. See Figure 19-16.

2. Select **Edit If**. The Add/Delete Interface to Zone Filter Web page is displayed. See Figure 19-18.

Figure 19-18. Add/Delete Interface to Zone Filter Web Page



3. Select the interface to be added from the **Add** field pull-down menu.

   Or

   Click the checkbox in the **Select** column for the interface that you want to be remove.

4. Click:

   — **Add** to add this zone filter to the selected interface.

   — **DELETE** to remove this zone filter from the selected interface.

   — **CANCEL** to restore previous settings.

CLI Commands

To add interfaces to a zone filter using the CLI, enter the following command from Interface mode:

`(config-if:<interface-name>)#` **appletalk access-list** <***access-list-number***> **{permit | deny} zone** <***string***>

To delete interfaces to an zone filter, enter the following command from Interface mode:

`(config-if:<interface-name>)#` **no appletalk access-list** <***access-list-number***> **{permit | deny} zone** <***string***>

Viewing AppleTalk Statistics

This section includes:

- Viewing AppleTalk Global Statistics
- Viewing the AppleTalk Interface Statistics Table
- Viewing the AppleTalk Route Table
- Viewing AppleTalk Route Table Statistics
- Viewing the AppleTalk ARP Cache Table
- Viewing the AppleTalk Zone Table
- Viewing AppleTalk Zone Table Statistics
- Viewing the AppleTalk NBP Table

Viewing AppleTalk Global Statistics

You can view AppleTalk global statistics using either the Web Agent or the CLI.

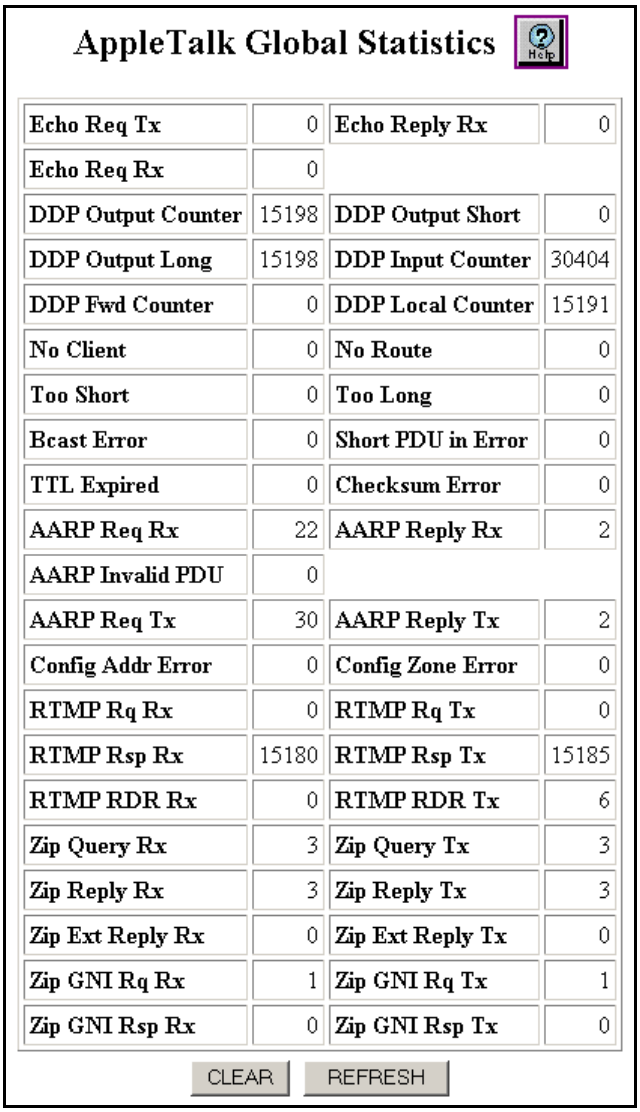
Web Agent Procedure

To view AppleTalk global statistics using the Web Agent:

1. In the navigation pane, expand the **Routing > AppleTalk > Display** folders, and then click **Global Statistics**.

   The AppleTalk Global Statistics Web page is displayed in the content pane. See Figure 19-19.

Figure 19-19. AppleTalk Global Statistics Web Page



AppleTalk Global Statistics

Echo Req Tx	0	Echo Reply Rx	0
Echo Req Rx	0		
DDP Output Counter	15198	DDP Output Short	0
DDP Output Long	15198	DDP Input Counter	30404
DDP Fwd Counter	0	DDP Local Counter	15191
No Client	0	No Route	0
Too Short	0	Too Long	0
Bcast Error	0	Short PDU in Error	0
TTL Expired	0	Checksum Error	0
AARP Req Rx	22	AARP Reply Rx	2
AARP Invalid PDU	0		
AARP Req Tx	30	AARP Reply Tx	2
Config Addr Error	0	Config Zone Error	0
RTMP Rq Rx	0	RTMP Rq Tx	0
RTMP Rsp Rx	15180	RTMP Rsp Tx	15185
RTMP RDR Rx	0	RTMP RDR Tx	6
Zip Query Rx	3	Zip Query Tx	3
Zip Reply Rx	3	Zip Reply Tx	3
Zip Ext Reply Rx	0	Zip Ext Reply Tx	0
Zip GNI Rq Rx	1	Zip GNI Rq Tx	1
Zip GNI Rsp Rx	0	Zip GNI Rsp Tx	0

CLEAR REFRESH

2. See Table 19-5 for information about the AppleTalk Global Statistics Web page parameters.

3. Click **REFRESH** to update all statistics, or **CLEAR** to reset all statistics to zero.

Table 19-5. AppleTalk Global Statistical Web Page Parameters

Parameter	Number of...
Echo Req Tx	Echo requests transmitted.
Echo Reply Rx	Echo replies received.
Echo Req Rx	Echo requests received.
DDP Output Counter	DDP packets sent from this node.
DDP Output Short	DDP packets sent using the short format.
DDP Output Long	DDP packets sent using the long format.
DDP Input Counter	DDP packets received at this node.
DDP Fwd Counter	DDP packets forwarded through this node.
DDP Local Counter	DDP packets received destined for this node.
No Client	Packets received for which the destination socket was not known.
No Route	Packets received for which the destination route was not known.
Too Short	Packets received that were smaller than the minimum size allowed for an AppleTalk packet.
Too Long	Packets received that were larger than the minimum size allowed for an AppleTalk packet.
Bcast Error	Broadcast errors detected.
Short PDU in Error	Packets received that had a short PDU in error.
TTL Expired	Packets dropped because they timed out.
Checksum Error	Packets which had checksum in error.
AARP Req Rx	AppleTalk ARP requests received.
AARP Reply Rx	Appletalk ARP replies received.
AARP Invalid PDU	AppleTalk ARP requests received which were invalid.
AARP Req Tx	AppleTalk ARP requests transmitted.
AARP Reply Tx	AppleTalk ARP replies transmitted.
	1 of 2

Table 19-5. AppleTalk Global Statistical Web Page Parameters

Parameter	Number of...
Config Addr Error	Configuration address errors detected.
Config Zone Error	Zone configuration errors detected.
RTMP Rq Rx	RTMP requests received.
RTMP Rq Tx	RTMP requests transmitted.
RTMP Rsp Rx	RTMP responses received.
RTMP Rsp Tx	RTMP responses transmitted.
RTMP RDR Rx	RTMP route data requests received.
RTMP RDR Tx	RTMP route data requests transmitted.
Zip Query Rx	ZIP queries received.
Zip Query Tx	ZIP queries transmitted.
Zip Reply Rx	ZIP replies received.
Zip Reply Tx	ZIP replies transmitted.
Zip Ext Reply Rx	ZIP extended replies received.
Zip Ext Reply Tx	ZIP extended replies transmitted.
Zip GNI Rq Rx	ZIP get net info request received.
Zip GNI Rq Tx	ZIP get net info request transmitted.
Zip GNI Rsp Rx	ZIP get net info response received.
Zip GNI Rsp Tx	ZIP get net info response transmitted.
	2 of 2

CLI Command

To view AppleTalk global statistics using the CLI, enter the following command from Privileged mode:

```
# show appletalk traffic
```

Viewing the AppleTalk Interface Statistics Table

You can view the AppleTalk interface statistics table using either the Web Agent or the CLI.

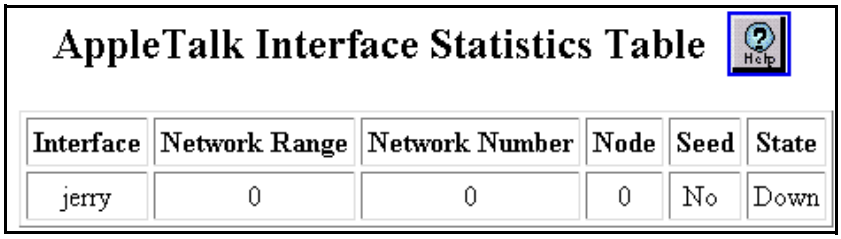
Web Agent Procedure

To view the AppleTalk Interface statistics table using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Display** folders, and then click **Interface Statistics**.

   The AppleTalk Interface Statistics Table is displayed in the content pane. See Figure 19-20.

Figure 19-20. AppleTalk Interface Statistics Table



AppleTalk Interface Statistics Table

Interface	Network Range	Network Number	Node	Seed	State
jerry	0	0	0	No	Down

2. See Table 19-6 for information on the **AppleTalk Interface Statistics Table** parameters:

Table 19-6. AppleTalk Interface Statistics Table Parameters

Parameter	Definition
Interface	Displays the name of the AppleTalk interface.
Network Range	Displays the network range associated with the AppleTalk interface.
Network Number	Displays the network number of this node.
Node	Displays the node number of this node.
Seed	Displays whether the AppleTalk interface seeded the network.
State	Displays the state of the AppleTalk interface. The options include: • **Up** - indicates that the AppleTalk interface is active. • **Down** - indicates that the AppleTalk interface is inactive.

CLI Command

To view the AppleTalk Interface statistics table using the CLI, enter the following command from Privileged mode:

show appletalk interface [brief] [<*interface-name*>]

Viewing the AppleTalk Route Table

You can view the AppleTalk route table using the Web Agent or the CLI. You can also delete single entries from the table, or flush all of the routes from the table.

Web Agent Procedure

To view the AppleTalk route table, and delete or flush entries from the table using the Web Agent:

1. In the navigation pane, expand the **Routing > AppleTalk > Display** folders, and then click **Route Table**.

   The AppleTalk Route Table is displayed in the content pane. See Figure 19-21.

Figure 19-21. AppleTalk Route Table Web Page



2. See Table 19-7 for information about the AppleTalk Route Table parameters.

3. Click **Delete Entries** to remove selected route table entries.

   Or

   Click **Flush Route Table** to empty the route table of all dynamic entries. Static and Local entries are not flushed.

Table 19-7. AppleTalk Route Table Parameters

Parameter	Definition
Select	Select the entry to be acted upon.
Network Range	Displays the network range.
Metric	Displays the AppleTalk metric for the network range.
State	Displays the state of the entry.The options include: • Good • Suspect • Going Bad • Bad
	1 of 2

Table 19-7. AppleTalk Route Table Parameters

Parameter	Definition
Owner	Displays the AppleTalk component responsible for the addition of the route.The options include: • Local • Static • RTMP
Next Hop	Displays the next hop address where forwarded packets are routed.
Interface	Displays the Appletalk interface associated with the route table entry.
Zones	Displays the zones associated with the selected AppleTalk route.
	2 of 2

CLI Command

To view the AppleTalk Interface statistics table using the CLI, enter the following command:

> show appletalk route

Viewing AppleTalk Route Table Statistics

You can view AppleTalk route table statistics using either the Web Agent or the CLI.

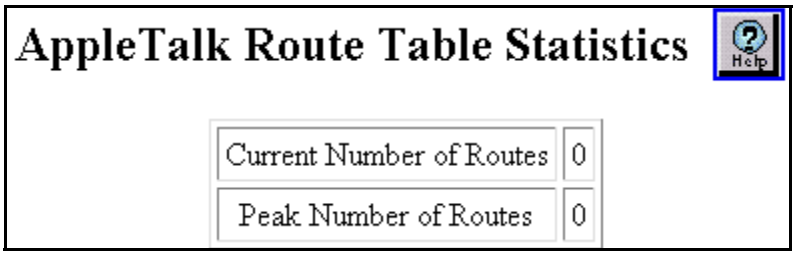
Web Agent Procedure

To view AppleTalk route table statistics using the Web Agent:

1. In the navigation pane, expand the **Routing > AppleTalk > Display** folders, and then click **Route Table Statistics**.

   The AppleTalk Route Table Statistics Web page is displayed in the content pane. See Figure 19-22.

Figure 19-22. AppleTalk Route Table Statistics Web Page



2. See Table 19-8 for information on AppleTalk Route Table Statistics Web page parameters:

Table 19-8. AppleTalk Route Table Statistics Web Page

Parameter	Definition
Current Number of Routes	Displays the current number of AppleTalk routes.
Peak Number of Routes	Displays the peak number of AppleTalk routes.

CLI Command

Currently there is no CLI command to show appleTalk Route Statistics.

Viewing the AppleTalk ARP Cache Table

You can view the AppleTalk ARP cache table using either the Web Agent or the CLI.

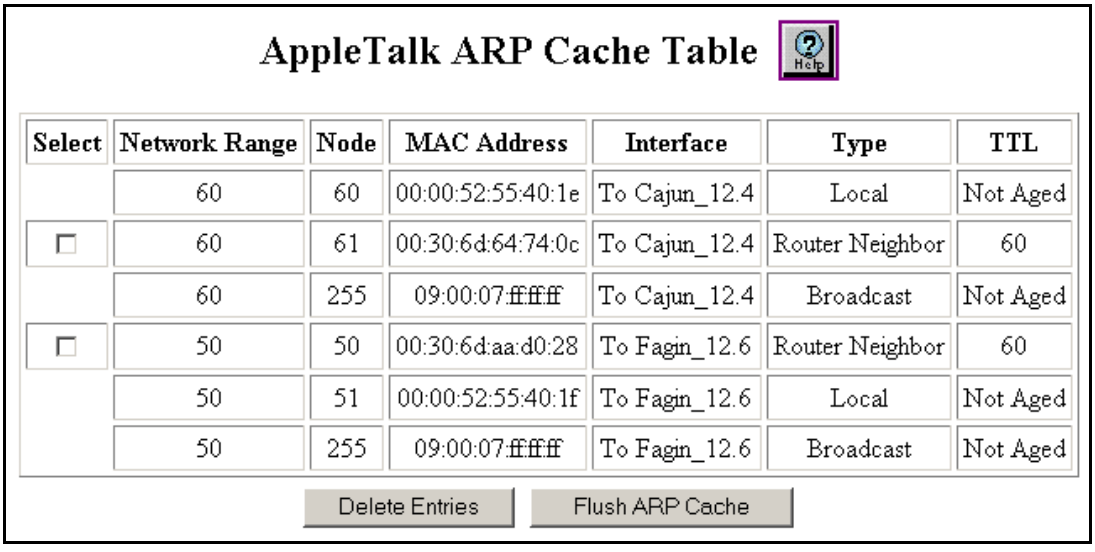
Web Agent Procedure

To view the AppleTalk ARP cache table using the Web Agent:

1. In the navigation pane, expand the **Routing > AppleTalk > Display** folders, and then click **ARP Table**.

   The AppleTalk ARP Cache Table is displayed in the content pane. See Figure 19-23.

Figure 19-23. AppleTalk ARP Cache Table



2. See Table 19-9 for a definition of the **AppleTalk ARP Cache Table** parameters.

3. Select the ARP entries that you want to delete and click **Delete Entries**.

   or

   Click **Flush Route Table** to reset the ARP cache table entries. Only Dynamic and Router Neighbor entries are flushed.

Table 19-9. AppleTalk ARP Cache Table Parameters

Parameter	Definition
Select	Select the table entry to be acted upon.
Network Range	Displays the network range.
Node	Displays the node number for the entry.
MAC Address	Displays the MAC address associated with the AppleTalk ARP cache table entry of the node.
Interface	Displays the AppleTalk interface associated with the AppleTalk ARP cache table entry.
Type	Displays the type of cache entries.The values are: • Local • Broadcast • Dynamic • Router Neighbor
TTL	Displays the time-to-live value for the selected AppleTalk ARP cache table entry. Local and Broadcast entries do not time out.

CLI Command

To view the AppleTalk ARP cache table using the CLI, enter the following command from **User** mode:

> show appletalk arp [all]

Viewing the AppleTalk Zone Table

You can view the AppleTalk Zone table using either the Web Agent or the CLI.

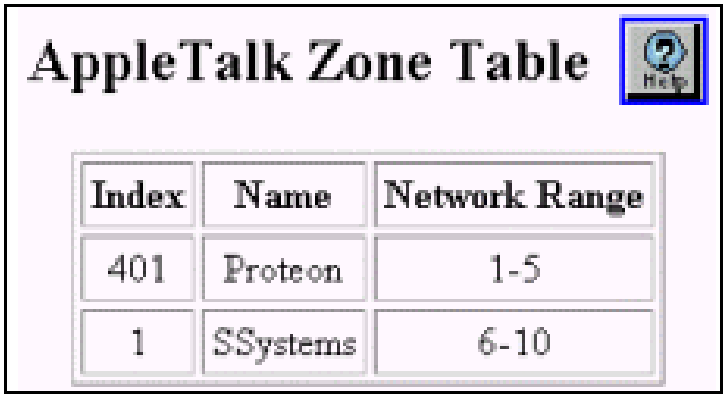
Web Agent Procedure

To view the AppleTalk zone table using the Web Agent:

1. In the navigation pane, expand the **Routing > AppleTalk > Display** folders, and then click **Zone Table**.

   The AppleTalk Zone Table is displayed in the content pane. See Figure 19-24.

Figure 19-24. AppleTalk Zone Table



2. See Table 19-10 for an explanation of the **AppleTalk Zone Table** parameters.

Table 19-10. AppleTalk Zone Table Parameters

Parameter	Definition
Index	Displays the zone index.
Name	Displays the zone name.
Network Range	Displays the network range associated with the zone.

CLI Command

The command to show AppleTalk Zone Statistics currently does not exist.

Viewing AppleTalk Zone Table Statistics

You can view the AppleTalk Zone Table Statistics using either the Web Agent or the CLI.

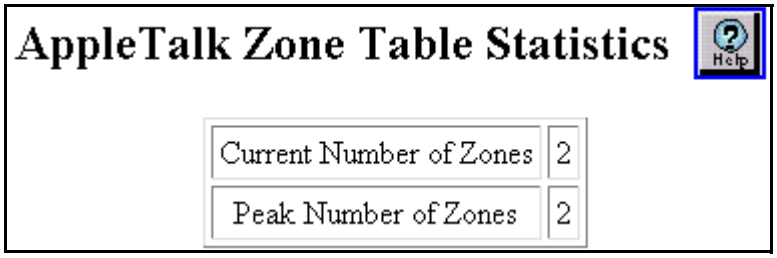
Web Agent Procedure

To view AppleTalk zone table statistics using the Web Agent:

1. In the navigation pane, expand the **Routing** > **AppleTalk** > **Display** folders, and then click **Zone Table Statistics**.

   The AppleTalk Zone Table Statistics Web page is displayed in the content pane. See Figure 19-25.

Figure 19-25. AppleTalk Zone Table Statistics Table



2. See Table 19-11 for an explanation of the **AppleTalk Zone Table Statistics** parameters.

Table 19-11. AppleTalk Zone Table Statistical Parameters

Parameter	Definition
Current Number of Zones	Displays the current number of AppleTalk zones.
Peak Number of Zones	Displays the peak number of AppleTalk zones.

CLI Command

To view AppleTalk zone table statistics using the CLI, enter the following command from **User** mode:

> show appletalk zone [<*zone-name*>]

Viewing the AppleTalk NBP Table

You can view the AppleTalk NBP Table using either the Web Agent or the CLI.

Web Agent Procedure

To view the AppleTalk NBP table using the Web Agent:

1. In the navigation pane, expand the **Routing > AppleTalk > Display** folders, and then click **Zone Table**.

   The AppleTalk Zone Table Web page is displayed in the content pane. See Figure 19-24.

2. See Table 19-12 for an explanation of the AppleTalk NBP Table parameters:

Table 19-12. AppleTalk NBP Table Statistical Parameters

Parameter	Definition
Index	Displays the index of the name binding protocol entry.
Name	Displays the name of the NBP entry.
Type	Displays the type of object named.
Interface	Displays the AppleTalk interface associated with the AppleTalk NBP table entry.
Zone	Displays the zone field associated with the NBP table entry.

CLI Command

To view the AppleTalk NBP table using the CLI, enter the following command from **User** mode:

> show appletalk nbp

20 Managing Intelligent Multicasting

Overview

The following information and procedures provided in this chapter pertain to layer 2 and layer 3 module configurations:

- Introduction
- Configuring Intelligent Multicasting
- Managing IGMP Snooping
- Managing the LGMP Server
- Managing the LGMP Client
- Managing CGMP Snooping

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Introduction

Intelligent multicasting refers to the forwarding of multicast traffic (packets with a multicast destination MAC address) within a VLAN to a subset of ports participating in that VLAN. It limits the forwarding of multicast traffic to only those ports on the VLAN with clients that want to receive this multicast traffic.

When disabled, intelligent multicasting floods layer 2 multicast traffic to all ports on the VLAN on which the traffic is received.

All traffic that is sent to a particular multicast MAC address is said to be in a multicast session. The switch supports 58 sessions per VLAN. Each multicast session keeps track of which ports must receive that session's multicast traffic within the VLAN. There are two types of ports: client ports and router ports.

A client port is a port with an attached host configured to receive a multicast session.

Router ports are ports that are attached to (or in the path to) multicast routers and must be treated specially. All multicast traffic on a VLAN must be forwarded to the router port.

Configuration of an Intelligent Multicast session first requires a session to be established. Once that session is established, client and router ports can be added to or removed from the session. Session and port configuration can be done either manually or dynamically. Dynamic intelligent multicasting is achieved through Internet Group Management Protocol (IGMP) Snooping, and may also involve Lucent Group Membership Protocol (LGMP), or Cisco Group Management Protocol (CGMP) Snooping. All of the dynamic mechanisms are based on the assumption that the client host is running IGMP, and is requesting membership in the IP multicast session.

*** Note:** If there is no multicast session created for a multicast flow in a VLAN, then that multicast flow will be flooded to all ports on the VLAN. This is the default behavior for a bridge as described in IEEE 802.1D. Intelligent multicasting must be enabled for any dynamic intelligent multicasting to be active.

By default, the switch rate limits inter-router multicast traffic on all modules that support rate limiting. If you need for your switch to support heavy multicast traffic, disable rate limiting on ports that are connected to routers. Also note that if high-bandwidth multicast streams are being used, rate limiting will affect directly connected clients if that stream is not part of an Intelligent Multicast session.

Manually Configured Intelligent Multicasting

Manual configuration of Intelligent Multicast sessions allows the network administrator to dictate which multicast streams will be intelligently multicasted. This method of configuration is also useful where dynamic Intelligent Multicasting can not be used. Dynamic Intelligent Multicasting can only be used in an environment that uses IGMP and an IP multicast routing protocol (optional) to distribute multicast streams.

Dynamic Intelligent Multicasting

In an environment that uses IGMP (and an IP multicast routing protocol, optionally), Dynamic Intelligent Multicasting may be used. The ability of the switch to dynamically set up sessions and add or remove client and router ports is particularly useful in a flexible multicast environment where there are many multicast sessions to administer. The protocols and mechanisms used to perform Intelligent Multicasting are enabled or disabled across all VLANs on a switch. The Dynamic Intelligent Multicasting process is split into three parts: Learning, Administration, and Dissemination.

The Learning Process, IGMP Snooping

To learn which sessions must be configured, or which client or router ports must be added or removed, a mechanism to draw that information from the layer 3 protocols must be enabled. For the Avaya Multiservice Switch, IGMP Snooping is that mechanism. IGMP Snooping can learn about multicast sessions from IP interfaces that reside on the same switch.

*** Note:** IGMP Snooping is only available for VLANs that have an IGMP interface enabled.

Enabling DVMRP on an interface also enables IGMP.

The Administration Process; Management and Pruning

At the core of all Intelligent Multicast functionality, the Administration Process is manipulated by manual configuration and dynamic configuration, and implements a pruning function. In a dynamically configured environment, the Administration Process takes the information from the Learning Process (from IGMP Snooping) or from one of the clients in the Dissemination Process and creates the AFT table entries which will perform the actual Intelligent Multicast functionality on the local switch. It also passes the new information to the servers in the Dissemination Process for distribution to other switches in the same VLAN.

The Administration Process views all data inputs as **applications**. Manual configuration is the MGMT (management) application, IGMP Snooping is the IGMP application, the LGMP client is the LGMP application, and the CGMP snooper is the CGMP application.

The pruning functionality of the Administration Process is a cleanup facility that prevents stale information from existing in the Intelligent Multicast tables in the event an application fails to perform its own cleanup. There are three types of pruning, each with their own timers: session, client port, and router port.

Automatic Session Pruning, if enabled, will remove a session if that session has not been active for Session Pruning Time.

Automatic Client Port Pruning, if enabled, will remove a client port from a session if no IGMP reports have been received on that port for that session in the time specified.

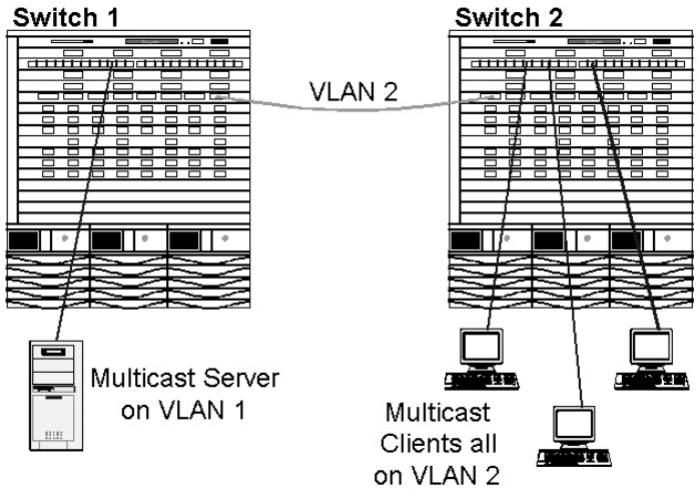
*** Note:** By enabling Automatic Client Port Pruning, there is a chance that the switch may disrupt multicast service temporarily to clients requesting to receive that multicast traffic. This is because of the report suppression mechanism in IGMP v1 and v2.

Automatic Router Port pruning, if enabled, will prune quiet router ports. If the switch has not been notified that there is a router on the port in Router Pruning Time, then the router port is removed.

The Dissemination Process, LGMP and CGMP Snooping

The Dissemination Process provides a method to dynamically configure multicast sessions on switches with VLANs that do not have IP interfaces. See Figure 20-1.

Figure 20-1. LGMP and CGMP Snooping



Switch 1 will route the multicast traffic from the Multicast Server on VLAN 1 to VLAN 2. Switch 2 does not have an IP interface on VLAN 2. Because the Learning Process requires an IP interface to perform learning, a different method must be used to create multicast sessions on Switch 2. Switch 1 must disseminate the Intelligent Multicast information to all switches on the attached VLANs that do not have IP interfaces. Two protocols are available for this use on the Switch, LGMP and CGMP.

LGMP and CGMP are similar protocols. Both protocols have a server implementation that runs on a switch that implements the Learning Process, and both protocols have a client implementation which runs on switches that do not have local IP interfaces on all involved VLANs. For the Avaya Multiservice switch, only the client implementation of CGMP is available in case the switch is connected to a Cisco router. This implementation is called CGMP Snooping.

*** Note:** IGMP Snooping must be enabled to learn the sessions that the LGMP server will disseminate. An LGMP server is active only for VLANs that have an IGMP interface enabled.

Configuring Intelligent Multicasting

This section provides the following procedures:

- Configuring Global Intelligent Multicasting
- Displaying Router Ports
- Configuring Static Router Ports
- Searching for Intelligent Multicast Sessions
- Deleting an Intelligent Multicast Session
- Deleting a Multicast Session Client Port
- Creating a Static Multicast Session
- Deleting Static Multicast Sessions
- Creating Static Client Ports
- Deleting Static Client Ports

Configuring Global Intelligent Multicasting

You can configure intelligent multicasting globally using either the Web Agent or the CLI.

Web Agent Procedure

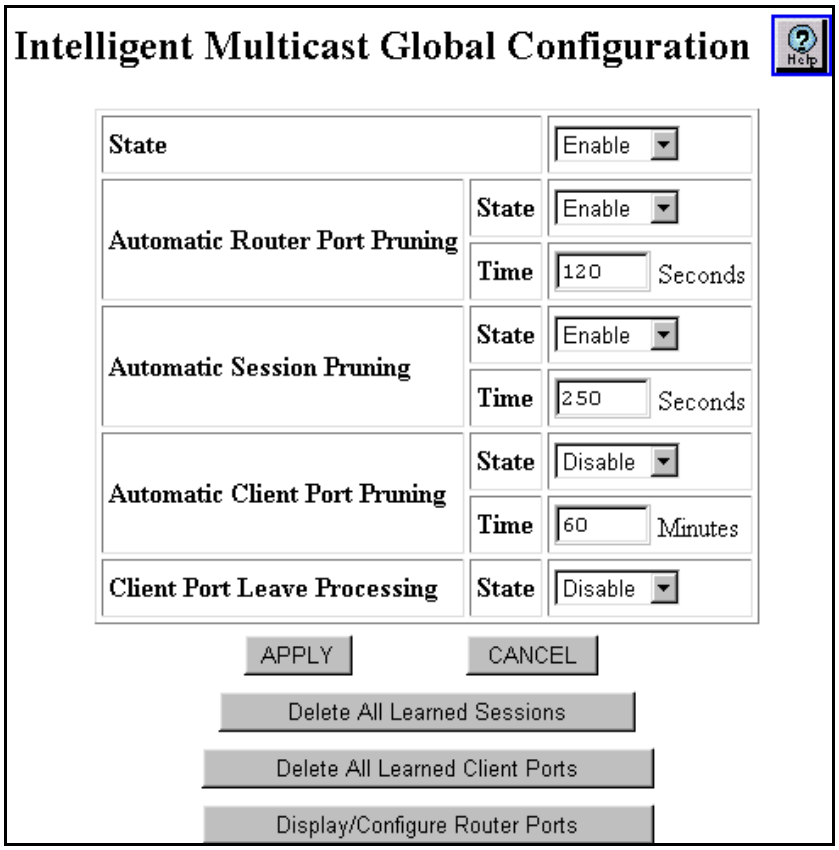
To configure intelligent multicasting globally using the Web Agent:

1. In the navigation pane, expand the **Layer 2 Switching** > **Intelligent Multicast** folders, and then click **Global Configuration**.

   The Intelligent Multicast Global Configuration Web page is displayed in the content pane. See Figure 20-2.

*** Note:** You cannot delete Static sessions and client ports from this Web page. See "Deleting a Multicast Session Client Port" and "Deleting Static Multicast Sessions" in this chapter to delete static sessions and client ports.

Figure 20-2. Intelligent Multicasting Global Configuration Web Page



***Note:** In order to route multicast traffic, IP multicast forwarding must be enabled on the switch. See "Enabling IP Routing Global Parameters" in Chapter 12, "Configuring IP Routing."

2. See Table 20-1 to configure the Intelligent Global Multicast Configuration Web page parameters.

3. Click

   — **APPLY** to save your changes.

   — **CANCEL** to restore previous settings.

   — **Delete All Learned Sessions** to remove all learned multicast sessions.

   — **Delete All Learned Client Ports** to remove all learned client ports from all multicast sessions.

   — **Display/Configure Router Ports** to display the router ports and configure your static router ports.

Table 20-1. Intelligent Global Multicast Configuration Web Page Parameters

Parameter	Definition
Enable State	Select **Disable** to prevent intelligent multicasting globally. The default value is **Enabled**.
Automatic Router Port Pruning	Select: • **Disable** - to prevent automatic router port pruning. The default value is **Enable**. • **Time** - Enter the time, in seconds, after which quiet learned router ports can be pruned. The valid range (in seconds) is 10 to 172800 (48 hours). The default value is **120 seconds**.
Automatic Session Pruning	Select: • **Disable** - To prevent automatic removal of stale sessions. The default value is **Enable**. • **Time** - Enter the time, in seconds, after which stale learned sessions can be removed. The valid range (in seconds) is 10 to 172800 (48 hours). The default value is **250 seconds**.
Automatic Client Port Pruning	Select: • **Disable** - to prevent the automatic removal of quiet learned client ports from a session. The default value is **Disable**. • **Time** - Enter the time, in minutes, after which quiet learned client ports can be automatically removed from a session. The valid range (in minutes) is 1 to 1440 (24 hours). The default value is **60 minutes**.

CLI Command

To configure intelligent multicasting globally using the CLI, enter the following command from **Configure** mode:

```
(configure)# set intelligent-multicast {enable}
```

Displaying Router Ports

You can display router ports using either the Web Agent or the CLI.

Web Agent Procedure

To display router ports using the Web Agent:

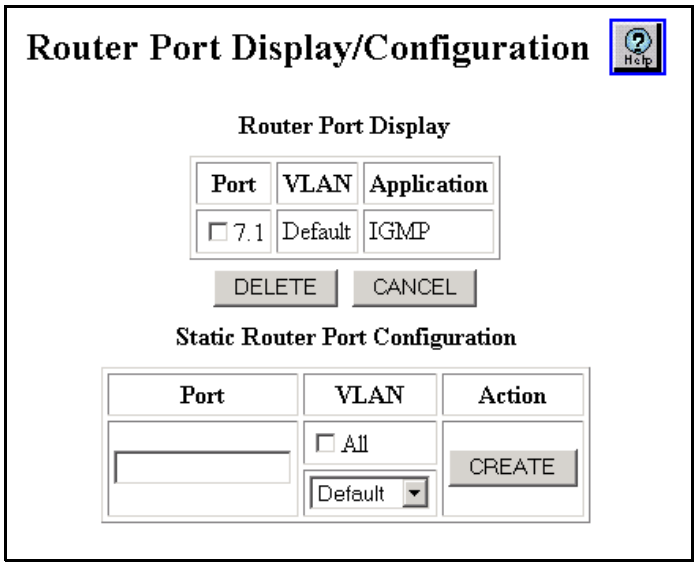
1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **Global Configuration**.

   The Intelligent Multicast Global Configuration Web page is displayed in the content pane. See Figure 20-2.

   ***Note:** You must enable intelligent multicasting on a global basis to make your router ports active. See "Configuring Global Intelligent Multicasting" earlier in this chapter.

2. Select **Display/Configure Router Ports**. The Router Port Display/Configuration Web page is displayed. See Figure 20-3.

Figure 20-3. Router Port Display/Configuration Web Page



3. See Table 20-2 for an explanation of the Router Port Display/Configuration Web page parameters.

4. Select a router port from the **Port** column and click **DELETE** to delete the port, or **CANCEL** to restore previous settings.

Table 20-2. Router Port Display/Configuration Web Page Parameters

Parameter	Definition
Port	Displays the switch port configured as a router port.
VLAN	Displays which VLAN the router port is bound to.
Application	Displays the active applications of the router port. The applications include: • Mgmt (Static) • IGMP • LGMP • CGMP

CLI Command

To display router ports using the CLI, enter the following command from User mode:

> show intelligent-multicast router-port

Configuring Static Router Ports

You can configure Static Router ports using either the Web Agent or the CLI.

Web Agent Procedure

To configure a static router port using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **Global Configuration**.

   The Intelligent Multicast Global Configuration Web page is displayed in the content pane. See Figure 20-2.

   ***Note:** You must enable intelligent multicasting on a global basis to make your static router port configuration active. See "Configuring Global Intelligent Multicasting" earlier in this chapter.

2. Select **Display/Configure Router Ports**. The Router Port Display/ Configuration Web page is displayed. See Figure 20-3.

3. Enter a port number in the **Port** field.

4. Select **All** from the **VLAN** column to add this router port to all VLANs

   Or

   Select a specific VLAN from the VLAN field pull-down menu.

   ***Note:** When adding a router port to all VLANs, the router port is added only to the VLANs bound to the switch port. To bind multiple VLANs to a switch port, see "Using VLANs, Hunt Groups, and VTP Snooping" in Chapter 6.

5. Click **CREATE** to save your changes.

CLI Command

To configure a static router port using the CLI, enter the following command from Configure mode:

`(configure)#` **set intelligent-multicast router-port vlan <*options*>**

Searching for Intelligent Multicast Sessions

You can search for intelligent multicast sessions using either the Web Agent or the CLI.

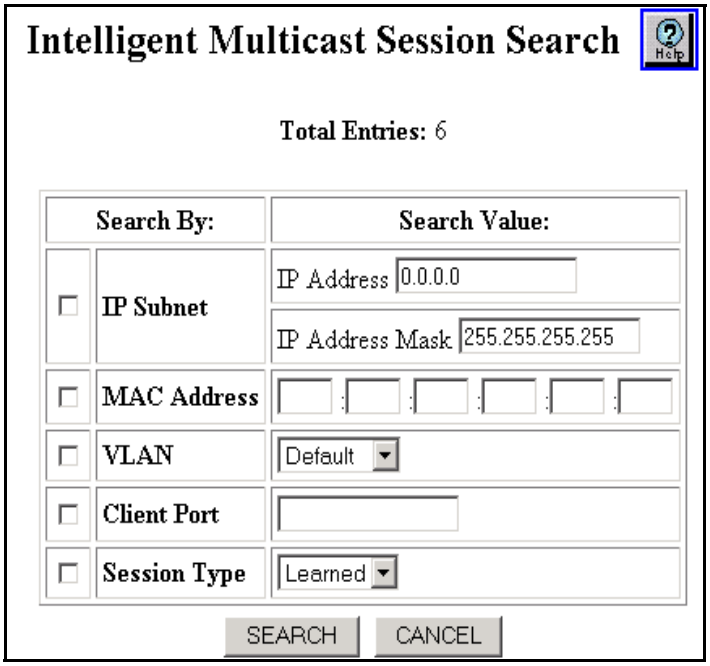
Web Agent Procedure

To perform a session search using the Web Agent:

1. In the navigation pane, expand the **L2 Switching** > **Intelligent Multicast** folders, and then click **Session Search**.

   The Intelligent Multicast Session Search Web page is displayed in the content pane. See Figure 20-4.

Figure 20-4. Intelligent Multicast Session Search Web Page



2. See Table 20-3 for an explanation of the Intelligent Multicast Session Search Web page parameters.

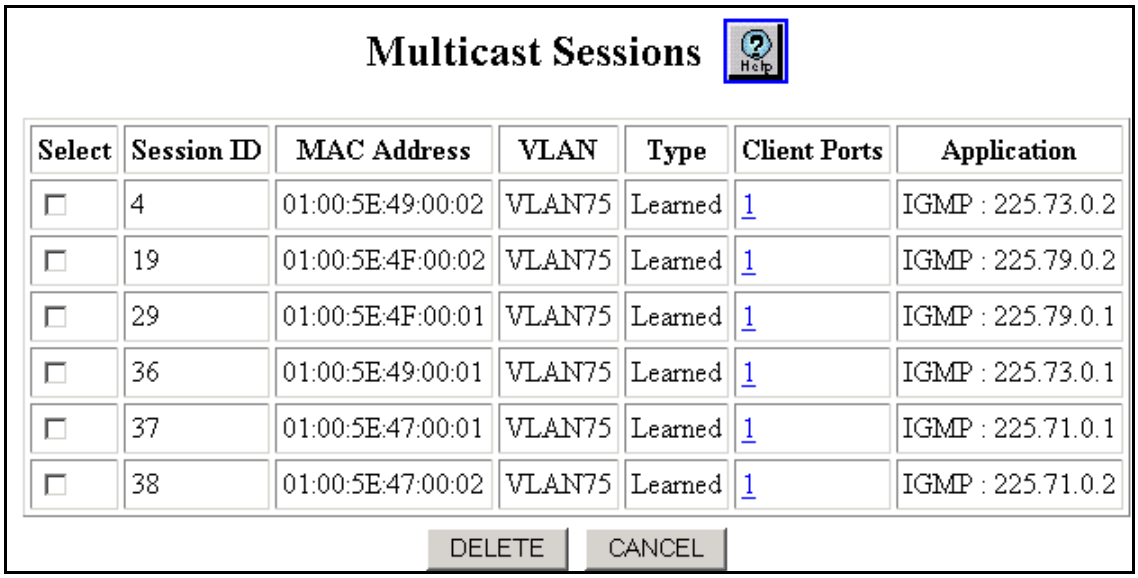
Table 20-3. Intelligent Multicast Session Search Web Page Parameters

Search By Parameter:	Search Value:
IP Subnet	• **IP Address** - enter the IP address on which you want to perform a search • **IP Address Mask** - Enter the associated IP Address mask
MAC Address	Enter the MAC address on which you want to perform a search
VLAN	Select a VLAN from the pull-down menu on which you want to perform a search.
Client Port	Enter a client port number on which you want to perform a search.
Session Type	Select a Session Type from the pull-down menu on which you want to perform a search. The options are Learned and Mgmt.

3. Select **SEARCH** to begin the search for the multicast session.

   The Multicast Sessions Web page is displayed.with the search results. See Figure 20-5.

Figure 20-5. Multicast Sessions Web Page



4. See Table 20-4 for an explanation of the Multicast Sessions Web page parameters:

Table 20-4. Multicast Sessions Web Page Parameters

Parameter	Definition
Select	Select the multicast session.
Session ID	Displays the multicast session identifier.
MAC Address	Displays the MAC address of the multicast session.
VLAN	Displays the VLAN on which the multicast session exists.
Type	Displays the type of multicast session. Options include: • **Learned** - Entry is dynamically learned. • **Mgmt** - Entry is configured statically by the user.
	1 of 2

Table 20-4. Multicast Sessions Web Page Parameters

Parameter	Definition
Client Ports	Displays the number of client ports in the multicast session and opens the Multicast Session Client Ports Web page.
Application	Displays the active application(s) that configured the multicast session. The applications include: • Mgmt • IGMP • LGMP • CGMP
	2 of 2

CLI Command

To perform a session search using the CLI, enter the following command from User mode:

> show intelligent-multicast session <*options*>

Deleting an Intelligent Multicast Session

You can delete an intelligent multicast session using either the Web Agent or the CLI.

Web Agent Procedure

To delete an intelligent multicast session using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **Session Search**.

   The Intelligent Multicast Session Search Web page is displayed in the content pane. See Figure 20-4.

2. Select **SEARCH** to begin the search for the multicast session. The Multicast Sessions Web page displays with the search results. See Figure 20-5.

3. Select the checkbox next to the multicast session that you want to delete and click **DELETE**.

   ***Note:** Static (Mgmt) sessions can only be deleted through the Static Sessions Web page.

CLI Command

To delete a multicast session using the CLI, enter the following command from Configure mode:

`(configure)#` **clear intelligent-multicast session <*session-id*>**

Deleting a Multicast Session Client Port

You can delete an intelligent multicast session client port by using either the Web Agent or the CLI.

Web Agent Procedure

To delete a multicast session client port by using the Web Agent:

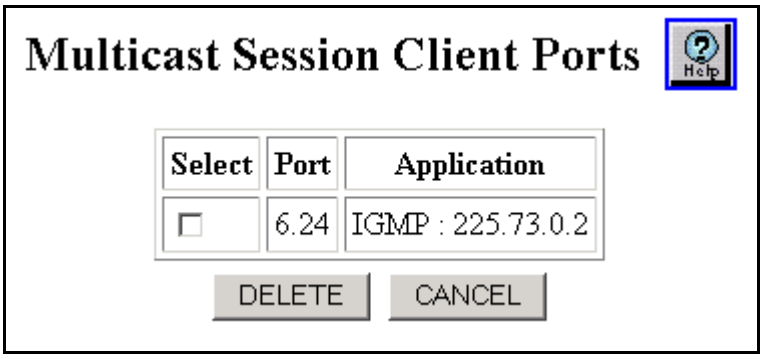
1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **Session Search**.

   The Intelligent Multicast Session Search Web page is displayed in the content pane. See Figure 20-4.

2. Select **SEARCH** to begin the search for the multicast session. The Multicast sessions Web page is displayed with the search results. See Figure 20-5.

3. Select the client port number from the **Client Ports** column. The Multicast Session Client Port Web page is displayed. See Figure 20-6.

Figure 20-6. Multicast Sessions Client Port



4. See Table 20-5 for an explanation of the Multicast Session Client Port Web page parameters.

5. Click **DELETE** to remove your selected multicast session client port, or **CANCEL** to restore previous settings.

*** Note:** Static (Mgmt) client ports can only be removed through the Static Sessions Web page.

Table 20-5. Multicast Session Client Port Web Page Parameters

Parameter	Definition
Select	Select the multicast session client port.
Port	Displays the client port number.
Application	Displays the application(s) associated with this client port. The applications are: • Mgmt • IGMP • LGMP • CGMP

CLI Command

To delete a multicast session client port using the CLI, enter the following command from Configure mode:

`(configure)#` **clear intelligent-multicast client-port <*session-id*> port <*mod-port-spec*>**

Creating a Static Multicast Session

You can create a static multicast session using either the Web Agent or the CLI.

Web Agent Procedure

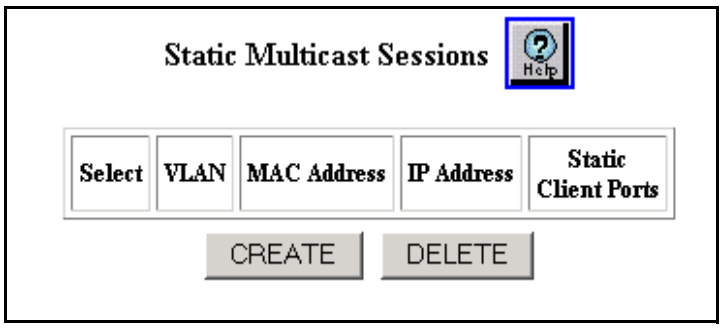
To create a new static multicast session using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **Static Sessions**.

   The Static Multicast Sessions Web page is displayed in the content pane. See Figure 20-7.

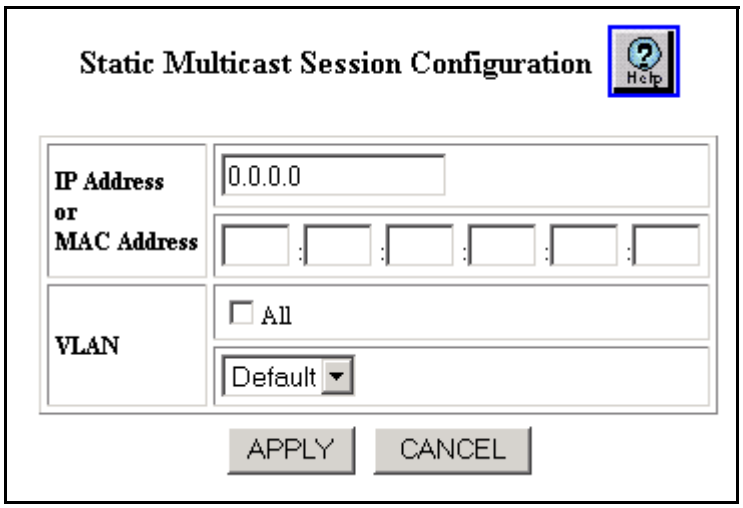
   ***Note:** You must enable intelligent multicasting on a global basis to make your static multicast session active. See "Configuring Global Intelligent Multicasting" earlier in this chapter.

Figure 20-7. Static Multicast Session Web Page



2. Select **Create** to create a new session. The Static Multicast Session Configuration Web page is displayed. See Figure 20-8.

Figure 20-8. Static Multicast Session Configuration Web Page



3. Configure the Static Multicast Session Configuration Web page parameters. See Table 20-6.

4. Click **APPLY** to create the new static multicast session, or **CANCEL** to restore previous settings.

Table 20-6. Static Multicast Session Configuration Web Page Parameters

Parameter	Definition
IP Address	Enter the IP address of the new static multicast session. The range must be between 224.0.1.0 to 239.255.255.255.
	1 of 2

Table 20-6. Static Multicast Session Configuration Web Page Parameters

Parameter	Definition
MAC Address	Enter the MAC address of the new static multicast session. Specifying the MAC address is not used for an IP multicast session.
VLAN	Enter the VLAN associated with the new multicast session. Click All to add all VLANs to the multicast session, or select a specific VLAN from the pull-down menu.
	2 of 2

CLI Command

To create a new static multicast session using the CLI, enter the following command from Configure mode:

`(configure)#` **set intelligent-multicast static-session <*options*>**

Deleting Static Multicast Sessions

You can delete static multicast sessions using either the Web Agent or the CLI.

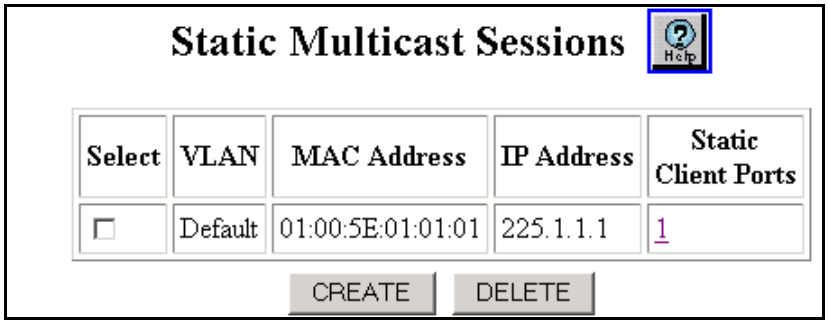
Web Agent Procedure

To delete a static multicast session using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **Static Sessions**.

   The Static Multicast Sessions Web page is displayed in the content pane. See Figure 20-9. See Table 20-7 for an explanation of the Static Multicast Sessions Web page parameters.

Figure 20-9. Static Multicast Sessions Web Page



2. Select the checkbox next to the static multicast session that you want to delete.

3. Click **DELETE** to remove the static multicast session.

Table 20-7. Static Multicast Sessions Web Page

Parameter	Definition
Select	Select the static multicast session to be deleted.
VLAN	Displays the VLAN on which the static multicast session exists.
MAC Address	Displays the MAC address of the static multicast session.
IP Address	Displays the IP address associated with the static multicast session, if available.
Static Client Ports	Displays the number of static client ports associated with the static multicast session and opens the Static Multicast Session Client Ports Web page.

CLI Command

To delete a static multicast session using the CLI, enter the following command from Configure mode:

`(configure)#` **clear intelligent-multicast static-session** <***options***>

Creating Static Client Ports

You can create and add static client ports using either the Web Agent or the CLI.

*** Note:** You can only add static client ports to static sessions on VLANs that the port is bound to or on sessions created for All VLANs.

Web Agent Procedure

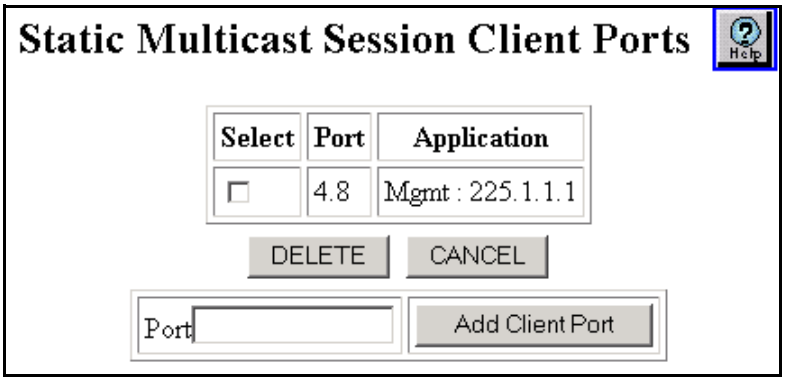
To create a static client port using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **Static Sessions**.

   The Static Multicast Sessions Web page is displayed in the content pane. See Figure 20-9.

2. Select the number from the **Static Client Ports** column. The Static Multicast Sessions Client Ports Web page is displayed. See Figure 20-10.

Figure 20-10. Static Multicast Sessions Clients Port Web Page



3. Enter the new port number to be added in the **Port** field,.

4. Click **Add Client Port**. The new port is added.

CLI Command

To create a static client port using the CLI, enter the following command from Configure mode:

`(configure)#` **set intelligent-multicast static-client-port** <***options***>

Deleting Static Client Ports

Static client ports can be deleted from either the Web or the CLI.

Web Agent Procedure

To delete static client ports using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **Static Sessions**.

   The Static Multicast Sessions Web page is displayed in the content pane. See Figure 20-9.

2. Select the number from the **Static Client Ports** column. The Static Multicast Sessions Client Ports Web page is displayed. See Figure 20-10.

3. Select a port and click **DELETE** to remove the static client port, or **CANCEL** to restore previous settings.

CLI Command

To delete static client ports using the CLI, enter the following command from Configure mode:

`(configure)#` **clear intelligent-multicast static-client-port** <***options***>

Managing IGMP Snooping

This section provides the following procedures for managing Internet Group Management Protocol (IGMP) snooping:

- Enabling IGMP Snooping
- Viewing IGMP Snooping Statistics

Enabling IGMP Snooping

You can enable IGMP snooping using either the Web Agent or the CLI.

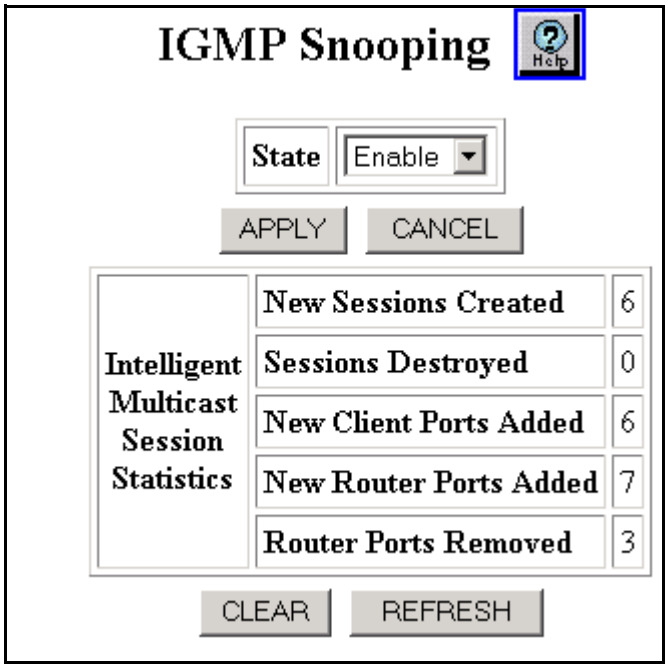
Web Agent Procedure

To enable IGMP Snooping using the Web Agent:

1. In the navigation pane, expand the **L2 Switching** > **Intelligent Multicast** folders, and then click **IGMP Snooping**.

   The IGMP Snooping Web page is displayed in the content pane. See Figure 20-11.

Figure 20-11. IGMP Snooping Web Page



*** Note:** IGMP snooping only works on VLANS that have an IGMP interface enabled. You must also enable intelligent multicasting on a global basis.

2. Select **Enable** from the State field.

3. Click **APPLY** to complete the operation or **CANCEL** to ignore the operation.

Table 20-8. IGMP Snooping Web Page Parameters

Parameter	Definition
Enable State	Select **Enable** to allow the IGMP snooping on the interface. The default value is **Disable**.
Intelligent Multicast Session Statistics	• **New Sessions Created** - Displays the number of new sessions created by IGMP snooping. • **Sessions Destroyed** - Displays the number of sessions removed by IGMP snooping. • **New Client Ports Added** - Displays the number of new client ports added by IGMP snooping. • **New Router Ports Added** - Displays the number of new router ports added by IGMP snooping. • **Router Ports Removed** - Displays the number of router ports removed by IGMP snooping.

CLI Command

To enable IGMP Snooping using the CLI, enter the following command from Configure mode:

```
(configure)# set igmp-snooping enable
```

Viewing IGMP Snooping Statistics

You can view IGMP snooping statistics using either the Web Agent or the CLI.

Web Agent Procedure

To view IGMP snooping statistics using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click IGMP Snooping.

   The IGMP Snooping Web page is displayed in the content pane. See Figure 20-11.

2. See Table 20-8 for an explanation of the IGMP Snooping Web page parameters:

3. Click:
   - **CLEAR** to clear the statistics.
   - **REFRESH** to refresh the contents of the table.

CLI Command

To view IGMP snooping statistics, use the following CLI command:

> **show igmp-snooping statistics**

Managing the LGMP Server

This section provides the following procedures for managing the Lucent Group Management Protocol (LGMP) Server:

- Configuring the LGMP Server
- Viewing the LGMP Server Statistics
- Viewing the LGMP Server Statistics per VLAN

Configuring the LGMP Server

You can configure the LGMP server using either the Web Agent or the CLI.

Web Agent Procedure

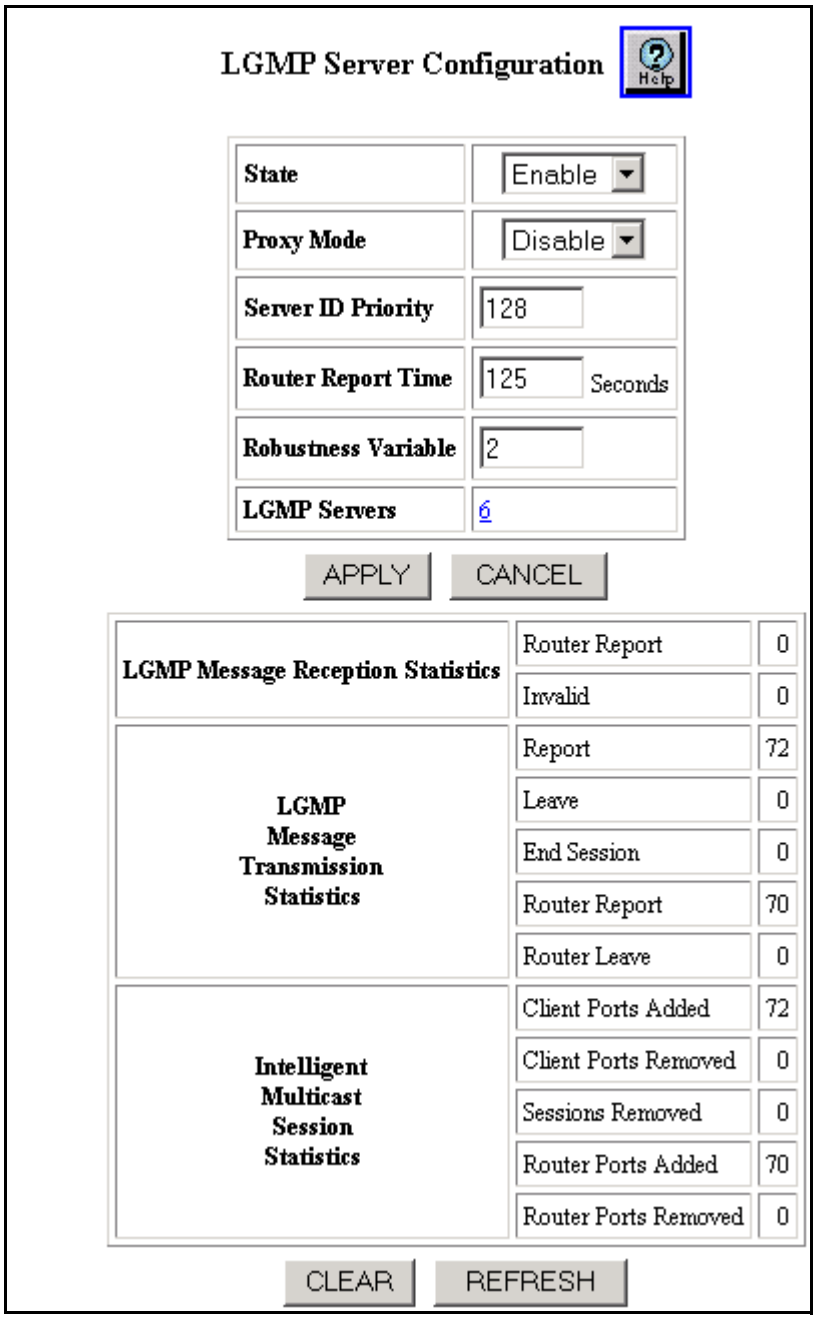
To configure the LGMP server using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **LGMP Server**.

   The LGMP Server Configuration Web page is displayed in the content pane. See Figure 20-12.

   ***Note:** You must enable intelligent multicasting on a global basis to activate the LGMP Server. LGMP serving is only active on VLANs with an IGMP interface enabled and IGMP snooping globally enabled.

Figure 20-12. LGMP Server Configuration Web Page



2. See Table 20-9 to configure the LGMP Server Configuration Web page parameters.

3. Click...

   — **APPLY** to save your changes.

   — **CANCEL** to restore previous settings.

Table 20-9. LGMP Server Configuration Web Page Parameters

Parameter	Definition
Enable State	Select to enable or disable LGMP server configuration. The default value is Disabled.
Proxy Mode	Select to enable or disable Proxy mode. Proxy mode allows the server to send LGMP router report and leave messages on behalf of neighbor routers on the same VLAN.
Server ID Priority	The priority of the LGMP server on this switch. The server ID priority and the IP address associated with the VLAN determine whether the LGMP server wins LGMP distribution election. The lower the number the more likely it will win the election. The valid range is 0 to 255. The default value is 128.
Router Report Time	The time interval (in seconds) between router reports sent by the LGMP server in distributor state. The valid range is 10 to 10,000. The default value is 125 seconds.
Robustness Variable	The scalar value used by non-distributor LGMP servers when timing out the LGMP server in the distributor state. The valid range is 2 to 10. The default value is 2.
LGMP Servers	Displays the number of LGMP servers and opens the LGMP Server Display per VLAN Web page. **Note: Only VLANs that have an active IGMP interface can be LGMP servers.**
LGMP Message Reception Statistics	• **Router Report** - Displays the number of LGMP router report messages received. • **Invalid** - Displays the number of LGMP messages received with an invalid payload.
	1 of 2

Table 20-9. LGMP Server Configuration Web Page Parameters

Parameter	Definition
LGMP Message Transmission Statistics	• **Report** - Displays the number of LGMP report messages transmitted. • **Leave** - Displays the number of LGMP leave messages transmitted. • **End Session** - Displays the number of LGMP end session messages transmitted. • **Router Report** - Displays the number of LGMP router report messages transmitted. • **Router Leave** - Displays the number of LGMP router leave messages transmitted.
Intelligent Multicast Session Statistics	• **Client Ports Added** - Displays the number of add client events generated by LGMP. • **Client Ports Removed** - Displays the number of remove client events generated by LGMP. • **Sessions Removed** - Displays the number of remove session events generated by LGMP. • **Router Ports Added** - Displays the number of add router events generated by LGMP. • **Router Ports Removed** - Displays the number of remove router events generated by LGMP.
	2 of 2

CLI Command

To configure the LGMP server using the CLI, enter the following command from Configure mode:

`(configure)#` **set lgmp server** <***options***>

Viewing the LGMP Server Statistics

You can view the LGMP server statistics using either the Web Agent or the CLI.

Web Agent Procedure

To enable/view the LGMP server using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **LGMP Snooping**.

   The LGMP Server Configuration Web page is displayed in the content pane. See Figure 20-12.

2. Select **Enable** from the **State** field, if it is not already displayed.

3. Click **CLEAR** to clear the statistics or **REFRESH** to refresh the contents of the table.

CLI Command

To view LGMP server statistics, use the following CLI command:

> **show lgmp server statistics**

Viewing the LGMP Server Statistics per VLAN

You can view the LGMP server statistics per VLAN using either the Web Agent or the CLI.

Web Agent Procedure

To modify the LGMP server display per VLAN using the Web Agent:

1. In the navigation pane, expand the **L2 Switching** > **Intelligent Multicast** folders, and then click **LGMP Server**.

   The LGMP Server Configuration Web page is displayed in the content pane.

*** Note:** You must enable intelligent multicasting on a global basis to activate your LGMP server per VLAN configuration

2. Select the number from the **LGMP Servers** column. The LGMP Server Display per VLAN Web page is displayed in the content pane. See Figure 20-13.

3. See Table 20-10 to configure the **LGMP Server Display per VLAN**.

Figure 20-13. LGMP Server per VLAN Web Page

LGMP Server Display per VLAN

Select	VLAN	State	LGMP Message Reception Statistics		LGMP Message Transmission Statistics					Intelligent Multicast Session Statistics				
			Router Report	Invalid	Report	Leave	End Session	Router Report	Router Leave	Client Ports Added	Client Ports Removed	Sessions Removed	Router Ports Added	Router Ports Removed
☐	VLAN34	Distributor	0	0	0	0	0	15	0	0	0	0	15	0
☐	VLAN52	Distributor	0	0	0	0	0	2	0	0	0	0	2	0
☐	VLAN54	Distributor	0	0	0	0	0	16	0	0	0	0	16	0
☐	VLAN55	Distributor	0	0	0	0	0	16	0	0	0	0	16	0
☐	VLAN74	Distributor	0	0	0	0	0	15	0	0	0	0	15	0
☐	VLAN75	Distributor	0	0	90	0	0	15	0	90	0	0	15	0

CLEAR CLEAR ALL REFRESH

4. Click:
   - **CLEAR** to reset selected row information.
   - **CLEAR ALL** to reset all statistics.
   - **REFRESH** to view the latest information.

Table 20-10. LGMP Server Display per VLAN Web Page Parameters

Parameter	Definition
Select	Select the LGMP server to modify.
VLAN	Displays the VLAN associated with the LGMP server.
State	Displays the current state of the LGMP server. • **Distributor** - The LGMP server serves LGMP messages to LGMP clients. • **Non-Distributor** - The LGMP server monitors the current distributor.
LGMP Message Reception Statistics	• **Router Report** - Displays the number of LGMP router report messages received per VLAN. • **Invalid** - Displays the number of LGMP messages received with an invalid payload per VLAN.
LGMP Messages Transmission Statistics	• **Report** - Displays the number of LGMP report messages transmitted per VLAN. • **Leave** - Displays the number of LGMP leave messages transmitted per VLAN. • **End Session** - Displays the number of LGMP end session messages transmitted per VLAN. • **Router Report** - Displays the number of LGMP router report messages transmitted per VLAN. • **Router Leave** - Displays the number of LGMP router leave messages transmitted per VLAN.
Intelligent Multicast Session Statistics	• **Client Ports Added** - Displays the number of add client events generated by LGMP per VLAN. • **Client Ports Removed** - Displays the number of remove client events generated by LGMP per VLAN. • **Sessions Removed**- Displays the number of remove session events generated by LGMP per VLAN. • **Router Ports Added** - Displays the number of add router events generated by LGMP per VLAN. • **Router Ports Removed** - Displays the number of remove router events generated by LGMP per VLAN.

CLI Command

To view the LGMP server statistics per VLAN, use the following CLI command:

> show lgmp server statistics vlan {all | <*vlan-id*> | name <*vlan-name*>}

Managing the LGMP Client

This section provides the following procedures for managing the Lucent Group Management Protocol (LGMP) client:

- Enabling the LGMP Client
- Viewing LGMP Clients Statistics
- Viewing LGMP Clients per VLAN statistics

Figure 20-14. LGMP Client Configuration Web Page

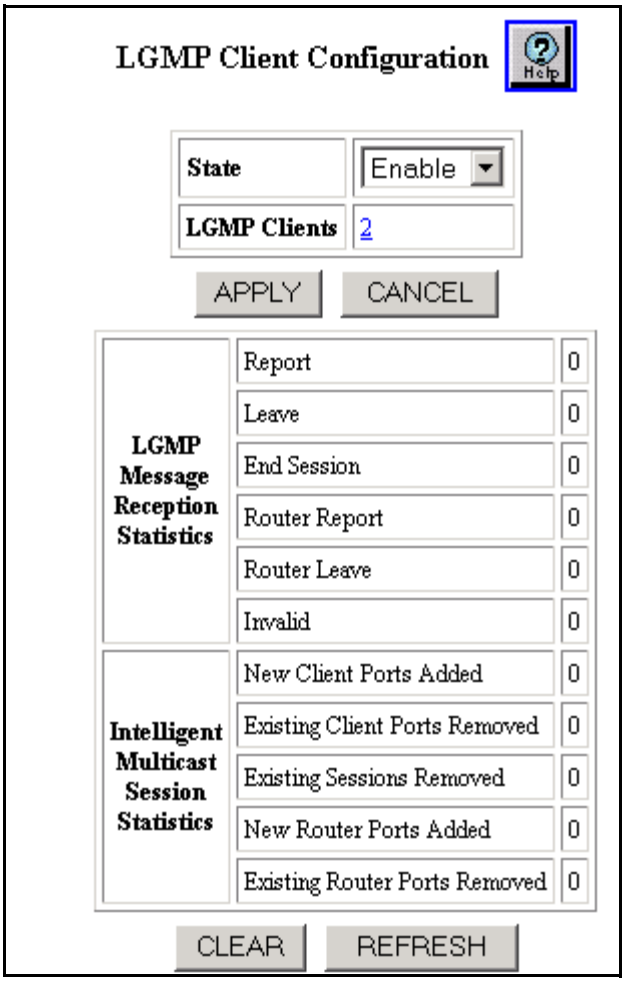


Table 20-11. LGMP Client Configuration Web Page Parameters

Parameter	Definition
Enable State	Select to enable or disable LGMP client. The default value is Disabled.
LGMP Clients	Displays the number of LGMP clients per VLAN and opens the LGMP Client Displays per VLAN Web page.
LGMP Message Reception Statistics	• **Report** - Displays the number of LGMP report messages received. • **Leave** - Displays the number of LGMP leave messages received. • **End Session** - Displays the number of LGMP end session messages received. • **Router Report** - Displays the number of LGMP router report messages received. • **Router Leave** - Displays the number of LGMP router leaves messages received. • **Invalid** - Displays the number of LGMP messages received with an invalid payload.
Intelligent Multicast Session Statistics	• **New Client Ports Added** - Displays the number of new clients added by LGMP. • **Existing Client Ports Removed** - Displays the number of clients removed by LGMP. • **Existing Sessions Removed** - Displays the number of sessions removed by LGMP. • **New Router Ports Added** - Displays the number of new routers added by LGMP. • **Existing Router Ports Removed** - Displays the number of new routers removed by LGMP.

Enabling the LGMP Client

You can enable the LGMP client statistics using either the Web Agent or the CLI.

Web Agent Procedure

To enable an LGMP client using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **LGMP Client**.

   The LGMP Client Configuration Web page is displayed in the content pane. See Figure 20-14.

   ***Note:** You must enable intelligent multicasting on a global basis to activate LGMP client configuration.

2. Select **Enable** from the State field pull-down menu, if it is not already displayed.

3. Select **APPLY** to complete the operation, or **CANCEL** to ignore the operation.

CLI Command

To enable an LGMP client using the CLI, enter the following command from Configure mode:

```
(configure)# set lgmp client enable
```

Viewing LGMP Clients Statistics

You can view LGMP clients statistics using either the Web Agent or the CLI.

Web Agent Procedure

To view LGMP Client statistics using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **LGMP Client**.

   The LGMP Client Configuration Web page is displayed in the content pane. See Figure 20-14.

2. See Table 20-11 for an explanation of the LGMP Client Configuration Web page parameters

3. Click **CLEAR** to clear the statistics, or **REFRESH** to refresh the contents of the table.

CLI Command

To view the LGMP client statistics, use the following CLI command:

```
> show lgmp client statistics
```

Viewing LGMP Clients per VLAN statistics

You can view LGMP clients per VLAN statistics using either the Web Agent or the CLI.

Web Agent Procedure

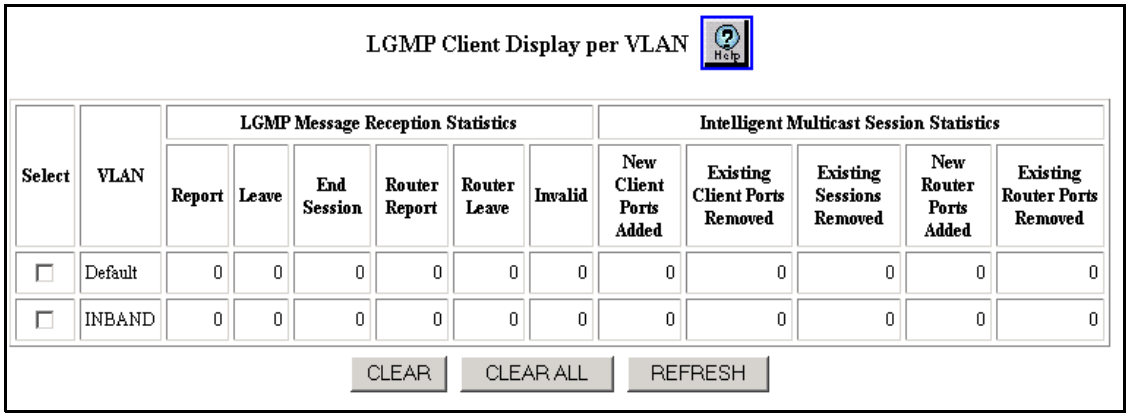
To view LGMP client statistics per VLAN using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **LGMP Client**.

   The LGMP Client Configuration Web page is displayed in the content pane. See Figure 20-14.

2. Select the number from the **LGMP Clients** field. The **LGMP Client Display per VLAN** Web page is displayed in the content pane. See Figure 20-15.

Figure 20-15. LGMP Client Display per VLAN Web Page



3. See Table 20-12 for an explanation of the LGMP Client Display per VLAN Web page parameters.

4. Click...

   — **CLEAR** to reset selected row information.

   — **CLEAR ALL** to reset all statistics.

   — **REFRESH** to view the latest information.

Table 20-12. LGMP Client Display per VLAN Web Page Parameters

Parameter	Definition
Select	Select the LGMP client statistics to clear.
VLAN	Displays the VLAN associated with the LGMP client.
LGMP Message Reception Statistics	• **Report** - Displays the number of LGMP report messages received per VLAN. • **Leave** - Displays the number of LGMP leave messages received per VLAN. • **End Session** - Displays the number of LGMP end session messages received per VLAN. • **Router Report** - Displays the number of LGMP router report messages received per VLAN. • **Router Leave** - Displays the number of LGMP router leaves messages received per VLAN. • **Invalid** - Displays the number of LGMP messages received with an invalid payload per VLAN.
Intelligent Multicast Session Statistics	• **New Client Ports Added** - Displays the number of new clients added by LGMP per VLAN. • **Existing Client Ports Removed** - Displays the number of clients removed by LGMP per VLAN. • **Existing Sessions Removed** - Displays the number of sessions removed by LGMP per VLAN. • **New Router Ports Added** - Displays the number of new routers added by LGMP per VLAN. • **Existing Router Ports Removed** - Displays the number of new routers removed by LGMP per VLAN.

CLI Command

To view LGMP client statistics per VLAN, use the following CLI command:

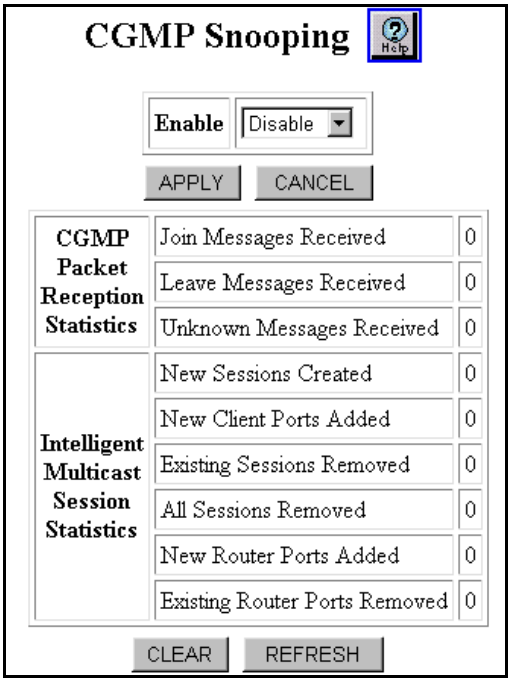
> **show lgmp client statistics vlan {all | <*vlan-id*> | name <*vlan-name*>}**

Managing CGMP Snooping

You can manage Cisco Group Management Protocol (CGMP) snooping by using either the Web Agent or the CLI. This section provides the following procedures:

- Enabling CGMP Snooping
- Viewing CGMP Snooping

Figure 20-16. CGMP Snooping Web Page



Enabling CGMP Snooping

You can enable CGMP Snooping using either the Web Agent or the CLI.

Web Agent Procedure

To enable CGMP snooping using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **CGMP Snooping**.

   The CGMP Snooping Web page is displayed in the content pane. See Figure 20-16.

   ***Note:** You must enable intelligent multicasting on a global basis to activate CGMP snooping configuration.

2. See Table 20-13 for an explanation of the CGMP Snooping Web page parameters.

3. Click...

   — **APPLY** to save your changes

   — **CANCEL** to restore previous settings.

Table 20-13. CGMP Snooping Web Page Parameters

Parameter	Definition
Enable State	Select to enable or disable CGMP snooping.
CGMP Packet Reception Statistics	• **Join Messages Received** - Displays the number of CGMP join messages received. • **Leave Messages Received** - Displays the number of CGMP leave messages received. • **Unknown Messages Received** - Displays the number of unknown CGMP messages received.
Intelligent Multicast Session Statistics	• **New Sessions Created** - Displays the number of new multicast sessions created by CGMP snooping. • **New Client Ports Added** - Displays the number of new client ports added to a multicast session. • **Existing Sessions Removed** - Displays the number of existing multicast sessions that have been removed by CGMP snooping. • **All Sessions Removed** - Displays the number of times that all multicast sessions created by CGMP snooping were removed. • **New Router Ports Added** - Displays the number of new router ports added by CGMP snooping. • **Existing Router Ports Removed** - Displays the number of router ports that were created by the CGMP snooper and were removed.

CLI Command

To enable CGMP snooping using the CLI, enter the following command from Configure mode:

```
(configure)# set cgmp enable
```

Viewing CGMP Snooping

You can view CGMP Snooping using either the Web Agent or the CLI.

Web Agent Procedure

To view CGMP snooping using the Web Agent:

1. In the navigation pane, expand the **L2 Switching > Intelligent Multicast** folders, and then click **CGMP Snooping**.

   The CGMP Snooping Web page is displayed in the content pane. See Figure 20-16.

2. View the CGMP Snooping Web page fields.

3. Click:

   — **CLEAR** to clear the statistics

   — **REFRESH** to refresh the contents of the table

CLI Command

To view CGMP snooping using the CLI, enter the following command from User mode:

> show cgmp statistics [detailed]

21 Monitoring the Avaya Multiservice Switch

Overview

The following information and procedures provided in this chapter pertain to layer 2 and layer 3 module configurations:

- Interpreting Front Panel LED Displays
- Checking Active Alarms
- Using the Event Subsystem
- Configuring Syslog Event Reporting
- Configuring Utilization Monitoring

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.

Interpreting Front Panel LED Displays

For an explanation of the switch front-panel LEDs, see Table 21-1.

Table 21-1. Front Panel LED Explanations

Module...	LED...	Behavior...	Indication...
All Modules	Module Status	Solid green	Normal operation.
		Flashing orange	Diagnostic failure.
		Off	Module not operational or not receiving power.
Gigabit Modules	TX/RX	Flashing orange	Port sending/receiving traffic.
		Off	Port not sending/receiving traffic.

1 of 2

Table 21-1. Front Panel LED Explanations

Module...	LED...	Behavior...	Indication...
	Port	Solid green	Port enabled with link up.
		Flashing green	Port disabled.
		Flashing orange	Hardware failure.
		Off	No link.
10/100 Modules	Port	Solid green, with orange flash	Port enabled and sending and receiving traffic Traffic indicated by yellow flashes.
		Flashing green	Port disabled with link up.
		Flashing orange	Hardware failure.
		Off	No link.
			2 of 2

Checking Active Alarms

Each switch stores a table of active alarms from which you can view a list of open issues without having to view the entire event log. This provides you with a quick snapshot of the switch's health.

You can view the active alarm table by using either the Web Agent or the CLI.

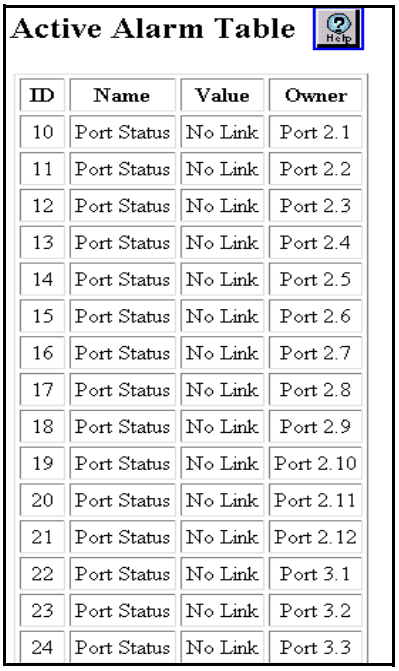
Web Agent Procedure

To view the **Active Alarm Table** using the Web Agent:

1. In the navigation pane, expand the **Events** folder, and then click **Active Alarms**.

   The Active Alarm Table is displayed in the content pane. See Figure 21-1.

Figure 21-1. Active Alarm Table Web Page



Active Alarm Table

ID	Name	Value	Owner
10	Port Status	No Link	Port 2.1
11	Port Status	No Link	Port 2.2
12	Port Status	No Link	Port 2.3
13	Port Status	No Link	Port 2.4
14	Port Status	No Link	Port 2.5
15	Port Status	No Link	Port 2.6
16	Port Status	No Link	Port 2.7
17	Port Status	No Link	Port 2.8
18	Port Status	No Link	Port 2.9
19	Port Status	No Link	Port 2.10
20	Port Status	No Link	Port 2.11
21	Port Status	No Link	Port 2.12
22	Port Status	No Link	Port 3.1
23	Port Status	No Link	Port 3.2
24	Port Status	No Link	Port 3.3

CLI Command

Use the **show alarms** CLI command to view the active alarm table. You can enter this command in User mode.

Using the Event Subsystem

Overview

The Avaya P580 and P882 Multiservice switches have two activity logs:

- **Event Log** - stores a large table of events. You can set the size of the table. Because these events are stored in switch memory, the list is cleared each time the switch reboots.
- **Shutdown Log** - stores the same information as the event Log, but generally in a smaller table because the table is stored in the switch's nonvolatile RAM (NVRAM). This log list is particularly useful in assessing the cause of a switch failure because the information is retained even after the switch restarts.

This sections contains procedures for the following tasks:

- Configuring Event Notification
- Configuring Notification of Protocol Events
- Setting Log Size
- Viewing the Event and Shutdown Logs
- Clearing the Event Log
- Viewing Event Statistics

Configuring Event Notification

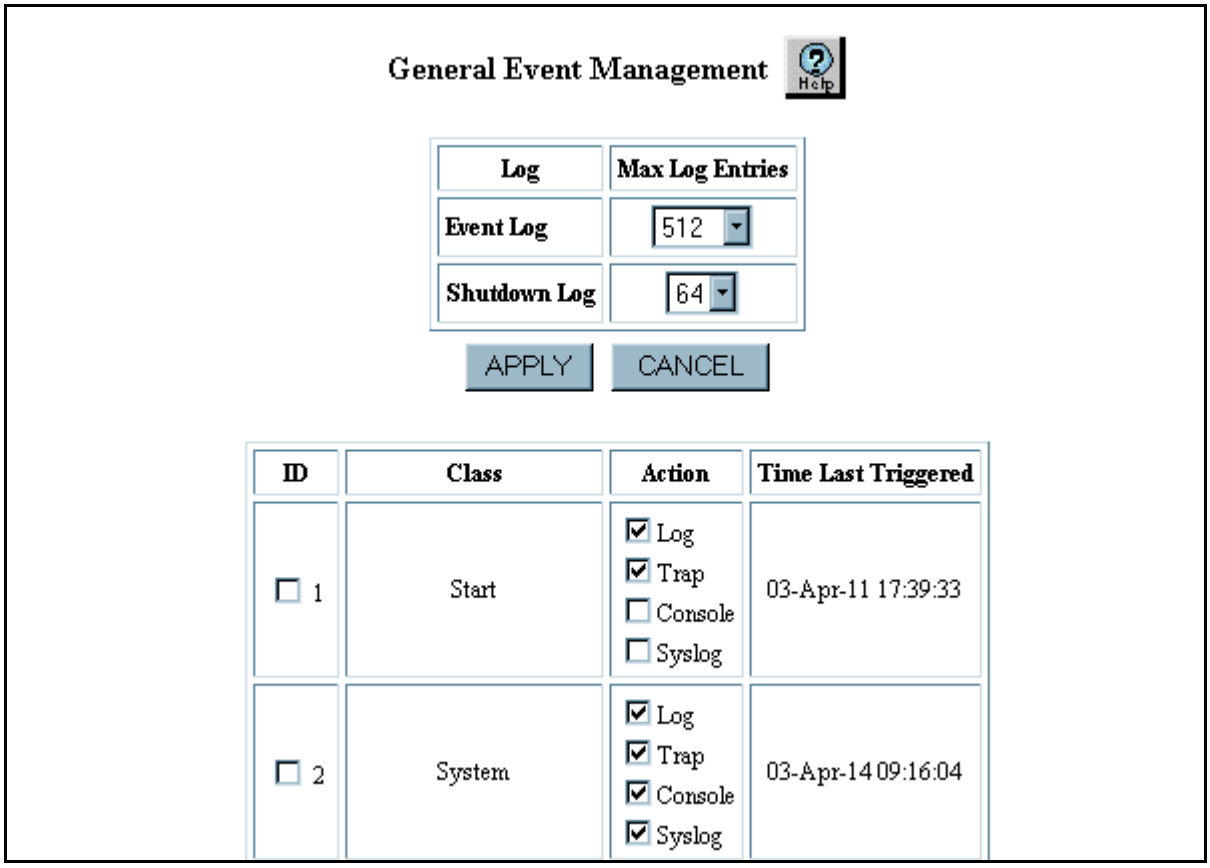
Web Agent Procedure

To configure event notification:

1. In the navigation pane, expand the **Events** folder and then click **General Events**.

   The General Event Management Web page is displayed in the content pane. See Figure 21-2.

Figure 21-2. General Event Management Web Page



2. In the **ID** field, select the classes of events that you want logged. Table 21-2 lists the classes of events that you can log and what events are logged for each class.

   If you enable event notification for CLI, SNMP, RIP, OSPF, DVMRP, LDAP, Apple Talk, or VRRP, you must set which categories of protocol events generate notifications. For information on how to set the categories of protocol events that generate notifications, see "Configuring Notification of Protocol Events" later in this chapter.

*** Note:** Event IDs 20 and 21 are currently not used and reserved for future use.

3. In the **Action** field, select the notifications that you want for each logged event class. Options are:

   - Log—Events are logged in the event log and in the shutdown log.
   - Trap—Events generate SNMP traps.
   - Console—Events generate a message on the console.
   - Syslog—Events are forwarded to syslog servers.

4. Click **APPLY**.

Table 21-2. Event Classes

Class	Determines whether the switch sends a notification for...
Start	Starts of the system.
System	System events.
Configuration	Each configuration change (for example, enabling and disabling ports).
Temperature Status	Changes in temperature status. Temperature status messages could precede a switch shutdown, and are often critical.
Resource	Changes in system resources.
Fan Status	Fan status. Fan failures will eventually lead to overheating the system. The fan status message provides a good early warning for a failure that could eventually cause the switch to shut down.
Service Port Status and **User Port Status**	Status changes in service ports and user ports. Set a port as either a *service port* or *user port* on the Detailed Physical Port Configuration Web page. This feature makes it possible for you to use different notification levels for critical (service ports), if desired.
Power Status	The addition or removal of a power supply
Bridge Status	Changes in bridge status.
Switch Fabric	Failures in the switch fabric. These failures are critical and should be monitored closely.
OSPF	OSPF events if OSPF protocol event logging is enabled.
RIP	RIP events if RIP protocol event logging is enabled.
LDAP	LDAP events if LDAP protocol event logging is enabled.
AppleTalk	AppleTalk events if AppleTalk protocol event logging is enabled.
Authentication Failure	Authentication failures. This is a security-related feature used to detect unauthorized SNMP activity.
	1 of 2

Table 21-2. Event Classes

Class	Determines whether the switch sends a notification for...
Redundant CPU	Changes in status of a redundant CPU. Notification is sent if: • The status changes from standby to active or vice versa. • The active supervisor loses or establishes contact with the standby supervisor.
DVMRP	DVMRP events, if DVMRP event logging is enabled.
CLI	CLI events, if CLI event logging is enabled.
SNMP	SNMP events, if SNMP protocol event logging is enabled.
Unknown Mac Received	Unknown MAC addresses if received.
VRRP	VRRP events, if VRRP protocol event logging is enabled.
User Login Status	User login or logout.
Access List Rule Match	Packets that match access control rules.
Secure Protocol SSL/ SSH	SSH events.
	2 of 2

CLI Command

Use the following CLI commands to configure event notification:

- To log events in the event log and in the shutdown log, `(configure)#` **logging history {start | system | config | temp | resource | fan | power | service_port | user_port | auth_failure | bridge_stat | switch_fabric | ospf | dvmrp | rip | ldap | cli | snmp | appletalk | redundant_cpu | vrrp | unknown_mac | login_stat us | acl_log | ssl_ssh}**
- To generate SNMP traps, `(configure)#` **logging traps {start | system | config | temp | resource | fan | power | service_port | user_port | auth_failure | bridge_stat | switch_fabric | redundant_cpu | unknown_mac | snmp | login_status}**
- To generate event notification to the console, `(configure)#` **logging console {start | system | config | temp | resource | fan | power | service_port | user_port | auth_failure | bridge_stat | switch_fabric | ospf | dvmrp | rip | ldap | cli | snmp | appletalk | redundant_cpu | vrrp | unknown_mac | login_status | acl_log | ssl_ssh}**

- To forward events to syslog servers, `(configure)#` **set syslog facility {start | system | config | temp | resource | fan | service_port | user_port | power | bridge_stat | switch_fabric | ospf | rip | ldap | appletalk | auth_failure | redundant_cpu | dvmrp | cli | snmp | unknown_mac | vrrp | login_status | acl_log | ssl_ssh | all}**

Configuring Notification of Protocol Events

Overview

***Important:** If enabled, protocol event logging displays system messages that help Avaya Technical Support troubleshoot network problems. Avaya recommends that logging of protocol events be enabled only during troubleshooting sessions. If protocol event logging is enabled during normal network operation, the switch may display messages that users may incorrectly interpret as indications of protocol failures.

If you enable event notification for CLI, SNMP, RIP, OSPF, DVMRP, LDAP, Apple Talk, or VRRP, you must set which categories of the protocol events generate notifications.

*** Note:** Enabling logging of protocol events may cause the event log to rapidly fill with protocol events.

You can use either the Web Agent or the CLI to configure event notification for protocol events.

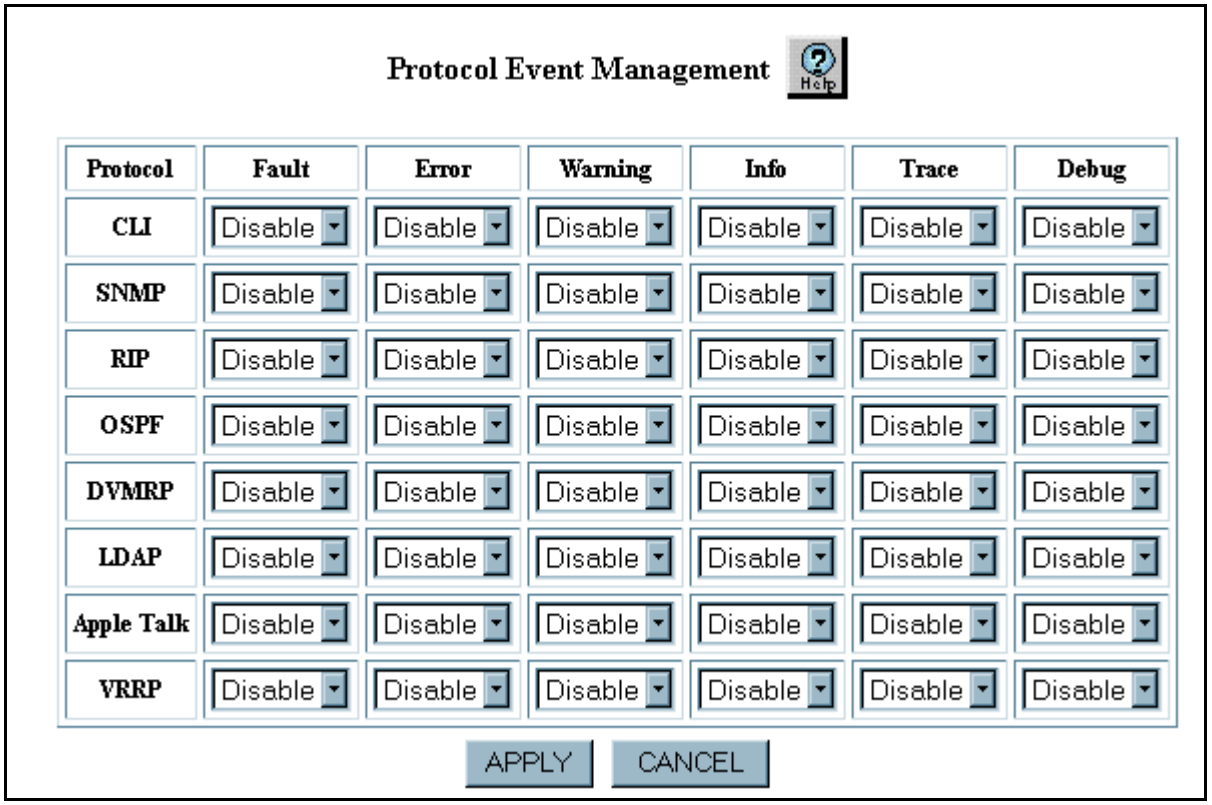
Web Agent Procedure

To set which categories of protocol events generate notifications:

1. In the navigation pane, expand the **Events** folder and then click **Protocol Events**.

   The Protocol Event Management Web page is displayed in the content pane. See Figure 21-3.

Figure 21-3. Protocol Event Management Web Page



2. Enable the categories of protocol events for which you want to generate notifications. See Table 21-3 for an explanation of each category.

3. Click **APPLY**.

Table 21-3. Protocol Event Categories

Event Category	Explanation
Fault	Serious errors that can cause a system crash, for example, panic. The default setting is **Disable**.
Error	Serious errors that will not cause a system crash but can contribute to protocol problems.
Warning	Noncritical errors. The default setting is **Disable**.
Info	Event details. The default setting is **Disable.**
Trace	Packet traces. The default setting is **Disable**. If you enable Trace logging, all protocol packets sent and received are logged as protocol events.
Debug	Event messages used to troubleshoot a network problem. The default setting is **Disable**.

CLI Command

To set which categories of protocol events generate notifications, use the following CLI command:

`(configure)#` **logging protocol event {rip | ospf | dvmrp | ldap | cli | snmp | appletalk | vrrp} {fault | error | warning | info | trace | debug}**

Setting Log Size

You can set the log size using either the Web Agent or the CLI.

Web Agent Procedure

To set the log size using the Web Agent:

1. In the navigation pane, expand the Events folder and then click **General Events**.

   The General Event Management Web page is displayed in the content pane. See Figure 21-2.

2. In the **Max Log Entries** field for the event log, select the number of entries that you want the event log to store.

   The default setting is 512 entries.

3. In the **Max Log Entries** field for the shutdown log, select the number of entries that you want the shutdown log to store.

   The default setting is 16 entries.

4. Click **APPLY**.

CLI Command

Use the following CLI commands to set the size of the event log and shutdown log:

- To set the size of the event log, `(configure)#` **logging history size {128 | 512 | 1024 | 2048}**
- To set the size of the shutdown log, `(configure)#` **logging shutdown size {16 | 32 | 64}**

Viewing the Event and Shutdown Logs

You can view the event log and shutdown log using either the Web Agent or the CLI.

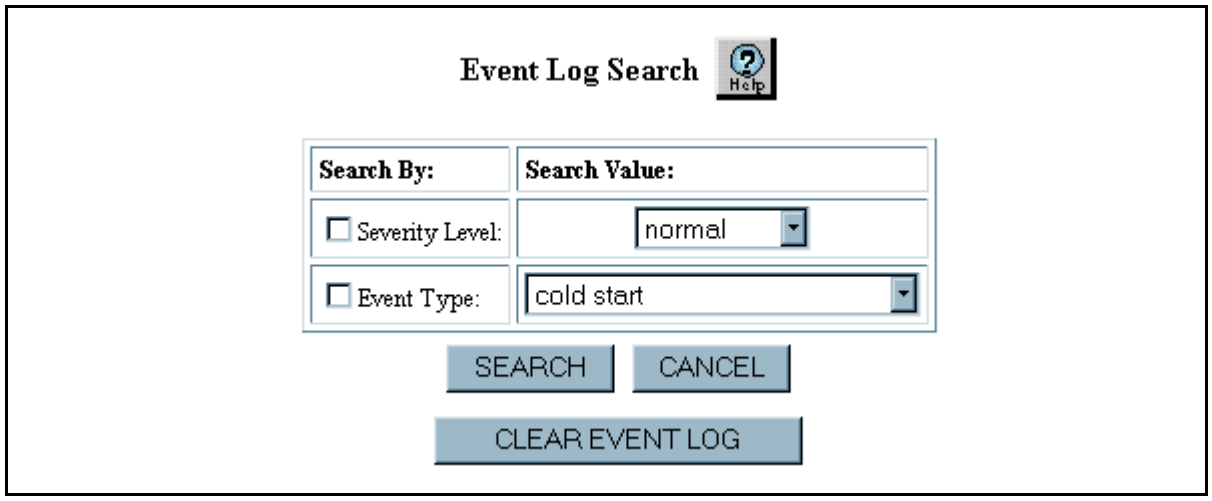
Web Agent Procedure

To view the event log or shutdown log using the Web Agent:

1. In the navigation pane, expand the **Events** folder and then click either **Event Log** or **Shutdown Log**.

   The Event Log Search Web page or Shutdown Log Search Web page is displayed in the content pane. See Figure 21-4 for the Event Log Search Web page. The Shutdown Log Search Web page is identical except for its title.

Figure 21-4. Event Log Search Web Page



2. To view only events of a specific severity:

   a. In the **Search By** column, select **Severity Level**.

   b. In the **Search Value** column, click the severity of events that you want to view.

3. To view only events of a specific type:

   a. In the **Search By** column, select **Event Type**.

   b. In the **Search Value** column, click the type of events that you want to view.

4. Click **Search**.

The event log or shutdown log is displayed in the content pane. See Figure 21-5. For an explanation of the event log fields, see Table 21-4.

*** Note:** If you do not select a severity level or event type, the entire event log or shutdown log is displayed.

Figure 21-5. Event Log Web page

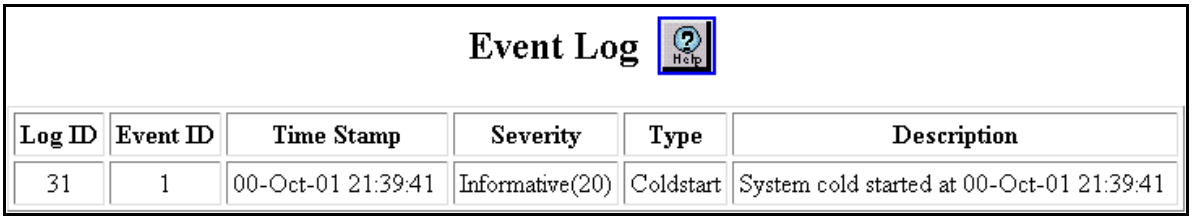


Table 21-4. Event Log Fields

Entry	Definition
Log ID	Displays the number of this event in the log FIFO (First In First Out).
Event ID	Displays an index that identifies the event class.
Time Stamp	Displays the date and time the event was recorded. The time is displayed in yy/mm/dd sequence.
Severity	Displays the severity of the event. The possibilities are: • Normal • Informative • Warning • Alarm • Error • Fatal
Type	Displays a description of the event type (for example System start and Status Change).
Description	Displays a text string that describes the specific event with the date and time of the event.

CLI Command

To view the event and shutdown logs, use the **> show logging [shutdown] [<*num-events*>]** CLI command.

Clearing the Event Log

You must have read-write access to clear the event log. This functionality is not available if you have read-only access.

Web Agent Procedure

To clear the event log:

1. In the navigation pane, expand the **Events** folder and then click **Event Log**.

   The Event Log Search Web page is displayed in the content pane. See Figure 21-4.

2. Click **CLEAR EVENT LOG**.

CLI Command

To clear the event log, use the `(configure)#` **logging clear** CLI command.

Viewing Event Statistics

Event statistics are available only in the Web Agent. You cannot use the CLI to view event statistics.

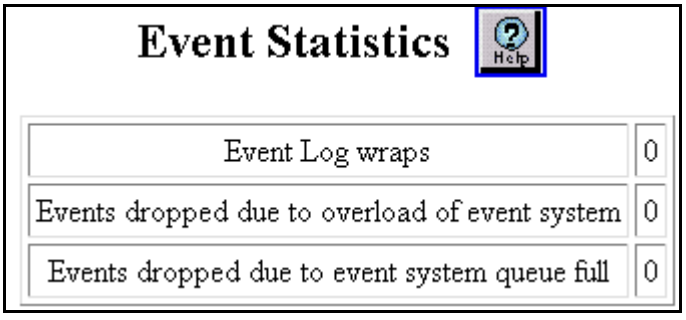
Web Agent Procedure

To view event statistics using the Web Agent:

1. In the navigation pane, expand the **Events** folder, and then click **Statistics**.

   The Event Statistics Web page is displayed in the content pane. See Figure 21-6.

Figure 21-6. Event Statistics Web Page



2. See Table 21-5 for an explanation on the Event Statistics Web page parameters.

Table 21-5. Event Statistics Web Page Parameters

Parameter	Definition
Event Log wraps	Displays the number of times the event log has wrapped. The event log may or may not wrap, depending on how many events have been sent to the event log and when it was last cleared. When the event log does wrap, the old events are discarded and replaced with the newest events.
Events dropped due to overload of event system	Displays the number of events that were dropped to prevent overloading.
Events dropped due to event system queue full	Displays the number of events that were dropped due to a full Event System queue.

Configuring Syslog Event Reporting

Overview

The Avaya Multiservice switches supports syslog event reporting. This feature makes it possible to forward specific system events to remote devices to be logged. You select the classes and severity of events that you want forwarded. The remote devices must run a syslog daemon and are called *syslog servers* or *collectors*.

Table 21-6 describes the different syslog severity levels. The switch logs error messages of the severity that you set and of all higher severities. For example, if you set the severity to **Warning**, error messages of severities Warning, Error, Alert, and Emergency are logged.

Table 21-6. Syslog Severity Levels

Severity Level	Description
Emergency	System Unusable
Alert	Immediate action needed
Error	Error Condition
Warning	Warning Condition
Normal	Normal but significant condition
Informational	Informational message only

You can specify a maximum of three remote syslog servers. Syslog messages from the Avaya Multiservice switch are unidirectional. No acknowledgement is expected from the syslog server.

*** Note:** If you enable syslog event reporting, Avaya recommends that you enable Simple Network Time Protocol (SNTP) on the switch and on the remote syslog servers. For information on SNTP, see "Enabling the Simple Network Time Protocol" in Chapter 3, "Configuring System Information."

You can also set the event types, also called "facilities," for which syslog events are generated. For more information, see "Configuring Event Notification."

Configuring Syslog

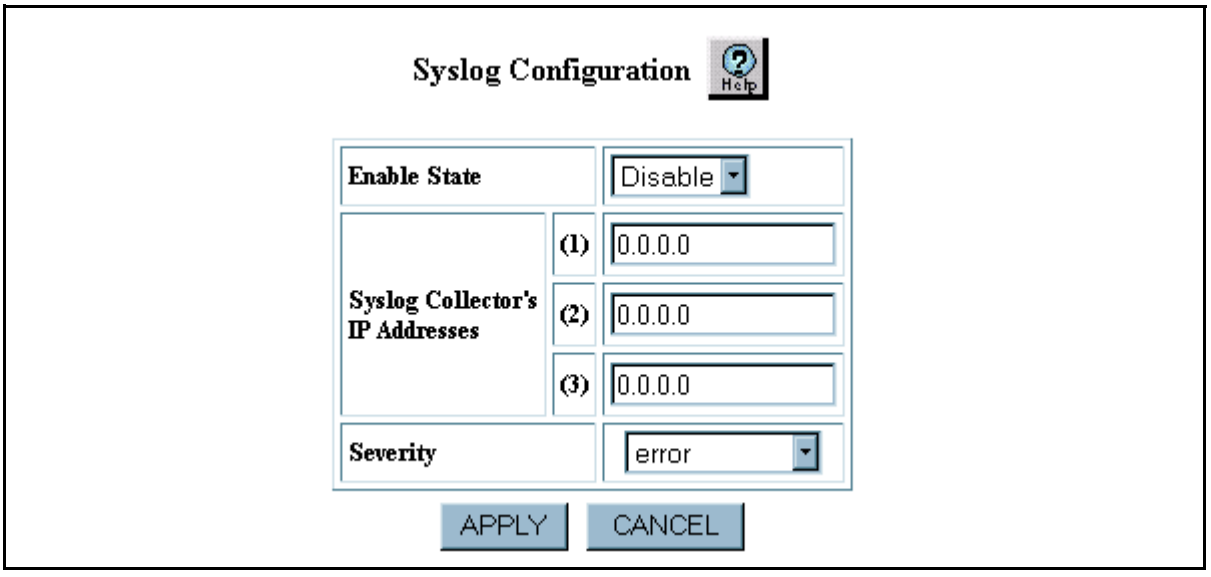
Web Agent Procedure

To configure syslog event reporting:

1. Expand the **Events** folder.
2. Click **Syslog**.

   The Syslog Configuration Web page is displayed.

Figure 21-7. Syslog Configuration Web Page



3. In the **Enable State** field, select **Enable** or **Disable**:
   - **Enable**—Enables syslog reporting to the IP addresses that you enter in the **Syslog Collector's IP Addresses** fields.
   - **Disable**—Disables syslog reporting. Syslog is disabled by default.
4. In the **Syslog Collector's IP Addresses** fields, enter the IP addresses of the remote syslog servers. You can specify a maximum of three syslog servers.
5. In the **Severity** field, select a severity level. See Table 21-6 for more information about each severity level. The default setting for this field is **error**.

   The switch logs error messages of the severity that you select and of all higher severities.
6. Click **Apply**.

CLI Command

Use the following CLI commands to configure syslog event reporting:

- To enable or disable syslog event reporting, `(configure)#` **set syslog {enable | disable}**
- To configure which events are logged, `(configure)#` **set syslog facility {start | system | config | temp | resource | fan | service_port | user _port | power | bridge_stat | switch_fabric | ospf | rip | ldap | appletalk | auth_failure | redundant_cpu | dvmrp | cli | snmp | unknown_mac | vrrp | login_status | acl_log | ssl_ssh | all}**
- To set IP addresses of remote syslog servers, `(configure)#` **set syslog server_ip <*ip_address*>**
- To set a severity level, `(configure)#` **set syslog severity {emergency | alert | error | warning | normal | informational}**
- To view the current configuration for syslog even reporting, `>` **show syslog config**
- To view events in the syslog buffer, `>` **show syslog buffer**

Configuring Utilization Monitoring

Overview

Utilization monitoring makes it possible for you to monitor utilization of the:

- CPU—the number of packets per second that the CPU on the supervisor module routes (slow path).
- Forwarding engine—the number of packets per second that the 80-series forwarding engines route. This statistic includes packets routed by the forwarding engines on 80-series media modules (FIRE) and on the supervisor module (FORE).

For more information about the slow path and FIRE and FORE paths, see "Routing Overview" in Chapter 1, "Introduction."

For each of these statistics, you can set a high threshold. If this threshold is exceeded, the switch logs the event in the event log, displays a message in the CLI, and sends an SNMP trap to the Trap receiver.

This section contains procedures for the following tasks:

- Enabling Utilization Monitoring
- Enabling Event Logging of Utilization
- Setting a Utilization Threshold
- Viewing Utilization Settings
- Viewing CPU Utilization Statistics
- Viewing Statistics for Forwarding Engine Utilization

You can configure utilization monitoring only by using the CLI. This feature is not available in the Web Agent.

Enabling Utilization Monitoring

To *enable* utilization monitoring, enter Global Configuration mode and use the following command:

set utilization monitoring {cpu | forwarding-engine}

The default setting for utilization monitoring is disabled.

Table 21-7. Keywords, Arguments, and Options

Keyword, Argument, or Option	Definition
cpu	Enables monitoring of CPU utilization.
forwarding-engine	Enables monitoring of 80-series forwarding engines.

To *disable* utilization monitoring, enter Global Configuration mode and use the following command:

clear utilization monitoring {cpu | forwarding-engine}

Enabling Event Logging of Utilization

To *enable* event logging of utilization, enter Global Configuration mode and use the following command:

set utilization threshold-event {cpu | forwarding-engine}

The default setting for event logging of utilization is disabled. If you enable event logging of utilization but do not specify a utilization threshold, the switch logs an event if the CPU utilization or forwarding engine utilization exceeds 95 percent.

Table 21-8. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
cpu	Enables event logging for CPU utilization.
forwarding-engine	Enables event logging for utilization of 80-series forwarding engines.

To *disable* event logging of utilization, enter Global Configuration mode and use the following command:

clear utilization threshold-event {cpu | forwarding-engine}

Setting a Utilization Threshold

To set a utilization threshold, enter Global Configuration mode and use the following command:

set utilization high-threshold {cpu | FIRE | FORE} <*utilization-percent*>

The default setting for the utilization threshold is 95 percent.

Table 21-9. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
cpu	Sets the high threshold for CPU utilization. 100% CPU utilization is the total capacity of the supervisor module to forward slow path traffic. When 100% utilization is reached, the performance of the switch may degrade.
FIRE	Sets the high threshold for utilization of the forwarding engines on 80-series media modules. 100% FIRE utilization is the total capacity of the forwarding engines on 80-series media modules to forward in band traffic. When 100% utilization is reached, the performance of the switch may degrade.
FORE	Sets the high threshold for utilization of the forwarding engine on the supervisor module. 100% FORE utilization is the total capacity of the supervisor module to forward out-of-band traffic. When 100% utilization is reached, the performance of the switch may degrade.
<*utilization-percent*>	The high threshold at which you want the switch to log an event in the event log. Enter a value from 50 through 99. The default setting is 95.

Clearing the utilization threshold resets it to 95 percent, its default setting. To clear the utilization threshold, enter Global Configuration mode and use the following command:

clear utilization high-threshold {cpu | FIRE | FORE} <*utilization-percent*>

Viewing Utilization Settings

To view the current utilization settings, use the following command:

show utilization settings

You can enter this command from user mode.

A sample of the display is as follows:

```
CPU monitoring is disabled
Forwarding Engine monitoring is disabled

CPU threshold level is 95 percent
Forwarding Engines:
FIRE threshold level is 95 percent
FORE threshold level is 95 percent

CPU threshold event generation is disabled
Forwarding Engine threshold event generation is
disabled
```

Viewing CPU Utilization Statistics

To display the CPU utilization statistics, use the following command:

show utilization results cpu

You can enter this command from user mode.

The command displays a history of utilization and the average utilization. A sample of the display is as follows:

```
Average CPU Utilization over the 60 second sample
window: 0 percent

Individual Sample Utilizations (sorted from most
recent to oldest):
Sample 0: 2 percent
Sample 1: 2 percent
Sample 2: 2 percent
Sample 3: 5 percent
Sample 4: 3 percent
Sample 5: 2 percent
Sample 6: 2 percent
Sample 7: 2 percent
```

```
Sample 8: 3 percent
Sample 9: 3 percent
Sample 10: 2 percent
Sample 11: 3 percent
```

Viewing Statistics for Forwarding Engine Utilization

To display the statistics for forwarding engine utilization, use the following command:

show utilization results forwarding-engine <*chip-fabport*> <*chip-index*>

You can enter this command from user mode.

Table 21-10. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*chip-fabport*>	The fabric port for which you want to view statistics.
<*chip-index*>	The forwarding chip for which you want to view statistics.

*** Note:** For an explanation of fabric ports and chip index numbers, see "Identifying the Ports," in Chapter 13, "Configuring Access Lists."

The command displays a history of utilization and the average utilization. A sample of the display is as follows:

```
Average statistics over the 120 second sample
window:
Forwarding Engine Utilization =1 percent.
Forwarding Engine Total Packet Rate = 28843 PPS.
Routing Component IPU Packet Rate   = 28835 PPS.
Routing Component IPM Packet Rate   = 0 PPS.
Routing Component IPX Packet Rate   = 0 PPS.

Statistic History (sorted from most recent to
oldest):

N     Utilizatio  Total PPS  IPU PPS  IPM PPS  IPX PPS
      n

0     1           28844      28836    0        0

1     1           28844      28836    0        0

2     1           28843      28836    0        0

3     1           28842      28834    0        0

4     1           28845      28837    0        0

5     1           28844      28836    0        0

6     1           28844      28836    0        0

7     1           28844      28836    0        0

8     1           28844      28836    0        0

9     1           28843      28836    0        0

10    1           28844      28836    0        0

11    1           28843      28836    0        0
```

22 Monitoring and Configuring the Forwarding Cache

Overview

Contents

The information and procedures provided in this chapter pertain to layer 3 module configuration only:

- Configuring the Forwarding Cache
- Monitoring the Forwarding Cache Statistics
- Displaying the Forwarding Cache

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

10-Gigabit Module and Forwarding Cache

Because of the volume of traffic that the 10-Gigabit module can process and its internal architecture, the module can cause the number of entries in the Layer 3 forwarding cache to reach the default maximum of 130,000 sooner than other modules. Once the maximum is reached:

- The switch routes all new traffic flows by means of the slow path.
- The speed at which traffic is forwarded is decreased.
- The role of master virtual router may be transferred from one VRRP router to another if VRRP is enabled.
- The following message is displayed in the CLI if debug mode is enabled and in the event log if logging of CLI debug messages is enabled:

  ```
  Log entry <number> by event 2 at <date and
  time> : Normal(0)

  Description: OS Error -- FILE: p2_ipenv.c
  LINE: 752 Value 0x2 : IP timer got skewed
  ```

  (Use the **logging protocol event cli debug** command to enable logging of CLI debug messages.)

This issue is most prevalent when the 10-Gigabit module as well as the rest switch is processing a high volume of layer 3 traffic flows.

To prevent this issue from occurring, monitor the forwarding cache and increase the maximum number of system entries as necessary. You may need to adjust the setting initially to accommodate peaks in network demand.

Configuring the Forwarding Cache

You can configure the forwarding cache using the Web Agent or the CLI.

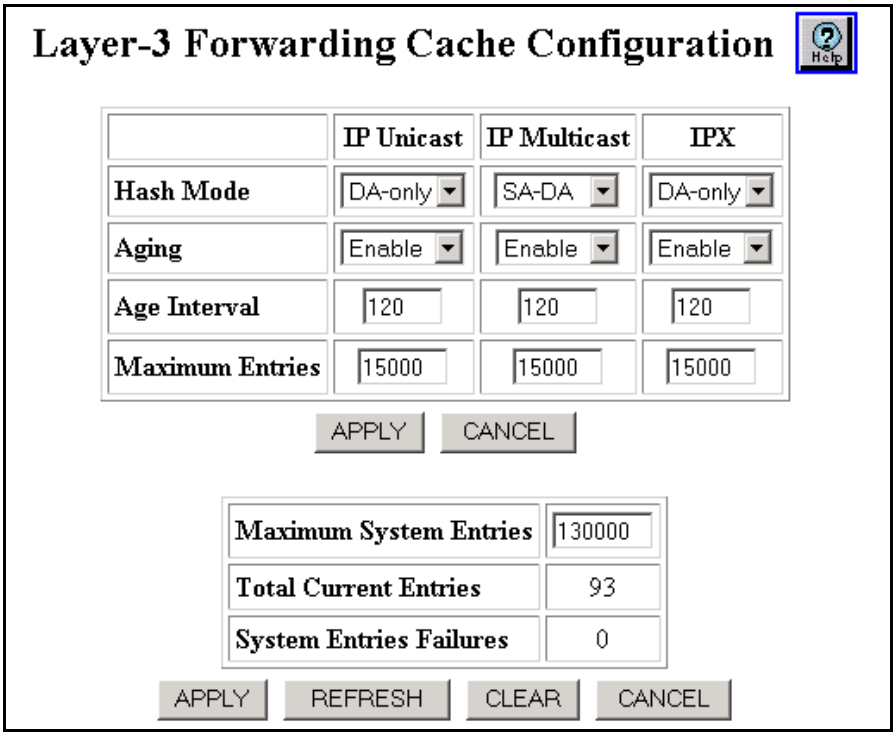
Web Agent Procedure

To configure the forwarding cache using the Web Agent:

1. In the navigation pane, expand the **Routing > L3 Forwarding Cache** folders, and then click **Cache Configuration**.

   The Layer 3 Forwarding Cache Configuration Web page is displayed in the content pane. See Figure 22-1.

Figure 22-1. Layer-3 Forwarding Cache Configuration Web Page



2. Configure the Layer 3 Forwarding Cache configuration parameters. See Table 22-1 for an explanation of the Layer-3 Forwarding Cache Configuration Web page parameters for tree configuration.

3. Click:

- **APPLY** to save your changes
- **CANCEL** to restore previous settings
- **REFRESH** to update your system configuration
- **CLEAR** to reset all configuration parameters to zero.

Table 22-1. Layer-3 Forwarding Cache Configuration Web Page Parameters for Tree Configuration

Field...	For the protocol...	Allows you to...
Hash Mode	IP Unicast (IPU) IP Multicast (IPM) IPX	Configure the table hash lookup mode for the IP unicast, IP multicast and IPX protocols. The options are: • **DA-only** - Forwarding entries input to the forwarding table are limited to using protocol destination address only. • **SA-DA** - Forwarding entries input to the forwarding table are limited to using destination and source address only.
Aging	IP Unicast IP Multicast IPX	Configure the IP unicast, IP multicast, or IPX protocols forwarding table entry aging status as either **Enable (default)** or **Disable**.
Age Interval	IP Unicast IP Multicast IPX	The IP unicast, IP multicast, or IPX protocols forwarding table entry aging period (in seconds). The valid range is **20-360** seconds. The default value is **120** seconds.
Maximum Entries	IP Unicast IP Multicast IPX	The number of active entries in the IP unicast, IP multicast, or IPX protocol forwarding table. This is the maximum number of active entries per fabric port. Additional flows are forwarded by the supervisor module. The default value is **15000**.

CLI Command

To configure the forwarding cache using the CLI, enter the following commands from Configure mode:

- `(configure)#` **ip multicast route-cache** <***options***>
- `(configure)#` **ip unicast route-cache** <***options***>
- `(configure)#` **ipx route-cache** <***options***>

*** Note:** There is no CLI command to clear the L3 FE cache.

Monitoring the Forwarding Cache Statistics

Monitoring the forwarding cache statistics includes:

- Displaying Frame Forwarding Statistics
- Displaying and Searching the L3 Forwarding Cache for an Entry

Displaying Frame Forwarding Statistics

The frame forwarding statistics indicate the performance of each of the multilayer media modules in respect to layer 3 routing and forwarding. You can display the frame forwarding statistics from either the Web Agent or the CLI.

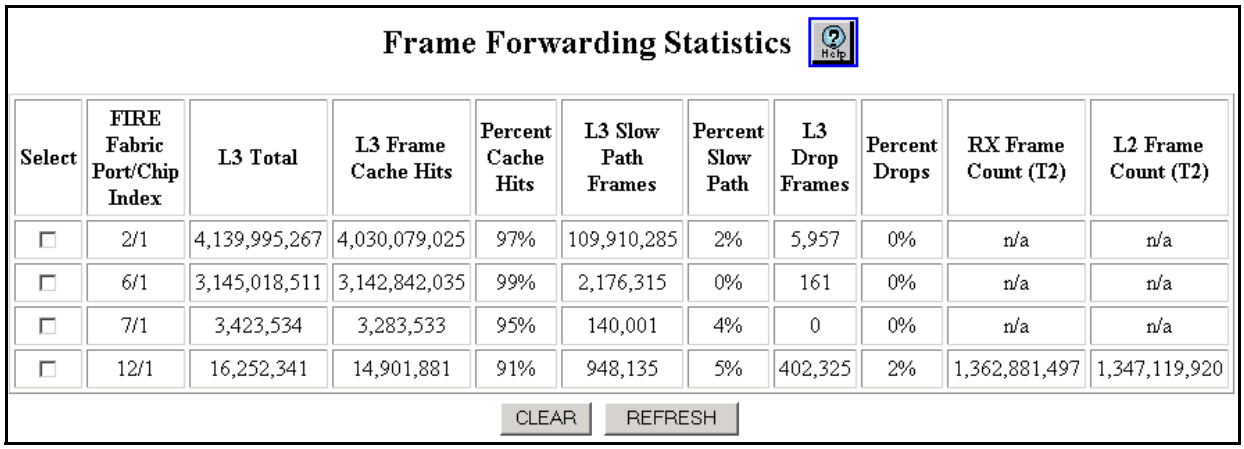
Web Agent Procedure

To display the frame forwarding statistics using the Web Agent:

1. In the navigation pane, expand the **Routing** > **L3 Forwarding Cache** folders, and then click **Forwarding Statistics**.

   The Frame Forwarding Statistics Web page is displayed in the content pane. See Figure 22-2.

Figure 22-2. Frame Forwarding Statistics Web Page



Frame Forwarding Statistics

Select	FIRE Fabric Port/Chip Index	L3 Total	L3 Frame Cache Hits	Percent Cache Hits	L3 Slow Path Frames	Percent Slow Path	L3 Drop Frames	Percent Drops	RX Frame Count (T2)	L2 Frame Count (T2)
☐	2/1	4,139,995,267	4,030,079,025	97%	109,910,285	2%	5,957	0%	n/a	n/a
☐	6/1	3,145,018,511	3,142,842,035	99%	2,176,315	0%	161	0%	n/a	n/a
☐	7/1	3,423,534	3,283,533	95%	140,001	4%	0	0%	n/a	n/a
☐	12/1	16,252,341	14,901,881	91%	948,135	5%	402,325	2%	1,362,881,497	1,347,119,920

CLEAR REFRESH

2. See Table 22-2 for an explanation for the Frame Forwarding Statistics Web page parameters.

Table 22-2. Frame Forwarding Statistics Web Page Parameters

Parameter	Defines the...
FIRE Fabric Port/ Chip Index	Fabric port associated with the layer 3 module. The P580 switch has 13 fabric ports: the P882 switch has 33 fabric ports. The supervisor module slot uses one fabric port: the media module slots each use two fabric ports. The supervisor module slot uses fabric port 1, the first media module slot uses fabric ports 2 and 3, and so on.
L3 Total (T2)	Total number of frames received on the fabric port.
L3 Frame Cache Hits	Number of packets received on the fabric port that were successfully matched against existing forwarding entries in the layer 3 (L3) address cache.
Percent Cache Hits	Total percentage of successful matches between packets received on a fabric port and the percent of those packets that matched the L3 address cache entries.
	1 of 2

Table 22-2. Frame Forwarding Statistics Web Page Parameters

Parameter	Defines the...
L3 Slow Path Frames	Number of frames received on a fabric port that were *not* successfully matched against existing forwarding entries in the layer 3 (L3) address cache. Subsequently, these frames were forwarded to the supervisor module (slow path). All frames sent to the supervisor module are routed in software. If this number is very large and the module associated with the fabric port is licensed, you may want to enable a short-lived protocol filter to conserve the forwarding engine resources. For information on how to enable a short-lived protocol filter, see "Configuring Short-Lived IP Protocol Filters" in Chapter 12, "Configuring IP Routing."
Percent Slow Path	Total percentage of unsuccessful matches between packets received on a fabric port and the percent of those packets that did not match the layer 3 address cache entries.
L3 Drop Frames	Number of layer 3 frames that were dropped. For example: They did not match the layer 3 address cache entries.
Percent Drops	Total percent of layer 3 frames dropped.
RX Frame Count (T2)	Total number of frames received on a fabric port.
L2 Frame Count (T2)	Number of layer 2 frames received on a fabric port that were forwarded on to an associated VLAN.
	2 of 2

Displaying and Searching the L3 Forwarding Cache for an Entry

Web Agent Procedure

You can display and search the Layer 3 Forwarding Cache for an entry using the Web Agent, To display the contents of the L3 forwarding Cache, you can use the search function to display:

- only the selected options,
- or all options by not selecting any options and clicking on the search button.

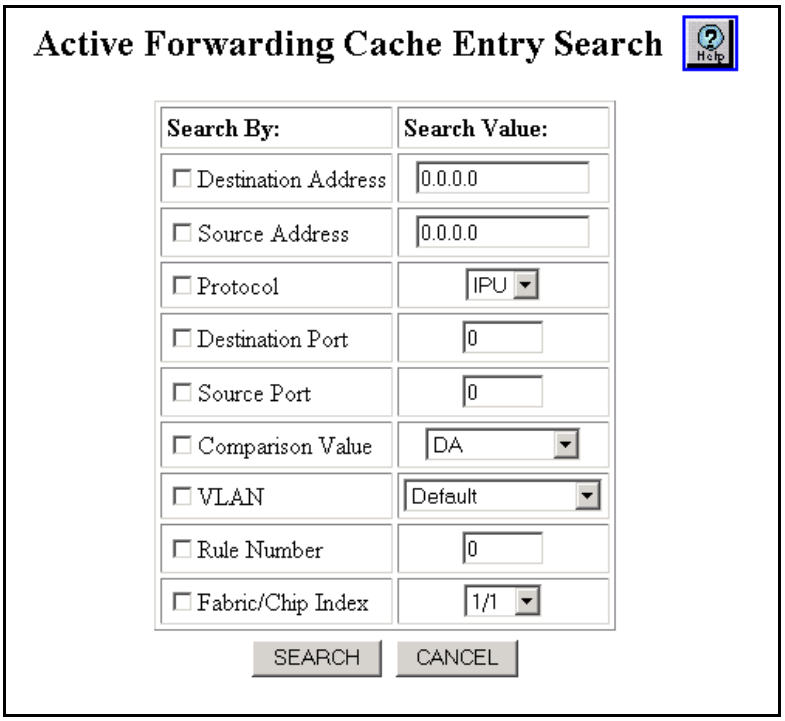
*** Note:** Layer 3 Active Forwarding Cache entry Search is available with the Web Agent only, no CLI version.

To search the L3 address cache using the Web Agent:

1. In the navigation pane, expand the **Routing > L3 Forwarding Cache** folders, and then click **Entry Search**.

   The Active Forwarding Cache Entry Search Web page is displayed in the content pane. See Figure 22-3.

Figure 22-3. Active Forwarding Cache Entry Search Web Page



2. Select the search criteria that you want to use to find more specific information on available routes from the Search By column.

   For example, if you want to find all entries in the Forwarding Cache by a VLAN, you would select the VLAN from the **Search By** column, and select the appropriate VLAN from the pull-down menu. Figure 22-4 is the actual display results from a Forwarding Cache search done by a VLAN.

   See Table 22-3 for an explanation of the Route Cache Entry Search Web page parameters:

Table 22-3. L3 Forwarding Cache Entry Search Web Page Parameters

Parameter	Allows you to search for all entries...
Destination Address	Within the routing cache that match the specified destination address (IP address), or IPX network number
Source Address	Within the routing cache that match the specified source address (IP address), or IPX network number
Protocol	By protocol identifier. By default, this field is 0 for IP.
Destination Port	Within the routing cache that match the specified destination port (IPU, IPM, IPX).
Source Port	Within the routing cache that match the specified source port (IPU, IPM, IPX).
Comparison Value	Within the routing cache that match the specified comparison value (TCP/UDP). Values include: • **DA** - Destination address • **DASA** - Destination and source address • **DAPROT** - Destination address and protocol • **DADP** - Destination address and destination port number • **DASAPROT** - Destination and source address and protocol • **DASADPSP** - Destination and source address, and the corresponding destination and source port numbers
VLAN	Within the routing cache that match the VLANs listed in the pull-down menu.
Rule Number	Within the routing cache that matches the specified rule number associated with an access list.
Fabric/Chip Index	Match the Forwarding Chip that is associated with the fabric port ID.

Figure 22-4. Forwarding Cache Search by VLAN

Active Forwarding Cache

Fabric/Chip Index	Tree	Access Rule	DA	SA	DP	SP	Comp	TTL	Age	Filter	Destination VLAN	Source VLAN	Mac Address	Use	Priority	Format
2/1	IP_UNI	None	10.0.1.22	0.0.0.0	0	0	DA	0	7	No	vlan10	n/a	02:e0:3b:e9:28:09	0	1	Eth2
4/1	IP_UNI	None	10.0.1.22	0.0.0.0	0	0	DA	0	7	No	vlan10	n/a	02:e0:3b:e9:28:09	0	1	Eth2
4/1	IP_MUL	None	239.255.0.1	10.0.4.8	0	0	DA	0	0	No	vlan10	vlan5	Derived from DA	1	1	Eth2
6/1	IP_UNI	None	10.0.1.22	0.0.0.0	0	0	DA	0	7	No	vlan10	n/a	02:e0:3b:e9:28:09	0	1	Eth2
7/1	IP_UNI	None	10.0.1.22	0.0.0.0	0	0	DA	0	7	No	vlan10	n/a	02:e0:3b:e9:28:09	0	1	Eth2

CLI Command

To display the L3 Forwarding Cache cache, use one of the following commands for a specific routing cache:

- > **show ip unicast cache**
- > **show ip multicast cache**
- > **show ipx cache**

Displaying the Forwarding Cache

You can display the Forwarding cache information using either the Web Agent or the CLI.

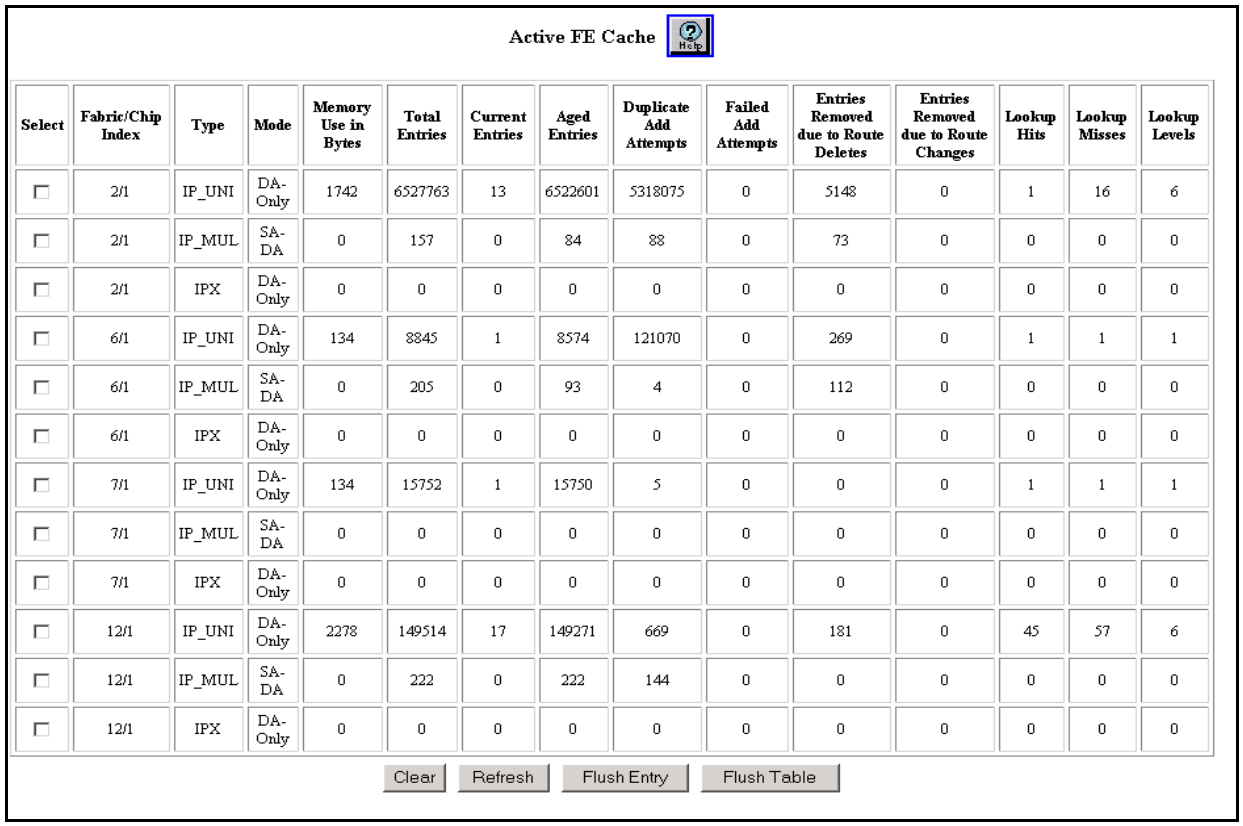
Web Agent Procedure

To display the forwarding cache information using the Web Agent:

1. In the navigation pane, expand the **Routing** > **L3 Forwarding Cache** folders, and then click **Cache Contents**.

   The Active FE Cache Web page is displayed in the content pane. See Figure 22-5.

Figure 22-5. Active FE Cache Web Page



Active FE Cache

Select	Fabric/Chip Index	Type	Mode	Memory Use in Bytes	Total Entries	Current Entries	Aged Entries	Duplicate Add Attempts	Failed Add Attempts	Entries Removed due to Route Deletes	Entries Removed due to Route Changes	Lookup Hits	Lookup Misses	Lookup Levels
☐	2/1	IP_UNI	DA-Only	1742	6527763	13	6522601	5318075	0	5148	0	1	16	6
☐	2/1	IP_MUL	SA-DA	0	157	0	84	88	0	73	0	0	0	0
☐	2/1	IPX	DA-Only	0	0	0	0	0	0	0	0	0	0	0
☐	6/1	IP_UNI	DA-Only	134	8845	1	8574	121070	0	269	0	1	1	1
☐	6/1	IP_MUL	SA-DA	0	205	0	93	4	0	112	0	0	0	0
☐	6/1	IPX	DA-Only	0	0	0	0	0	0	0	0	0	0	0
☐	7/1	IP_UNI	DA-Only	134	15752	1	15750	5	0	0	0	1	1	1
☐	7/1	IP_MUL	SA-DA	0	0	0	0	0	0	0	0	0	0	0
☐	7/1	IPX	DA-Only	0	0	0	0	0	0	0	0	0	0	0
☐	12/1	IP_UNI	DA-Only	2278	149514	17	149271	669	0	181	0	45	57	6
☐	12/1	IP_MUL	SA-DA	0	222	0	222	144	0	0	0	0	0	0
☐	12/1	IPX	DA-Only	0	0	0	0	0	0	0	0	0	0	0

Clear | Refresh | Flush Entry | Flush Table

2. To modify your FE Cache table, do one of the following:

   — Select the entry and click **Flush Entry** to delete one or more entries

   — Click **Flush Table** to clear the entire table

*** Note:** Flushing the table will disrupt traffic flow and clear all Forwarding Entries, and all flows must be relearned.

   — Click **REFRESH** to refresh the contents of the table

   — Click **Clear** to remove the statistics only (the entry in the table remains).

3. See Table 22-4 for an explanation of the Active FE Cache Web page fields.

Table 22-4. Active FE Cache Web Page Fields

Parameter	Defines...
Fabric/Chip Index	The Packet Routing Engine in question.
Type	The type of packet (for example, IP unicast, IP multicast).
Mode	The table mode which includes: • DA-Only • SA & DA The table mode indicates what values are used in a comparison to determine whether or not a packet is forwarded.
Memory Use in Bytes	Memory used by each entry.
Total Entries	The number of cumulative entries since the last time the statistics were cleared.
Current Entries	The number of active entries. If this number reaches the **Maximum Entries** setting for the protocol, you may want to enable a short-lived protocol filter to conserve the forwarding engine resources. For information on how to enable a short-lived protocol filter, see "Configuring Short-Lived IP Protocol Filters" in Chapter 12, "Configuring IP Routing."
Aged Entries	The number of entries aged out.
Duplicate Add Attempts	The number of attempts at adding the same slow path entries.
Failed Add Attempts	The number of failed attempts at adding a forwarding entry.
Entries Removed Due to Route Deletes	The number of entries removed because of route deletions.
Entries Removed Due to Route Changes	The number of entries removed because of route changes.
Lookup Hits	The number of incoming packets that was found in the cache.
Lookup Misses	The number of incoming packets that was not found in the cache.
Lookup Levels	Cache depth.

23 Using RMON and Ethernet Statistics to Analyze Network Performance

Overview

The following information and procedures provided in this chapter pertain to layer 2 and layer 3 module configurations:

- Viewing Network Statistics
- Setting Up Port Mirroring

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Viewing Network Statistics

You can view a variety of statistics from the switch interface that allows you to monitor network performance and troubleshoot network problems.

You can view statistics using either the Web Agent or the CLI.

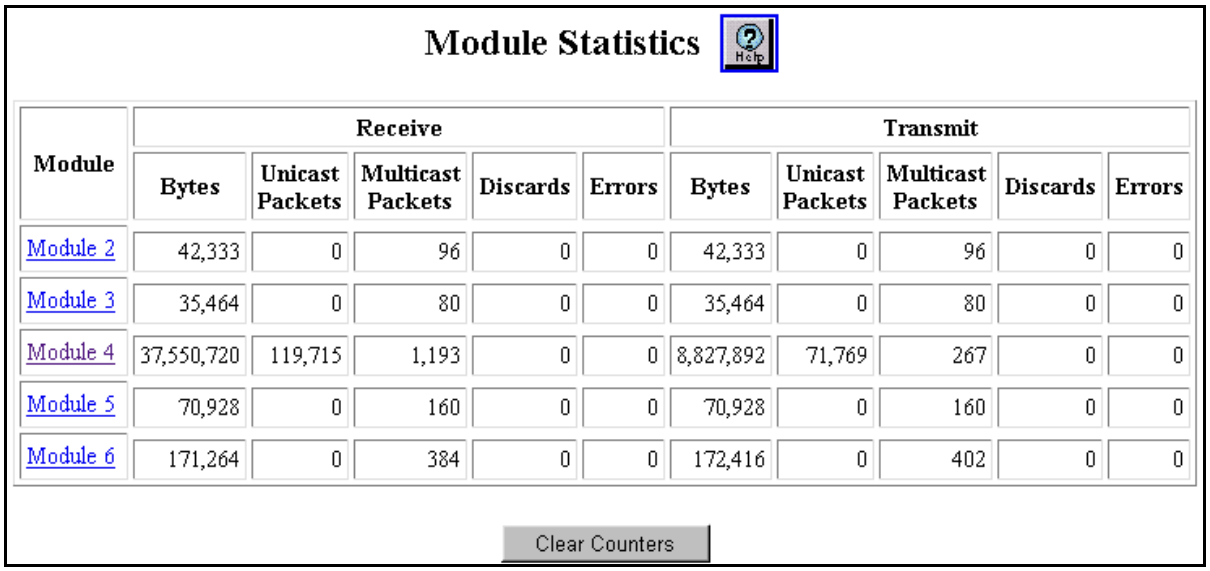
Web Agent Procedure

To view network statistics by using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Statistics**.

   The Module Statistics Web page is displayed in the content pane. See Figure 23-1.

Figure 23-1. Module Statistics Web page



Module Statistics

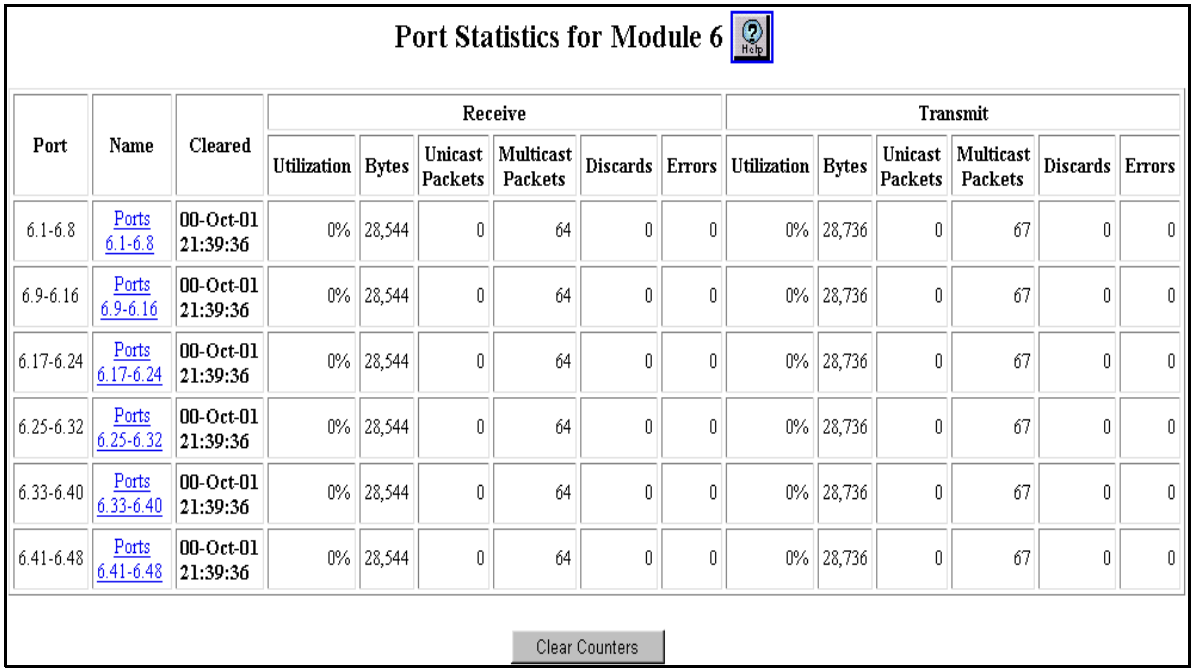
Module	Receive					Transmit				
	Bytes	Unicast Packets	Multicast Packets	Discards	Errors	Bytes	Unicast Packets	Multicast Packets	Discards	Errors
Module 2	42,333	0	96	0	0	42,333	0	96	0	0
Module 3	35,464	0	80	0	0	35,464	0	80	0	0
Module 4	37,550,720	119,715	1,193	0	0	8,827,892	71,769	267	0	0
Module 5	70,928	0	160	0	0	70,928	0	160	0	0
Module 6	171,264	0	384	0	0	172,416	0	402	0	0

Clear Counters

2. Select **Clear Counters** to get a fresh view of the statistics being gathered. This resets all of the counters to zero, so that you can track the counters from a specific point forward.

3. Select a module from the **Module** column to view statistics for that specific module. The Port Statistics Web page is displayed. See Figure 23-2.

Figure 23-2. Port Statistics Web page



Port Statistics for Module 6

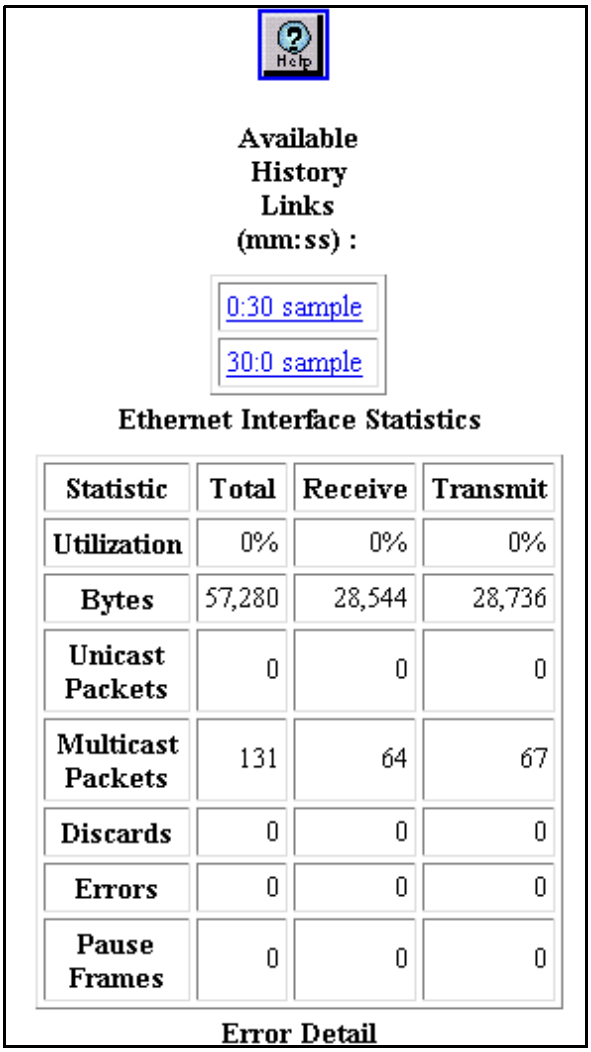
Port	Name	Cleared	Receive						Transmit					
			Utilization	Bytes	Unicast Packets	Multicast Packets	Discards	Errors	Utilization	Bytes	Unicast Packets	Multicast Packets	Discards	Errors
6.1-6.8	Ports 6.1-6.8	00-Oct-01 21:39:36	0%	28,544	0	64	0	0	0%	28,736	0	67	0	0
6.9-6.16	Ports 6.9-6.16	00-Oct-01 21:39:36	0%	28,544	0	64	0	0	0%	28,736	0	67	0	0
6.17-6.24	Ports 6.17-6.24	00-Oct-01 21:39:36	0%	28,544	0	64	0	0	0%	28,736	0	67	0	0
6.25-6.32	Ports 6.25-6.32	00-Oct-01 21:39:36	0%	28,544	0	64	0	0	0%	28,736	0	67	0	0
6.33-6.40	Ports 6.33-6.40	00-Oct-01 21:39:36	0%	28,544	0	64	0	0	0%	28,736	0	67	0	0
6.41-6.48	Ports 6.41-6.48	00-Oct-01 21:39:36	0%	28,544	0	64	0	0	0%	28,736	0	67	0	0

Clear Counters

4. Click **Clear Counters** to get a fresh view of the statistics being gathered. This resets all of the counters to zero so that you can track the counters from a particular point forward.

5. Select a port from the **Name** column to view statistics for that specific port. The Ethernet Interface Statistics Web page is displayed. See Figure 23-3.

Figure 23-3. Ethernet Interface Statistics Web page



6. Select either **0:30 sample** (30 second sample) or **30:0 sample** (30 minute sample) from the **Available History Links** field.

   The Ethernet Interface Statistics Web page is displayed with the sample you selected.

7. See Table 23-1 for an explanation of the Ethernet Interface Statistics Web page fields:

Table 23-1. Ethernet Interface Statistics Web Page Fields

Statistic	Indicates	Actions
Sample	The sample number.	N/A
Interval Start	The date and time this log entry was made.	N/A
Utilization	Percentage of utilization.	The percentage of available bandwidth used by traffic.
Bytes	Raw number of octets received at the interface. Provides some indication of the amount of network bandwidth being used.	A sharp increase could indicate a need to reconfigure the network.
Packets	Counts the raw number of readable Ethernet packets of legal length received at the interface.	A sharp increase could indicate a need to reconfigure the network. (However, octets are a better indication of bandwidth utilization.)
Broadcasts	Broadcast packets are a normal part of network operation. For example, IP networks use broadcasts as part of Address Resolution Protocol (ARP) to resolve network addresses.	Uses monitoring to recognize oncoming broadcast storms. Broadcast storms occur when stations are creating traffic that generates more traffic. **Possible cause:** Broadcasts cause every host on a network segment to process the packet. Possible actions: • To prevent broadcast storms, use VLANs to limit the area of the network that each broadcast packet affects. In general, each VLAN creates a separate broadcast domain. More VLANs mean less proliferation of broadcast packets. • Monitor the broadcast rate of your network during normal operation. • Establish a baseline. • Use Rate Limiting to reduce broadcasts.
		1 of 4

Table 23-1. Ethernet Interface Statistics Web Page Fields

Statistic	Indicates	Actions
Multicasts	Normal during network operation. For example, multicast packets are to send target video streams to selected stations on the network, and are part of the operation of the Spanning Tree Protocol.	**Possible causes:** • Too many multicast frames can consume valuable network bandwidth. **Possible actions:** • Using Intelligent Multicasting can significantly reduce multicast traffic on individual ports. • Segmenting the network into smaller VLANs and routing between them can also help control proliferation of multicasts.
CRC (Cyclic Redundancy Check) or Alignment Errors	Counts of the number of times that the number of bits in a frame cannot be divided by 8 (that is, cannot be broken into legal octets), and that contain a Frame Check Sequence validation error. Typically caused by turning equipment on or off, and by noise on twisted pair segments. These errors can also result from configuring a network that does not comply with 802.3 standards. In a standards-compliant Ethernet network, CRC or alignment errors represent transit and receive bit errors. The Ethernet standard allows 1 in 10^8 bit error rate, but you should expect performance to be less than 1 in 10^{12} packets. Rates in excess of one error per one thousand packets indicate a serious problem.	**Possible causes**: • Defect at the transmitting station. • Turning equipment on or off. This should cause only a few errors. • Damaged cables. • Interference on network cabling. **Possible action**s (respectively): • Use port error statistics to isolate the problem. Check the transceiver or adapter card connected to the port where the problem seems to originate. Also check the cable and cable connections for damage. • Normal operation, no action required. • Check cables for damage. • Inspect cable runs to see if they are too close to noisy devices, and check for problems with network devices.
Undersized Packets	Count of packets with a valid CRC that violate the minimum Ethernet packet size. These malformed packets are most often the result of software errors.	**Possible cause:** Device or application creating non-compliant packets. **Possible action:** Use a network analyzer to identify the which transceiver which is at the source of the problem. Replace the transceiver, network adapter, or station.
		2 of 4

Table 23-1. Ethernet Interface Statistics Web Page Fields

Statistic	Indicates	Actions
Oversized Packets	Count of packets with a valid CRC that violate the maximum Ethernet packet size. These malformed packets are most often the result of software errors.	**Possible cause:** Device or application creating non-compliant packets. **Possible action:** Use a network analyzer to identify the transceiver which at the source of the problem. Replace transceiver, network adapter, or station.
Fragments	Fragments or runts result from normal collision activity in Ethernet networks. A runt packet is an incomplete packet that is long enough to be detected by an Ethernet interface.	**Possible causes**: • Interference on network cabling. • A Transceiver attached to the Repeater is generating Signal Quality Errors (SQE). **Possible actions (respectively)**: • Inspect cable runs to see if they are too close to noisy devices, and check for problems with network devices. • Disable SQE on the Transceiver.
Jabbers	Jabbers indicate that devices on the networks are sending improper electrical signals. Because Ethernet uses electrical signalling to determine whether or not it is okay to transmit, a jabber condition can halt all traffic on a segment. Jabbers do not occur on fiber optic cable and thus do not occur on the 10-Gigabit module.	**Possible causes**: • Bad network interface card • Repeater network with looped traffic **Possible actions (respectively):** • Replace the network interface card. • Rewire network to remove the loop.
		3 of 4

Table 23-1. Ethernet Interface Statistics Web Page Fields

Statistic	Indicates	Actions
Collisions (half-duplex links only)	Counts number of times that packets have collided on the network. Collisions increase as network use of shared segments increases. Therefore, if the collision rate increases without an increase of network use, it might indicate a problem. Guidelines for appropriate collision rates are: • 10 percent: Normal collision rate for shared Ethernet segment. • 30 percent: Collisions begin to interfere with performance. • 70 percent: Practical limit for network to remain functioning. A full-duplex link should not show collision activity. Collisions are rare in a switched network, unless your switched segments attach to multiple ends stations (a legal configuration option). Collisions do not occur on fiber optic cable and thus do not occur on the 10-Gigabit module.	**Possible causes**: • Busy network • Broken adapter (not listening before broadcasting) • Network loop **Possible actions (respectively)**: • If you have multiple stations on a switch segment, reconfigure network into segments with fewer stations. • Isolate each adapter to see if the problem ceases. • Activate spanning tree to resolve loops automatically. • Ensure that there are no connections to the same station where both connections are simultaneously active.
		4 of 4

CLI Command

To view network statistics using the CLI, enter the following command in Privileged mode:

`(configure)#` **show ethernet counters <cr | mod-num | mod-swport-spec>**

Setting Up Port Mirroring

Configuring an RMON mirror port allows you to mirror traffic from a port or set of ports to a specific mirror port, where you can attach a sniffer or RMON probe. The switch supports a single mirror port and a single source port for each switch fabric port. For example, 20-port Fast Ethernet cards have two fabric ports (one for ports 1 through 10, one for ports 11 through 20). You can set up a single source port and a single mirror port for each set of ports associated with a fabric port. You can also choose to mirror all traffic from a particular fabric port to the mirror port, or set up multiple source ports to mirror traffic to a single mirror port.

Packets addressed to the CPU, such as pings, are duplicated out of the mirror port. Tagged packets that are sent into a source port with a VLAN ID to which the source port is not bound, are not transmitted out the mirror port. VLAN tag information is not propagated to the mirror port.

*** Note:** To prevent unnecessary traffic flooding on a mirror port, put the mirror port on the same VLAN as the source port.

Two port mirroring features are available for the Avaya Multiservice switches: Fabric mode 1 port mirroring and Fabric mode 2 port mirroring. The switch must be in Fabric mode 2 to perform Fabric mode 2 port mirroring.

Fabric mode 1 port mirroring works with 50-series modules. Fabric mode 2 port mirroring works with 80-series modules.

*** Note:** If you use port mirroring, the **Frame Tags** field on the Switch Port Configuration Web page for the port must be set to **Use** (default).

This section contains the following procedures:

- Setting Up a Port Mirror on a Switch in Fabric Mode 1
- Setting Up Port Mirroring on a Switch in Fabric Mode 2
- Removing a Fabric Mode 2 Port Mirror
- Viewing Information about a Fabric Mode 2 Port Mirror

Setting Up a Port Mirror on a Switch in Fabric Mode 1

You can set up a port mirror on a switch in Fabric Mode 1 by using either the Web Agent, the CLI, or SNMP. Fabric mode 1 port mirroring works with 50-series modules.

Restrictions

- Do not use a router port as a piggyback port.
- Do not set port mirroring on disabled ports.
- You must establish a link to a mirror port before you setup a port mirror with a piggyback port. Otherwise, the port mirror will not work.

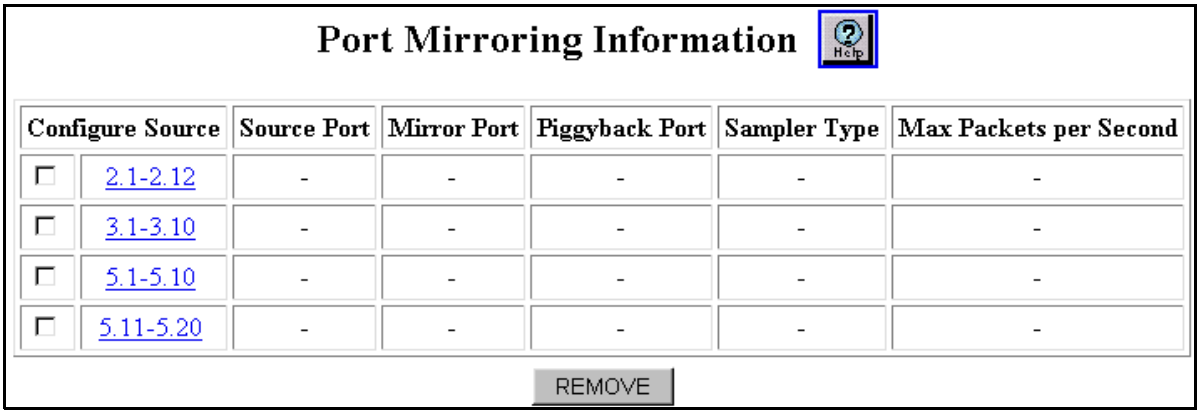
Web Agent Procedure

To set up a port mirror on a switch in Fabric mode 1 by using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Port Mirroring**.

   The Port Mirroring Information Web page is displayed in the content pane. See Figure 23-4.

Figure 23-4. Port Mirroring Information Web Page



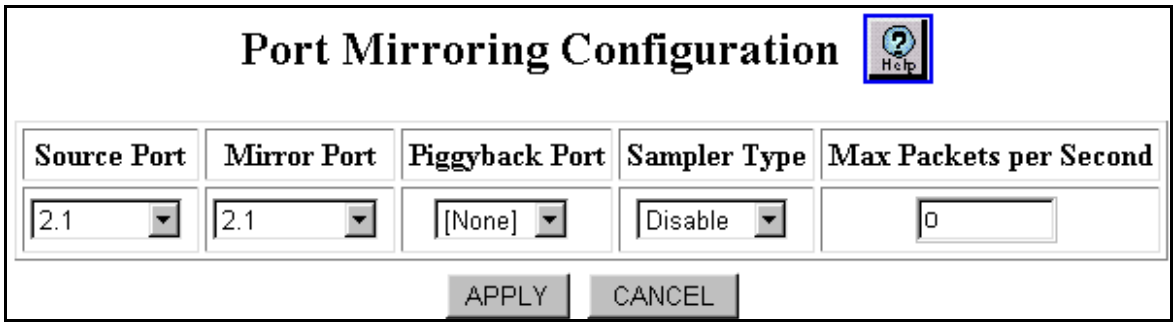
2. See Table 23-2 for an explanation of the fields on the Port Mirroring Information Web page.

Table 23-2. Port Mirroring Information Web Page Parameters

Parameter	Definition
Configure Source	Select the configuration source port. Provides a link to the Port Mirroring Configuration Web page.
Source Port	Displays the port under investigation.
Mirror Port	Displays the port that transmits mirrored data.
Piggyback Port	Displays the name of the port that is used for bi-directional port mirroring. When used as a mirror port, it is unavailable for other uses. **Note:** 80-series modules do not support piggyback ports.
Sampler Type	Displays the speed of sampling that is performed for source port traffic.
Max Packets per Second	Displays the maximum number of packets per second that are served by the mirror port.

3. Select a source port for traffic from the **Configure Source** column. The Port Mirroring Configuration Web page is displayed. See Figure 23-5.

Figure 23-5. Port Mirroring Configuration Web page



4. See Table 23-3 for an explanation of the fields on the Port Mirroring Configuration Web page.

Table 23-3. Port Mirroring Configuration Web Page Parameters

Parameter	Definition
Source Port(s)	List of available selections. You can select a particular source port associated with the selected fabric port. You can also select all ports. **Note:** To mirror inbound traffic only, select a source port and a mirror port, not a piggyback port.
Mirror Port	Displays the port from which you want to send the traffic. This port can be on another module in the switch. Once a specific port associated with a fabric port has been designated a mirror port, other ports associated with that fabric port no longer appear on the selection list. **Note:** The source port and the mirror port must be different physical ports.
	1 of 2

Table 23-3. Port Mirroring Configuration Web Page Parameters

Parameter	Definition
Piggyback Port	Port used to enable bi-directional port mirroring. If no piggyback port is specified, only received traffic from the source port will go to the mirror port. The piggyback port should have the same bandwidth as the source port. Only one port per fabric can be used as a piggyback port. Sampling rates have no effect on transmitted information. **Note:** To mirror outbound traffic only, select a source, mirror, and piggyback port. Set the sampling to disable. Since disabling sampling only applies to inbound traffic, only outbound traffic is received. **Considerations:** • You cannot use a port that has been designated as a piggyback port. • Once a port has been designated as a piggyback port, the link light is turned on, even if there is no connection on the selected port. • The piggyback port is displayed in place of the source port in the VLAN menu. • If the source and piggyback ports are at a higher bandwidth than the mirror port's bandwidth, the traffic on the source port may exceed the bandwidth that the mirror port can handle. • The piggyback port will always show that it is using multi-layer tagging. • A spanning tree topology change occurs when you change a piggyback port. • 80-series modules do not support piggyback ports. • Fabric mode 1 port mirroring works with 50-series modules. Fabric mode 2 port mirroring works with 80-series modules. • If you set up a piggyback port, multicast traffic that is routed through the mirrored port may be lost for a few minutes. This loss of multicast traffic is due to the DVMRP routes that pass through the mirrored port entering hold-down state. The default setting for the DVMRP hold-down period is 120 seconds. Once this time expires, multicast traffic resumes. • If you set up a port mirror with a piggyback port and the source port is a client port for an intelligent multicast session, the port is removed from the multicast session. Hosts that are attached to the client port stop receiving multicast traffic until the port is either statically or dynamically added to the multicast session again. To avoid this interruption in multicast traffic, after setting up the port mirror: • If the multicast session is a *static* session, add the client port to the session again. • If the multicast session is a *dynamic* session, the client port is added to the session again when the switch receives the next IGMP membership report.
Sampler Type*	Selects how often you want the mirror port to receive traffic samples: • **Always** - sends all samples. • **Periodic** - sends samples at the interval described below. • **Disabled** - shuts off traffic samples to the mirror port, but keeps the association intact.
Max Packets per Second	Displays the number of packets per second that are served by the mirror port.

* Sampling only applies to inbound traffic.

2 of 2

CLI Command

To set up an RMON mirror port by using the CLI, enter the following command from **Enable/Configure** mode:

```
(configure)# set port mirror <mod-port-range>
```

Setting Up Port Mirroring on a Switch in Fabric Mode 2

Fabric mode 2 port mirroring provides the ability for 80-series modules to monitor traffic that a port transmits and receives. Four channels are provided for monitoring traffic, so you can monitor four different ports or port ranges simultaneously. You can set up Fabric mode 2 port mirrors by using either the Web Agent or the CLI. The switch must be in Fabric mode 2 to perform Fabric mode 2 port mirroring.

Restrictions

- If you attempt to set up port mirroring on a port that is administratively disabled, and you configure the disabled port as a source or mirror port, note the following:

  Traffic sent to a disabled source port will be lost. A disabled port set up as a mirror port will not be able to send any traffic to a monitoring device.

  Avaya recommends that you check any ports to be used in port mirroring to ensure that the ports are active.

- When using Fabric mode 2 port mirroring, you cannot mirror transmit traffic of multiple source ports to one mirror port.

- When you use Fabric mode 2 port mirroring, both the source port and mirror port must either:

  — Be on the same VLAN and have the same VLAN binding

  or

  — Have VLAN binding set to **bind to all**

- Avaya recommends that you do *not* set up a port mirror with a source port or mirror port that is bandwidth limited. This is because packets are dropped during the bandwidth limiting process. Since both source and mirror packets are subject to drop, and because packets are dropped at random, if packets are dropped from either the source or the mirror traffic (or both), there is no guarantee that the mirror traffic will match the source traffic.

- Fabric mode 2 port mirroring is not supported for hunt group ports.

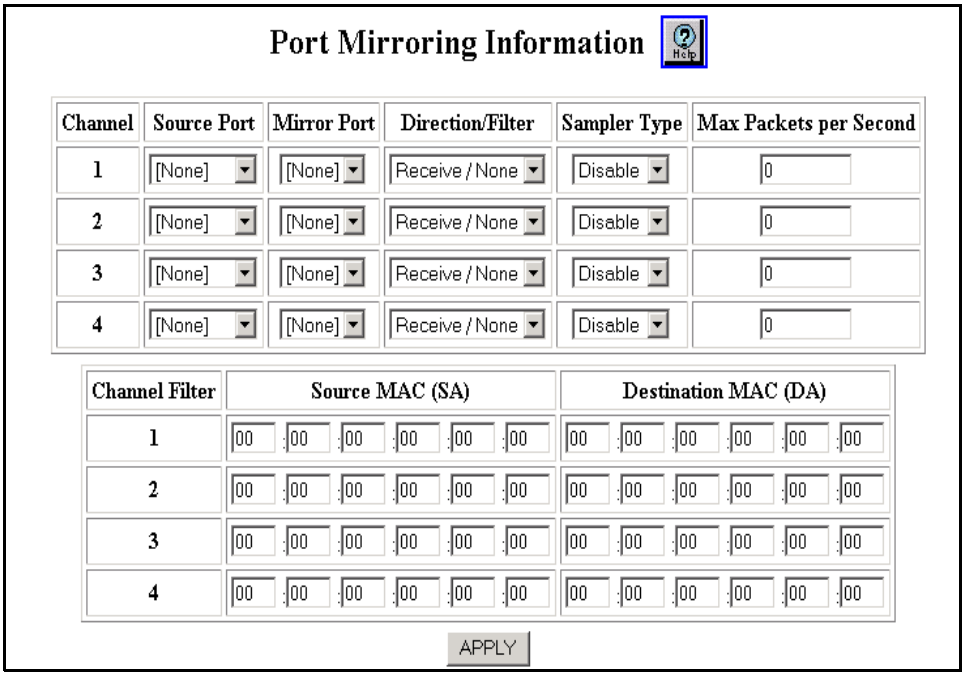
Web Agent Procedure

To use the Web Agent to set up port mirroring on a switch in Fabric mode 2:

1. Expand the **Modules & Ports** folder.
2. Click **Port Mirroring**.

   The Port Mirroring Information Web page is displayed in the content pane. See Figure 23-6. Table 23-5 provides an explanation of each field on the Web page.

Figure 23-6. Port Mirroring Information Web Page



   See Table 23-4, for the specific port ranges that you can mirror on each type of module.

Table 23-4. Port Ranges for Fabric Mode 2 Port Mirroring

Module	Port ranges that you can mirror
4-port gigabit modules	• 1–2 • 3–4 • Any single port You can mirror any four single ports simultaneously (one port per channel). However you cannot mirror a port range and a single port within that range simultaneously. **Example:** You *can* mirror port 1 on channel 1, port 2 on channel 2, and port 3 on channel 3 simultaneously. However, you *cannot* mirror ports 1 through 2 on channel 1 and port 2 on channel 2 simultaneously.
8-port gigabit modules	• 1–4 • 5–8 • Any single port You can mirror any four single ports simultaneously (one port per channel). However you cannot mirror a port range and a single port within that range simultaneously. **Example:** You *can* mirror port 1 on channel 1, port 2 on channel 2, and port 3 on channel 3 simultaneously. However, you *cannot* mirror ports 1 through 4 on channel 1 and port 2 on channel 2 simultaneously.
24-port 10/100 modules	• 1–12 — any 1 port or the entire range. • 13–24 — any 1 port or the entire range. If you mirror a single port, you can mirror only 1 port per range at a time. **Example:** You *can* mirror port 1 on channel 1 and port 13 on channel 2 simultaneously. However, you *cannot* mirror port 1 on channel 1 and port 2 on channel 2 simultaneously.
	1 of 2

Table 23-4. Port Ranges for Fabric Mode 2 Port Mirroring

Module	Port ranges that you can mirror
48-port 10/100 modules	• 1–12 — any 1 port or the entire range. • 13–24 — any 1 port or the entire range. • 25–36 — any 1 port or the entire range. • 37–48 — any 1 port or the entire range. If you mirror a single port, you can mirror only 1 port per range at a time. **Example:** You *can* mirror port 1 on channel 1 and port 13 on channel 2 simultaneously. However, you *cannot* mirror port 1 on channel 1 and port 2 on channel 2 simultaneously.
	2 of 2

3. Select the Channel and the associated port from the **Source Port** column pull-down menu.

4. In the **Mirror Port** column pull-down menu, select the port to which you want to mirror traffic.

Both the source port and mirror port must either:

— Be on the same VLAN and have the same VLAN binding

or

— Have VLAN binding set to **bind to all**

***Note:** The source port and the mirror port must be different physical ports.

5. In the **Direction/Filter** column, select the direction of traffic that you want to mirror and the filter that you want to set up, if any.

You can set up a MAC address filter to monitor only traffic with a specific source MAC address or destination MAC address.

*** Note:** Do not set the source MAC address in the **Destination/Filter** field.Setting a Source MAC address in the Destination MAC filter causes traffic to be improperly monitored.

6. In the **Sampler Type** column, select how often you want the mirror port to receive traffic samples.

7. Do one of the following:

IF you select...	THEN...
Periodic in the **Sampler Type** column	In the **Max Packets per Second** column, enter the maximum number of packets per second that you want the mirror port to receive.
Disable or **Always** in the **Sampler Type** column	Go to step 8.

8. If you selected **Transmit/DA Filter** in the **Direction/Filter** field, enter the MAC address that you want to monitor in the **DA Filter** field for the port mirror.

9. If you selected **Receive/SA Filter** in the **Direction/Filter** field, enter the MAC address that you want to monitor in the **SA Filter** field for the port mirror.

10. Click **Apply**.

Table 23-5. Port Mirroring Information Web Page Parameters

Parameter	Definition
Channel	Displays the channel number. Four channels are provided for monitoring traffic, so you can monitor four different ports or port ranges simultaneously.
Source Port	The port that you want to mirror.
Mirror Port	The port to which you want to mirror traffic. Both the source port and mirror port must either: • Be on the same VLAN and have the same VLAN binding **or** • Have VLAN binding set to **bind to all**
	1 of 2

Table 23-5. Port Mirroring Information Web Page Parameters

Parameter	Definition
Direction/Filter	The direction of traffic that you want to mirror and the filter that you want to set up, if any. Options are: • **Receive/None**-Mirrors all receive traffic. • **Transmit/None**-Mirrors all transmit traffic. • **Both/None**-Mirrors all transmit and receive traffic. • **Receive/SA Filter**-Mirrors receive traffic that has the source MAC address that you specify in the **Source MAC (SA)** field. • **Transmit/DA Filter**-Mirrors transmit traffic that has the destination MAC address that you specify in the **Destination MAC (DA)** field.
Sampler Type	Specifies how often you want the mirror port to receive traffic samples. Options are: • Disable • Always • Periodic
Max Packets per Second	The maximum number of packets per second that you want the mirror port to receive. Enter a number in this field *only* if you selected **Periodic** in the **Sampler Type** field.
Channel Filter	Displays the channel number.
Source MAC (SA)	The source MAC address that you want to mirror traffic for. Enter a MAC address in this field *only* if you selected **Receive/SA Filter** in the **Direction/Filter** field.
Destination MAC (DA)	The destination MAC address that you want to mirror traffic for. Enter a MAC address in this field *only* if you selected **Transmit/DA Filter** in the **Direction/Filter** field.
	2 of 2

CLI Command

Use the following CLI commands to configure Fabric mode 2 port mirroring. You must be in Global Configuration command mode to enter these commands.

To set up port mirroring, use the following CLI command:

`(configure)#` **set port mirror Fabric_mode2 source-port <*mod-port-range*> mirror-port <*mod-port-spec*> channel <*channel*> direction {tx | rx | both | sa | da} sampling {always | disable | periodic} [sa <*MAC-address*>] [da <*MAC-address*>] [max-packets-sec <*max-packets-sec-value*>]**

Removing a Fabric Mode 2 Port Mirror

Web Agent Procedure

To remove a port mirror on a switch in Fabric mode 2 by using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Port Mirroring**.

   The Port Mirroring Information page is displayed in the content pane.

2. In the **Source Port** column, select **None** for the port mirror that you want to remove.

3. Click **Apply**.

CLI Command

To remove a port mirror from a switch in Fabric mode 2, use the following CLI command from **configure** mode:

```
(configure)# clear port mirror Fabric_mode2 channel <channel>
```

Viewing Information about a Fabric Mode 2 Port Mirror

Web Agent Procedure

To view information about a Fabric mode 2 port mirror:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Port Mirroring**.

   The Port Mirroring Information page is displayed in the content pane. This page displays information about all port mirrors that are setup.

CLI Command

To view information about a Fabric mode 2 port mirror, use the following CLI command:

```
(configure)# show port mirror Fabric_mode2
```

This command displays the source ports, mirror port, direction being mirrored, sampler type, and maximum packet per second for all port mirrors that are currently set up.

24 Managing Buffers and Queues on 50-Series Modules

Overview

The following information and procedures provided in this chapter pertain to layer 2 and layer 3 module configurations. Buffer management features help you to optimize traffic throughput through the switch fabric.

This chapter provides the following information:

- How Queues Work
- Managing Buffers and Queues

For more information about the CLI commands that are mentioned in this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

How Queues Work

Frames are buffered in the I/O modules, before and after traversing the switch. Each queue can hold 256K bytes. (Architecturally they can support up to 1 MB each).

Each buffer is divided into two queues, one for High-priority Traffic and one for Normal-priority Traffic. The factory default is for the high-priority queue uses 20% (51K) of the buffer. The normal-priority queue uses the remaining 80% (205K). These values can be modified using either the Web Agent or SNMP.

*** Note:** When you change these values, you must reboot the switch before they can take effect.

Less buffer memory gets assigned to the high-priority queue because the high-priority queue gets serviced more frequently than the normal-priority queue. Since a frame spends less time on the high-priority queue, less buffer space is required for the queue.

The Service Ratio can be chosen to match traffic patterns and performance requirements using a weighted round robin scheduling algorithm. The available service ratios of the algorithm are defined in "Managing Buffers and Queues". The factory default service ratio for fabric ports is 999/1. The factory default service ratio for physical ports is 1023 to 1. If there is traffic to be serviced from both the high- and normal-priority queues, 999 packets of high-priority traffic will be processed for each normal-priority packet.

When the high-priority queue fills up, incoming frames are dropped. The philosophy is if a high-priority frame is going to be late, it is not worth sending it at all. The normal-priority queue uses either IEEE 802.3X PAUSE (variable timed XOFF) flow control or Half Duplex collisions to shut off incoming frames before the queue overflows.

The switch implements two flow control disciplines along the entire path that frames travel. The default case is that when output buffers fill up, frames destined for a particular buffer will be dropped. This should only occur in a case where the output port is very congested. However, there is an optional mode in which normal-priority frames are never dropped inside the switch. In this mode, input buffers may fill up. If they do, the affected input ports may use flow control to temporarily halt traffic from neighboring switches.

*** Note:** 802.1p packets that are received with a tag priority of 0 on a 50-series layer 2 (non-routing) module, and that must be routed via the FORE path on an 80-series supervisor, are queued and transmitted with a priority of 4. This priority change is due to the conversion from the high-low priority system that 50-series modules use to the 8-level priority system that 80-series modules use.

Managing Buffers and Queues

You can manage buffers and queues using either the Web Agent or the CLI.

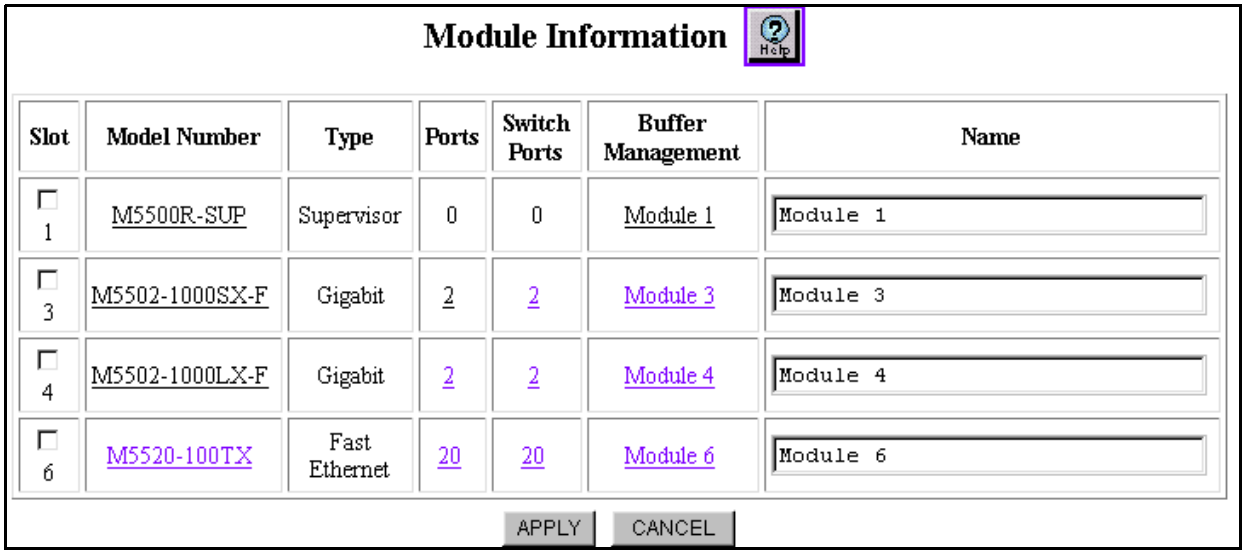
Web Agent Procedure

To manage buffers and queues using the Web Agent:

1. In the navigation pane, expand the **Modules & Ports** folder, and then click **Configuration**.

   The Module Information Web page is displayed in the content pane. See Figure 24-1.

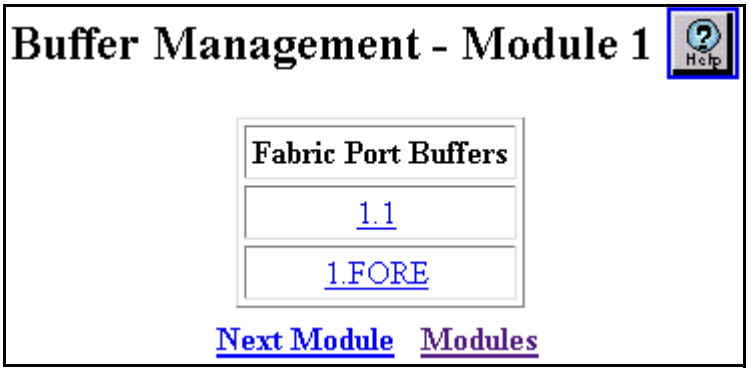
Figure 24-1. Module Information Web Page



2. Select the module whose buffers you want to manage from the **Select** column.

3. Select the Module number for that module from the Buffer Management column. The Buffer Management Web page is displayed. See Figure 24-2.

Figure 24-2. Buffer Management Web Page



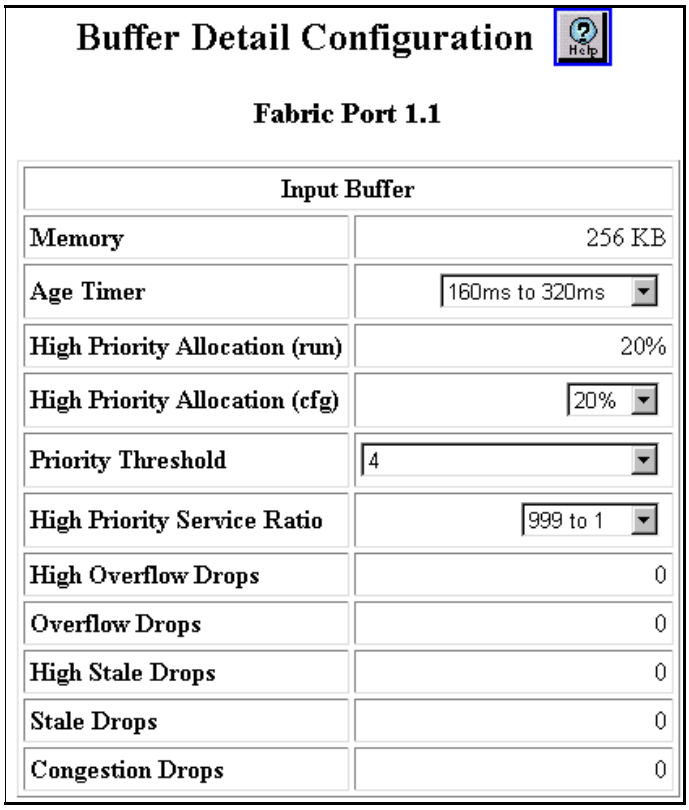
4. See Table 24-1 for an explanation of the Buffer Management Web page parameters.

Table 24-1. Buffer Management Web Page Parameters

Parameter	Definition
Fabric Port Buffers	Displays the port's fabric port buffers and allows you to open the Buffer Detail Configuration Web page for the selected module. Service ratios: • 3 to 1 • 99 to 1 • 999 to 1 • 9999 to 1
Physical Port Buffers	Displays the port's physical port buffers if available. The buffer ratios are: • 31 to 1 • 63 to 1 • 127 to 1 • 255 to 1 • 511 to 1 • 1023 to 1 • 2047 to 1 • 4095 to 1 • 8191 to 1 • 16383 to 1 • 32767 to 1

5. Select the Fabric Port Buffer number whose associated buffers you want to manage. The Buffer Detail Configuration Web page for that fabric port is displayed. See Figure 24-3.

Figure 24-3. Buffer Detail Configuration Web Page



6. See Table 24-2 for an explanation of the Buffer Detail Configuration Web page input and output fields:

Table 24-2. Buffer Detail Configuration Web Page Parameters

Parameter	Definition...
Memory	Displays the amount of physical memory associated with this buffer.
Age Timer	Displays the amount of time a packet remains in the queue before being discarded as a stale packet. You may want to increase the timer value for ports connected to 10 MB/s ports, particularly 10 MB/s shared media, because you may want to queue packets longer before discarding them.
	1 of 2

Table 24-2. Buffer Detail Configuration Web Page Parameters

Parameter	Definition...
High Priority Allocation	Displays the percent of the buffer's queuing space allotted to high priority traffic. Because the high-priority queue is serviced more frequently than the normal priority queue, raising this value may not necessarily provide better service. In fact, if you are using the high-priority queue for delay-sensitive traffic, you may want to reduce the amount of memory devoted to the high-priority queue. This ensures that packets that cannot be delivered in a timely manner are discarded. If you want the high priority queue to guarantee delivery of as many packets as possible, regardless of delay, increase this value. The change does not take effect until you reset the switch.
Priority Threshold	Allows you to set this parameter to the value at which the switch starts sending packets to the high-priority queue. The default value (4) causes all traffic with a priority greater than or equal to 4 (4, 5, 6, and 7) to be assigned to the high-priority queue. Priority schemes have more than two queues (the IEEE allows up to 8, numbered 0 through 7). Avaya recommends that you do not change this parameter.
High Priority Service Ratio	Allows you to set how many times the high priority queue is serviced for each time the low priority queue is serviced. The ideal value changes from queue to queue, but the goal is to ensure that traffic mix guarantees optimal mix between high-priority and best effort traffic.
High and Normal Overflow Drops	Displays the number of packets dropped because the associated buffer is full. Indicates that the device immediately before the queue is processing traffic faster than the next downstream element can process the same volume of traffic. For example, overflow drops on the input buffer indicate that traffic is arriving faster than the switch matrix can process it. Overflow drops on the output buffers indicates that the output port cannot handle the volume of the load being offered.
High and Normal Stale Drops	Displays the number of packets dropped because they timed out waiting for service (using the age timer value). In the high-priority queue, this can help determine how efficiently the switch is processing "better never than late" traffic. Excessive stale drops on the high-priority queue may indicate the need to increase the service ratio on the high-priority queue.
Congestion Drops	Displays the number of packets dropped because the switch controller has sensed congestion at the outbound port.
	2 of 2

7. Repeat Steps 1-4 to tune Physical Port (Fast Ethernet) buffers. Physical Port ports have additional buffers on both the input and output ports.

8. Click **APPLY** to save your changes, or **CANCEL** to restore previous settings.

CLI Command

To manage buffers and queues using the CLI, enter the following command from **Configure** mode:

`(configure)#` **set buffering port** <***mod-swport-spec***>

25 80-Series QoS

Overview

Quality of Service (QoS) is a set of tools that make it possible for you to manage traffic across a switch or a network. These tools protect specific traffic from the effects of network congestion. You can enable these features for a module or individual ports on a module to control the flow of traffic across your network. This control makes it possible for you to guarantee that delay-sensitive traffic such as voice over IP (VoIP) receives the priority it requires, while also ensuring that the switch services other low priority data.

The QoS features are supported only on 80-series modules. Fifty-series modules do not support these features.

This chapter contains the following sections:

- Why implement QoS?
- How Does QoS Work?
- Classification of Traffic
- Ingress Policing
- Queue-Servicing Algorithms
- QoS Statistics

*** Note:** In this release, you can configure the QoS features only by using the CLI or Avaya Policy Manager Version 2.2. These features are not available in the Web Agent or SNMP.

For more information about the CLI commands that are discussed this chapter, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Why implement QoS?

Purpose of QoS

In a network that has time-sensitive traffic (VoIP) or bandwidth-intensive traffic (real-time or near-real-time streaming-video), QoS makes it possible for you to prioritize the time-sensitive traffic and assign larger amounts of bandwidth to those applications that require it.

VoIP traffic has relatively low bandwidth requirements, but cannot tolerate latency or frame loss. Therefore, this traffic needs a high priority to ensure its timely delivery. On the other hand, streaming video is bandwidth-intensive but has large "jitter buffers" so can tolerate some latency. Thus, you can assign streaming video traffic a lower priority than voice, but must assign streaming video more bandwidth than voice.

Prerequisites

To successfully implement QoS, you must have a thorough knowledge of the traffic patterns in the network. You need this information to:

- Classify traffic and assign it the required priority and bandwidth.
- Identify the areas of the network where bottlenecks might occur and that therefore need bandwidth limiting.
- Identify the areas of the network where time-sensitive traffic is being delayed and needs to be prioritized better.

Implementation Example

An example of managing QoS across the network is to define traffic classes and manage these on a network-wide basis. The four classes and their priorities might look like those outlined in Table 25-1.

*** Note:** For information about DSCP (DiffServ Code Point), see "Diffserv" later in this chapter.

Table 25-1. Examples of Classes of Service

Service Class	Priority	DSCP Value	Type of Traffic
Highest Priority	7	56	Network Management Traffic, OSPF, Spanning Tree, etc.
Time Sensitive Traffic	5	40	Real-time voice, video conferences.
High Priority Data	1	8	SAP, Web, etc. Dependent upon your business.
Best Effort	0	0	Everything else.

How Does QoS Work?

The QoS process starts at the point where a frame enters the switch and ends when the frame exits the switch. This section describes the QoS process from start to finish:

- QoS Process for Ingress Traffic
- QoS Process for Egress Traffic
- Diagram of QoS Process

QoS Process for Ingress Traffic

Process

The QoS process for ingress traffic involves the following steps:

1. Identifying the priority, also called *class*, of the frame or packet. The switch can identify the priority of the frame or packet by using one or more of the following criteria:
   - The priority of the physical port that the switch received the frame or packet on
   - Cisco ISL tag priority
   - 802.1p tag priority (default)
   - The source or destination MAC address
   - The DiffServ code point
   - The IP protocol (assigned by means of an ACL rule)
   - The source or destination IP address (assigned by means of an ACL rule)
   - The source or destination TCP or UDP port (assigned by means of an ACL rule)

   For more information on identifying the priority of traffic, see "Classification of Traffic" later in this chapter.

2. Storing the frame or packet in one of eight ingress queues.

   The switch stores the frame or packet in the queue that matches the priority that was identified in Step 1.

3. Forwarding the frame or packet from the ingress queue to its destination.

   If you enable *policing* for the queue, the switch forwards ingress traffic that falls within the maximum bit rate that you set and drops ingress traffic that exceeds the maximum bit rate. For more information on policing, see "Ingress Policing" later in this chapter.

Example

You want to assign a priority of 5 to a VoIP flow that is destined to an IP 600 phone switch. You also want to police the port that receives the VoIP data to 5 Mbps.

1. You set up an ACL rule that associates a priority of 5 with the destination IP address of the VoIP flow. (VoIP traffic cannot tolerate latency or frame loss, so it needs a high priority to ensure its timely delivery.)

   ***Note:** Priority 5 serves as an example only. Actual implementations may vary.

2. You enable policing on the port that will receive the VoIP flow and set the guaranteed bit rate to 5 Mbps.
3. The switch stores packets that match the ACL rule in queue 5. (The switch stores packets in the queue that matches their priority.)
4. The switch forwards the VoIP traffic in queue 5 as long as its bit rate does not exceed 5 Mbps. If the bit rate of the queue exceeds 5 Mbps, the switch drops the excess traffic.

QoS Process for Egress Traffic

Process

The QoS process for egress traffic involves the following steps:

1. Storing the frame or packet in one of eight egress queues.

   The switch stores the frame or packet in the queue that matches the priority that was identified on the ingress port.

2. Forwarding the frame or packet from the queue to the egress port for transmission.

   The switch uses algorithms that allocate bandwidth among the egress queues to forward the traffic in the queues. The Avaya Multiservice switch software supports the following queue-servicing algorithms for egress ports:

   - Weighted fair queuing (WFQ)
   - Class-based queuing (CBQ)
   - Class-based weighted fair queuing (CBWFQ)
   - Strict Priority Queuing

Example

You want to assign a priority of 5 to a VoIP flow that is destined to an IP 600 phone switch. You then want the switch to use the CBWFQ queue-servicing algorithm to forward frames from queue 5. The IP 600 phone switch is connected to an Avaya P882 Multiservice switch on port 5.5.

1. You set up an ACL rule that associates a priority of 5 with the destination IP address of the VoIP flow. (VoIP traffic cannot tolerate latency or frame loss, so it needs a high priority to ensure its timely delivery.)

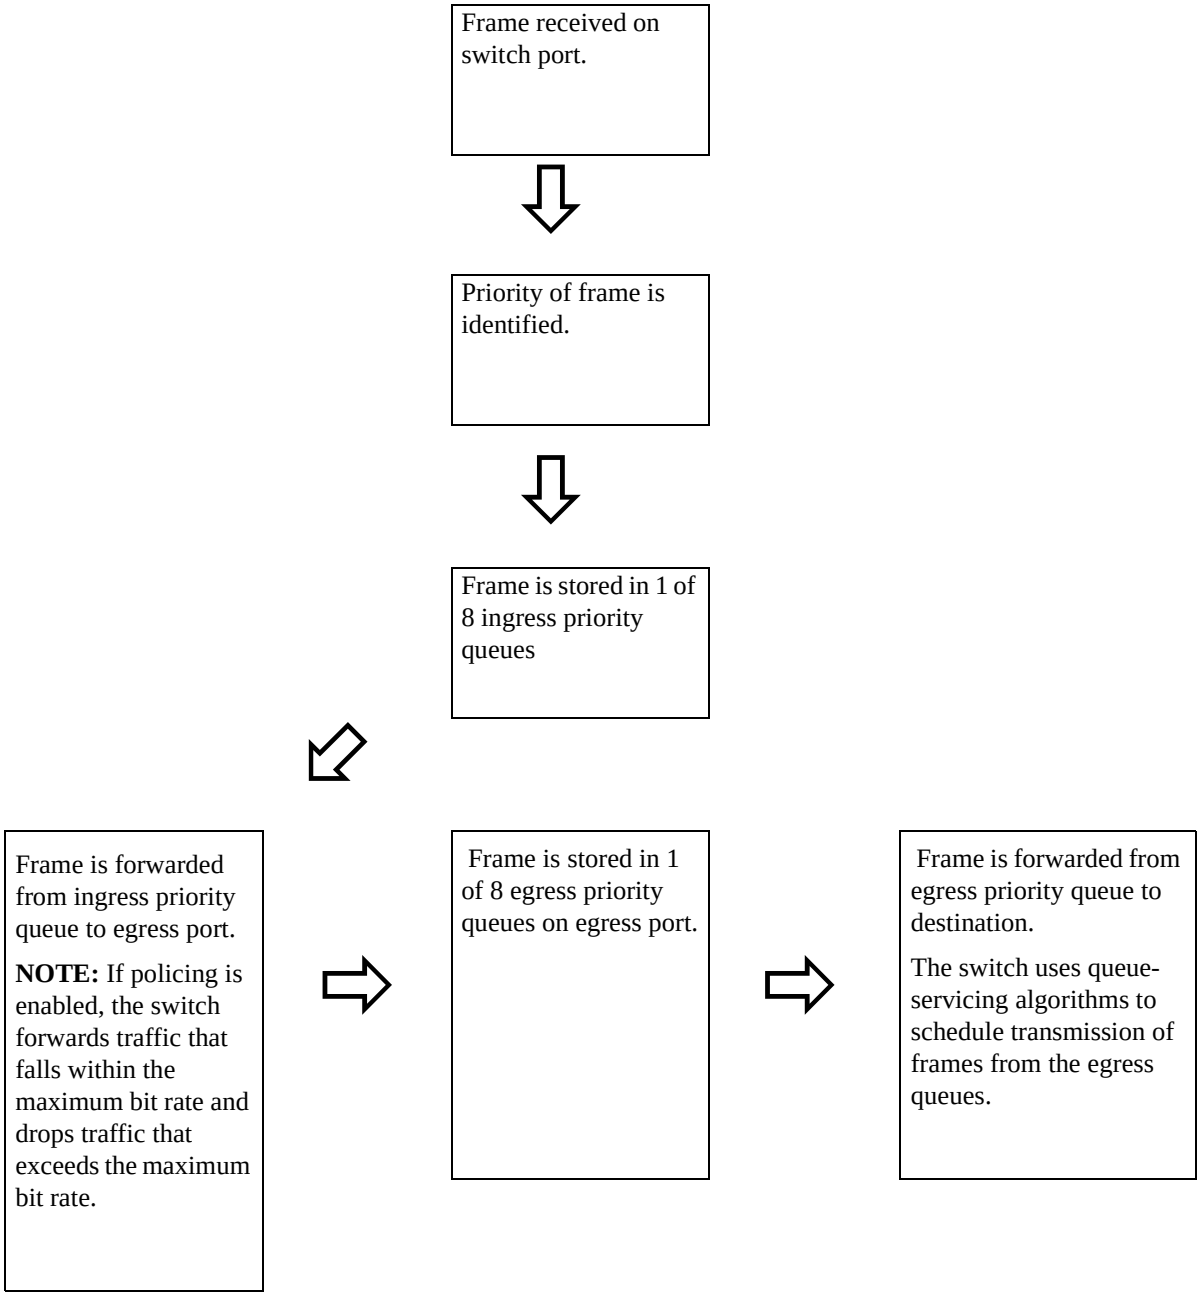
   ***Note:** Priority 5 serves as an example only. Actual implementations may vary.

2. You enable CBWFQ for queue 5 on port 5.5.
3. The switch assigns a priority of 5 to the ingress VoIP packets and forwards them through ingress queue 5 to the egress port.
4. The switch stores the VoIP packets in egress queue 5 on port 5.5. (The switch places packets in the queue that matches their priority.)
5. The switch uses the CBWFQ algorithm to remove the packets from queue 5 and forward them to the IP 600 phone switch.

Diagram of QoS Process

Figure 25-1 illustrates the QoS process from when a frame enters the switch to when the frame exits the switch.

Figure 25-1. QoS Process



Classification of Traffic

The switch assigns traffic to one of eight queues according to the priority, or "class," of the traffic. Priorities range 0 to 7, 7 being the highest priority. You can set the switch to classify traffic by the priority assigned to the following characteristics:

- Layer 2 characteristics:
  - Physical port that the frame or packet is received on
  - Cisco ISL tag or 802.1p tag
  - Source MAC address
  - Destination MAC address
- Layer 3 characteristics:
  - DSCP in the packet
  - New DSCP that replaces the original DSCP. You specify this new DSCP.
  - IP protocol (assigned by means of an ACL rule)
  - Destination IP address (assigned by means of an ACL rule)
  - Source IP address (assigned by means of an ACL rule)
- Layer 4 characteristics:
  - Destination TCP or UDP port (assigned by means of an ACL rule)
  - Source TCP or UDP port (assigned by means of an ACL rule)

This section contains the following information and procedures:

- Default Priority
- Classifying Traffic by Layer 2 Characteristics
- Classifying Traffic by Layer 3 or Layer 4 Characteristics
- Diffserv
- Precedence of Priorities
- Supported Number of Queues
- Setting the Priority of a Physical Port

- Setting a Physical Port to Ignore Tag Priority
- Setting the Priority of a MAC Address
- Displaying the Priority of a MAC Address
- Setting a Physical Port to Use DiffServ
- Setting a Physical Port to Mask DiffServ Bits
- Assigning a Priority to a DSCP
- Displaying the DiffServ Table
- Displaying the QoS Settings for a Physical Port
- Setting Up an ACL Rule
- Setting Up a Default ACL Rule
- Displaying ACL Rules

Default Priority

By default, the switch uses the priority from the 802.1p tag field, if present, to classify a frame.

If you do not change any of the QoS default settings and the frame does not have an 802.1 tag or Cisco ISL tag, the switch assigns the priority of the physical port to the packet. Each physical port has a default priority of 3. For information on how to change the priority for a physical port, see "Setting the Priority of a Physical Port" later in this chapter.

However, the priority of the 802.1 tag and Cisco ISL tag take precedence over the priority of the physical port, so the switch uses the priority of the physical port *only* if:

- No tags are present in the frame

  or

- You have set the physical port to ignore priorities in tags.

  For information on how to set a port to ignore priorities in tags, see "Setting a Physical Port to Ignore Tag Priority" later in this chapter.

*** Note:** 802.1p packets that are received with a tag priority of 0 on a 50-series layer 2 (non-routing) module, and that must be routed via the FORE path on an 80-series supervisor, are queued and transmitted with a priority of 4. This priority change is due to the conversion from the high-low priority system that 50-series modules use to the 8-level priority system that 80-series modules use.

Classifying Traffic by Layer 2 Characteristics

In addition to Cisco ISL tag, 802.1p tag, and physical port priority, the switch can classify traffic by:

- Source MAC address
- Destination MAC address

For information about how to set a priority for a source or destination MAC address, see “Setting the Priority of a MAC Address” later in this chapter.

In addition to these layer 2 characteristics, you can classify bridged IP traffic by DiffServ code point. For more information on classifying bridged IP traffic by DiffServ code point, see “Diffserv” later in this chapter.

Classifying Traffic by Layer 3 or Layer 4 Characteristics

ACL Rules

You can, alternately, configure the switch to classify traffic by the IP characteristics of packets, instead of 802.1p tag priority, physical port priority, or other layer 2 characteristics.

To assign priorities to packets by their IP characteristics, you create a rule in an access control list (ACL). The rule can:

- Set an ACL rule priority
- Use the DiffServ code point
- Mask the three least significant bits of the DSCP. The switch recognizes the remaining bits as the IP precedence field.
- Replace the existing DSCP with a DSCP that you specify
- Use the layer 2 priority

You can specify the TCP/IP traffic that you want the rule to apply to. The rule can apply to traffic with a specific:

- IP destination address
- IP source address
- IP Protocol
- Destination TCP or UDP port
- Source TCP or UDP port

The priority that is specified by an ACL takes precedence over all other priorities. Because of this precedence, the switch determines whether a rule in an ACL exists for an IP packet in the final stage of classification. If an ACL exists, the priority associated with the ACL replaces the current priority of the frame.

For information on how to set a rule in an ACL, see "Setting Up an ACL Rule" later in this chapter.

Using a Default ACL Rule

By default, the switch classifies packets by their layer 2 priority, if they do not match an ACL rule. However, you can set up a *default ACL rule* that the switch will apply to all packets that do not match any other ACL rules. This default ACL rule sets a default characteristic, other than the layer 2 priority, that the switch will use to classify the packets.

For example, if you set up the following ACL rules:

- access-list List1 1 fwd2 10.10.60.0 0.0.0.255
- access-list List1 2 fwd4 10.10.70.0 0.0.0.255
- access-list List1 3 permit use-priority 4 10.10.80.0 0.0.0.255
- access-list List1 4 permit use-priority 6 10.10.90.0 0.0.0.255
- access-list List1 5 fwd1 10.10.100.0 0.0.0.255
- access-list List1 6 fwd4 10.10.110.0 0.0.0.255
- access-list List1 7 fwd7 10.10.120.0 0.0.0.255
- access-list List1 8 fwd8 10.10.130.0 0.0.0.255
- access-list List1 512 permit use-priority 6 any (default ACL rule)

All packets that have a source IP address in the 10.10.60 subnet are assigned a priority of 1(fwd*x* is 1 based, but priorities are 0 based.) All packets that have a source IP address in the 10.10.70 subnet are assigned a priority of 3, and so on. Any packets whose source IP addresses do not match the IP addresses in ACL rules one through eight, are assigned a priority of 6.

For information on how to set up a default ACL rule, see "Setting Up a Default ACL Rule" later in this chapter.

Diffserv

RFC 2475 defines a field in the layer 3 header of IP packets, called the DiffServ code point (DSCP). Typically, hosts or routers sending traffic into a DiffServ network mark each transmitted packet with the appropriate DSCP. The switch then uses the DSCP to classify packets. You can alternately set the switch to replace the DSCP in a packet with a different DSCP. The switch then uses the new DSCP to classify the packet.

To set the switch to classify IP packets by their DSCP or to replace the DSCP with a different DSCP, you must set up an ACL rule. For information on how to set up an ACL rule to enable DiffServ functionality, see "Setting Up an ACL Rule" later in this chapter.

The Differentiated Services (DiffServ) Mapping Table associates specific DSCP values with specific priorities. You create these associations by assigning priorities to DSCPs. For information on how to assign priorities to DSCPs, see "Assigning a Priority to a DSCP" later in this chapter.

Because the DSCP is located in the layer 3 header, the switch does not typically use the DSCP to classify bridged IP traffic. However, you can set a physical port to use the DSCP to classify bridged IP traffic. For information on how to set a physical port to use DiffServ, see "Setting a Physical Port to Use DiffServ" later in this chapter.

Precedence of Priorities

If multiple priorities are associated with a frame or packet, the switch classifies the frame or packet according to the priority of highest precedence. See Table 25-2 for the precedence of each priority.

The switch then assigns the frame or packet to the appropriate priority queue based on the priority of the frames.

Table 25-2. Precedence of Priorities

Layer	Precedence of Priority	Priority Used for Classification
Layer 3		
	High	ACL rule priority
		DSCP in the packet or DSCP that the switch replaces the original DSCP with
		1 of 2

Table 25-2. Precedence of Priorities

Layer	Precedence of Priority	Priority Used for Classification
Layer 2		
		DSCP (for bridged IP traffic)
		Destination MAC address priority
		Source MAC address priority
		802.1p tag
		Cisco ISL tag
	Low	Physical port priority
		2 of 2

Supported Number of Queues

Table 25-3 specifies the number of ingress and egress queues that are available on each module.

Table 25-3. Modules and Available Queues

Module	Ingress Queues	Egress Queues
4-port gigabit modules	1 set of 8 queues per port **Total:** 4 sets of 8 queues	1 set of 8 queues per port **Total:** 4 sets of 8 queues
8-port gigabit modules	1 set of 8 queues per port **Total:** 8 sets of 8 queues	1 set of 8 queues per port **Total:** 8 sets of 8 queues
24-port 10/100 modules	• 1 set of 8 queues for ports 1–12. • 1 set of 8 queues for ports 13–24. **Total:** 2 sets of 8 queues	1 set of 8 queues per port. **Total:** 24 sets of 8 queues
		1 of 2

Table 25-3. Modules and Available Queues

Module	Ingress Queues	Egress Queues
48-port 10/100 modules	• 1 set of 8 queues for ports 1–12. • 1 set of 8 queues for ports 13–24. • 1 set of 8 queues for ports 25–36 • 1 set of 8 queues for ports 37–48 **Total:** 4 sets of 8 queues	1 set of 8 queues per port. **Total:** 48 sets of 8 queues.
		2 of 2

Setting the Priority of a Physical Port

CLI Command

Use the **set port default-priority** command to set the priority of a physical port. Each physical port has a default priority of 3. The syntax of the command is:

`(configure)#` **set port default-priority {{<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] | all-ports} <*priority*>**

Table 25-4. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, the priority is set for all ports on the module.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py.* For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the priority is set for the port or range of ports on the module that you specify.
	1 of 2

Table 25-4. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
all-ports	All ports in the chassis. If you specify **all-ports**, all ports on all modules in the chassis are set with the same priority.
<*priority*>	The priority that you want to assign to the port or port range. Enter a number between 0 and 7. The highest priority is 7. Each physical port has a default priority of 3.
	2 of 2

Examples

Table 25-5. Examples: set port default-priority

To...	Enter...
Set the port priority to 0 for all ports on the module in slot 3	**set port default-priority 3 0**
Set the port priority to 5 for ports 1 through 5 on the module in slot 3	**set port default-priority 3/1-5 5**
Set the port priority to 2 for ports 1 through 5 on the module in slot 3 and for port 1 on the module in slot 6	**set port default-priority 3/1-5,6/1 2**

Setting a Physical Port to Ignore Tag Priority

CLI Command

Use the **set port ignore-tag priority** command to set a port to ignore any layer 2 tag priority (including 802.1p tags). The syntax of the command is:

`(configure)#` **set port ignore-tag-priority {{<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] | all-ports} {on | off}**

Table 25-6. Keywords, Arguments, or Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, the switch ignores tag priorities on all ports of the module.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py.* For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the switch ignores tag priorities on the port or range of ports on the module in the slot that you specify.
all-ports	All ports in the chassis. If you specify **all-ports**, all ports on all modules in the chassis are set with the same priority.
{on | off}	Indicates whether you want the port to ignore tag priority. Enter **on** for the port to ignore the tag priority. The default setting is **off**.

Examples

Table 25-7. Examples: set port ignore-tag-priority

To...	Enter...
Set all ports on the module in slot 3 to ignore the 802.1p tag priority	**set port ignore-tag-priority 3 on**
Set ports 1 through 5 on the module in slot 3 to *not* ignore the 802.1p tag priority	**set port ignore-tag-priority 3/1-5 off**
Set ports 1 through 5 on the module in slot 3 and port 1 on the module in slot 6 to ignore the 802.1p tag priority	**set port ignore-tag-priority 3/1-5,6/1 on**

Setting the Priority of a MAC Address

CLI Command

Use the **set aft entry** command to set the priority of a source MAC address or destination MAC address. The syntax of the command is:

`(configure)#` **set aft entry <*mac-address*> vlan {<*vlan-id*> | name <*vlan-name*>} port-binding {filter | forward <*mod-port-spec*>} [persistence {ageout | permanent}] [priority {normal | high}] [sa-priority {port | aft <*entry-priority*> | max-port-aft <*entry-priority*>}] [da-priority {port | aft <*entry-priority*> | max-port-aft <*entry-priority*>}]**

Table 25-8. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
sa-priority port	Uses the priority of the physical port, Cisco ISL tag, or 802.1p tag to determine the layer 2 priority of frames.
sa-priority aft	Uses the priority that is assigned to the source MAC address in the Address Forwarding Table (AFT) to determine the layer 2 priority of frames.
<*entry-priority*>	The priority that you want to assign to the source MAC address. Enter a number between 0 and 7. This priority is stored in the AFT entry for the MAC address that you specify.
sa-priority max-port-aft	Determines the priority of a frame by using the higher of the: • Physical port priority or tag priority • Source MAC address priority
da-priority port	Uses the priority of the physical port, Cisco ISL tag, 802.1p tag, or source MAC address to determine the layer 2 priority of frames.
da-priority aft	Uses the priority that is assigned to the destination MAC address in the AFT to determine the priority of the frame.
	1 of 2

Table 25-8. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<entry-priority>	The priority that you want to assign to the destination MAC address. Enter a number between 0 and 7.
da-priority max-port-aft	Determines the priority of the frame by using the higher of the: • Physical port priority or tag priority • Destination MAC address priority
	2 of 2

For definitions of all other keywords, arguments, and options in this command, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*.

Examples

Table 25-9. Examples: set aft entry

To...	Enter...
• Associate MAC address 00:00:00:00:00:55 with port 1 on the module in slot 3 and with VLAN 50. • Forward frames that have a source or destination MAC address of 00:00:00:00:00:55. • Assign a priority of 7 to frames that have a source MAC address of 00:00:00:00:00:55.	**set aft entry 00:00:00:00:00:55 VLAN 50 port-binding forward 3/1 sa-priority aft 7**
• Associate MAC address 00:00:00:00:00:55 with port 1 on the module in slot 3 and with VLAN 50. • Forward frames that have a source or destination MAC address of 00:00:00:00:00:55. • Associate a priority of 5 with the source MAC address of 0:00:00:00:00:55. • Assign the higher of the port priority, tag priority, or source MAC address priority (5) to frames that have a source MAC address of 00:00:00:00:00:55.	**set aft entry 00:00:00:00:00:55 VLAN 50 port-binding forward 3/1 sa-priority max-port-aft 5**
	1 of 2

Table 25-9. Examples: set aft entry

To...	Enter...
• Associate MAC address 00:00:00:00:00:55 with port 1 on the module in slot 3 and with VLAN 50. • Forward frames that have a source or destination MAC address of 00:00:00:00:00:55. • Assign a priority of 7 to packets that have a destination MAC address of 00:00:00:00:00:55.	**set aft entry 00:00:00:00:00:55 VLAN 50 port-binding forward 3/1 da-priority aft 7**
• Associate MAC address 00:00:00:00:00:55 with port 1 on the module in slot 3 and with VLAN 50. • Forward frames that have a source or destination MAC address of 00:00:00:00:00:55. • Associate a priority of 5 with the destination MAC of address 0:00:00:00:00:55. • Assign the higher of the port priority, tag priority, or destination MAC address priority (5) to frames that have a destination MAC address of 00:00:00:00:00:55.	**set aft entry 00:00:00:00:00:55 VLAN 50 port-binding forward 3/1 da-priority max-port-aft 5**
	2 of 2

Displaying the Priority of a MAC Address

CLI Command

Use the **show aft entry** command to display the priority of a source MAC address or destination MAC address. The syntax of the command is:

> show aft entry [mac <*wildcard-mac-address*>] [VLAN {<*VLAN-id*> | name <*VLAN-name*>}] [port-binding {cpu | filter | forward [<*mod-port-spec*>]}] [status {learned | management | self | multicast}]

Setting a Physical Port to Use DiffServ

CLI Command

Use the **set port use-diffserv** command to set a port to classify bridged IP traffic by its DiffServ code point (DSCP). The syntax for the command is:

(configure)# **set port use-diffserv {{<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] | all-ports} {on | off}**

Table 25-10. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, all ports on the module are set to use DiffServ.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the port or range of ports on the module that you specify is set to use DiffServ.
{all-ports}	All ports in the chassis. If you enter **all-ports**, all ports in the chassis are set to use DiffServ.
{on | off}	Indicates whether the switch uses the DSCP to classify bridged IP traffic on the ports that you specify: • Enter **on** to classify traffic by DSCP. • Enter **off** to ignore the DSCP. The default setting is **off**.

Example

To set ports 4 through 12 on the module in slot 6 to classify bridged IP traffic by DSCP, enter:

set port use-diffserv 6/4-12 on

Setting a Physical Port to Mask DiffServ Bits

CLI Command

Use the **set port mask-diffserv** command to mask the three least significant bits of the DSCP when the switch is using the DSCP to classify bridged IP traffic. If you mask the three least significant bits of the DSCP, the switch recognizes the remaining bits as the precedence field of the type of service (TOS) field and classifies the packets accordingly.

The syntax of this command is:

`(configure)#` **set port mask-diffserv {{<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] | all-ports} {on | off}**

Table 25-11. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, all ports on the module mask the three least significant bits of the DSCP.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the port or range of ports on the module that you specify mask the three least significant bits of the DSCP.
{all-ports}	All ports in the chassis. If you enter **all-ports**, all ports in the chassis are set to mask the three least significant bits of the DSCP.
{on | off}	Indicates whether the switch masks the three least significant bits of the DSCP: • Enter **on** to mask the bits. • Enter **off** to not mask the bits. The default setting is **off**.

Example

To set all ports on the module in slot 3 to mask the three least significant bits of the DSCP, enter:

set port mask-diffserv 3 on

Assigning a Priority to a DSCP

CLI Command

Use the **set diffserv priority** command to assign a priority to a DiffServ code point (DSCP) in the DiffServ Mapping Table. The syntax of the command is:

`(configure)#` **set diffserv priority <*priority*> dscp <*dscp-start-range*> [<*dscp-end-range*>]**

Table 25-12. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*priority*>	The priority that you want to assign. Enter a number between 0 and 7.
<*dscp-start-range*>	The first DSCP in the range of DSCPs that you want to assign the priority to. DSCPs range from 0 to 63.
[<*dscp-end-range*>]	The last DSCP in the range of DSCPs that you want to assign the priority to. DSCPs range from 0 to 63.

Example

To assign a priority of 7 to DSCPs 15 through 63, enter:

set diffserv priority 7 dscp 15 63

Displaying the DiffServ Table

CLI Command

Use the **show diffserv table** command to display the priority that is assigned to each DSCP.

The syntax of the command is:

> show diffserv table

*** Note:** The **show diffserv table** CLI command displays the packet loss probability (PLP) for each DSCP. However, the switch does not currently support PLP.

Displaying the QoS Settings for a Physical Port

CLI Command

Use the **show port** command to display the QoS settings for a physical port. This command displays the priority of the port, if the port is set to ignore 802.1p tag priority, and if the port is set to use the DSCP for bridged IP traffic.

The syntax of the command is:

> show port [{<*mod-num*> | <*mod-swport-range*>} [...,{<*mod-num*> | <*mod-swport-range*>}]]

Table 25-13. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, the switch displays the QoS settings for all ports on the module in the slot that you specify.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the switch displays the QoS settings for the port or range of ports on the module in the slot that you specify.

Setting Up an ACL Rule

CLI Command

Use the **access list** command to set a rule in an access control list. The rule that you set up will be applied on all of the ports on the switch. For more information about using ACL rules, see “Classifying Traffic by Layer 3 or Layer 4 Characteristics” earlier in this chapter.

*** Note:** You must enable the access control list on which you want to set a rule. Only one access control list can be enabled at a time.

The command syntax is:

Standard ACL

`(configure)#` **access-list <*access-list-name*> <*access-list-index*> {permit [{use-priority <*priority*> | use-diffserv [mask] | remark-diffserv <*dscp*> [mask] | use-l2}] | deny | fwd1 | fwd2 | fwd3 | fwd4 | fwd5 | fwd6 | fwd7 | fwd8} {<*source-ip-addr*> <*source-wildcard*> | any | host <*source-ip-addr*>}**

Extended ACL

`(configure)#` **access-list <*access-list-name*> <*access-list-index*> {permit [{use-priority <*priority*> | use-diffserv [mask] | remark-diffserv <*dscp*> [mask] | use-l2}] | deny | fwd1 | fwd2 | fwd3 | fwd4 | fwd5 | fwd6 | fwd7 | fwd8} <*protocol-id*> {<*source-ip-addr*> <*source-wildcard*> | any | host <*source-ip-addr*>} [{lt <*port*> | eq <*port*> | gt <*port*> | range <*port*> <*port*>}] {<*dest-ip-addr*> <*dest-wildcard*> | any | host <*dest-ip-addr*>} [{lt <*port*> | eq <*port*> | gt <*port*> | range <*port*> <*port*>}] [established]**

Keywords, Arguments, and Options

Table 25-14. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*access-list-name*>	A unique name that identifies the access control list.
<*access-list-index*>	The unique rule number within the access list.
permit	Forwards the packet without changing its priority.
use-priority	Assigns the priority that you define in the following <*priority*> parameter to the packet.
<*priority*>	The priority that you want to assign to packets that match this ACL. Enter a number between 0 and 7.
use-diffserv	Classifies traffic by the DSCP in the packet.
[mask]	Masks the three least significant bits of the DSCP. If you mask the three least significant bits of the DSCP, the switch recognizes the remaining bits as the precedence field of the type of service (TOS) field and classifies the packets accordingly.
remark-diffserv	Replaces the DSCP in the packet with the DSCP that you enter for the following <*dscp*> parameter. The switch uses the DSCP that you enter for the <*dscp*> parameter to classify the packet.
	1 of 3

Table 25-14. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*dscp*>	The DSCP that you want to replace the DSCP of the packet.
use-l2	Classifies traffic by the layer 2 priority of the packet. If you enter **use-l2**, the switch ignores the ACL rule priority and DiffServ priority.
deny	Blocks the packet.
fwd1 | fwd2 | fwd3 | fwd4 | fwd5 | fwd6 | fwd7 | fwd8	The priority that you want to set. The number following the **fwd** specifies the priority. The **fwd*x*** arguments are 1-based, while the queue priorities are 0-based. Consequently, the 1-based priorities are converted to 0-based priorities by the QoS features. For example, to specify a priority of 0, enter **fwd1**. These keywords are retained from earlier versions of software for backward compatibility. The **use-priority <*priority*>** keyword and argument serve the same function.
<*protocol-id*>	The ID of the protocol that you want to assign a priority to. RFC791 defines the protocol IDs.
<*source-ip-addr*>	The source IP address of the subnet that you want to assign a priority to.
<*source-wildcard*>	The inverse of a network mask. Enter a 32-bit number in four-part, dotted decimal format. Place ones in the bit positions that you want to mask. This parameter specifies a range of IP address. For example, to specify all IP addresses in the 10.10.70 subnet, enter **10.10.70.0 0.0.0.255**.
any	A source of 0.0.0.0 and a source-wildcard of 255.255.255.255
host <*source-ip-addr*>	The source IP address that you want to assign a priority to.
[{lt <*port*> | eq <*port*> | gt <*port*> | range <*port*> <*port*>}]	A source port or range of source ports that pass between two hosts or switches using the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP).
	2 of 3

Table 25-14. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*dest-ip-addr*>	The destination IP address of the subnet that you want to assign a priority to.
<*dest-wildcard*>	The inverse of a network mask. Enter a 32-bit number in four-part, dotted decimal format. Place ones in the bit positions that you want to mask. This parameter specifies a range of IP address. For example, to specify all IP addresses in the 10.10.70 subnet, enter **10.10.70.0 0.0.0.255**.
any	A destination of 0.0.0.0 and a destination-wildcard of 255.255.255.255
host <*dest-ip-addr*>	The destination IP address that you want to assign a priority to.
[{lt <*port*> | eq <*port*> | gt <*port*> | range <*port*> <*port*>}]	A destination port or range of destination ports that pass data between two hosts or switches using the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP).
[established]	Permits TCP connections to be established that match the rule.
	3 of 3

*** Note:** You must use the CLI to set up ACLs for QoS. Do not attempt to use the IP Access Control Web page in the Web Agent to set up ACLs for QoS.

Examples: Standard ACL Rules

Table 25-15. Examples: Standard ACL Rules

To...	Enter...
• Use the DSCP in the packet to classify all traffic that has a source IP address in the 10.10.60 subnet. • Mask the three least significant bits of the DSCP.	**access-list MyAccessList1 4 permit use-diffserv mask 10.10.60.0 0.0.0.255**
Assign a priority of 7 to all traffic that has a source IP address in the 10.10.70 subnet.	**access-list MyAccessList1 5 permit use-priority 7 10.10.70.0 0.0.0.255**
	1 of 2

Table 25-15. Examples: Standard ACL Rules

To...	Enter...
• Replace the existing DSCP with a DSCP of 5 for all traffic that has a source IP address in the 10.10.80 subnet. • Mask the three least significant bits of the DSCP	**access-list MyAccessList1 6 permit remark-diffserv 5 mask 10.10.80.0 0.0.0.255**
Use the layer 2 priority of the packet to classify all traffic that has a source address in the 11.11.11 subnet	**access-list MyAccessList1 7 permit use-l2 11.11.11.0 0.0.0.255**
Use the DSCP in the packet to classify all traffic that has a source IP address of 199.93.239.168	**access-list MyAccessList1 8 permit use-diffserv host 199.93.239.168**
• Use the DSCP in the packet to classify all traffic that has a source IP address of 3.3.3.3 • Mask the three least significant bits of the DSCP	**access-list MyAccessList1 9 permit use-diffserv mask host 3.3.3.3**
Assign a priority of 2 to all traffic that has a source IP address of 1.1.1.1	**access-list MyAccessList1 10 permit use-priority 2 1.1.1.1**
Block all traffic that has a source IP address of 10.1.0.55	**access-list MyAccessList1 11 deny 10.1.0.55**
	2 of 2

Examples: Extended ACL Rules

Table 25-16. Examples: Extended ACL Rules

To...	Enter...
Use the DSCP in the packet to classify all traffic that has a: • Source IP address of 199.93.239.168 • Destination address in the 1.1.1 subnet	**access-list MyAcessList2 1 permit use-diffserv ip host 199.93.239.168 1.1.1.0 0.0.0.255**
Use the DSCP in the packet to classify all traffic that has a: • Source IP address in the 3.0 subnet • Destination address in the 5.0 subnet • Mask the three least significant bits of the DSCP	**access-list MyAccessList2 2 permit use-diffserv mask ip 3.0.0.0 0.255.255.255 5.0.0.0 0.255.255.255**
	1 of 3

Table 25-16. Examples: Extended ACL Rules

To...	Enter...
Assign a priority of 2 to all TCP traffic that has a: • Source IP address in the 1.1 subnet • Source port that is greater than 24 • Destination IP address in the 6.6 subnet • Destination port of 23	**access-list MyAccessList2 3 permit use-priority 2 tcp 1.1.0.0 0.0.255.255 gt 24 6.6.0.0 0.0.255.255 eq 23**
• Replace the existing DSCP of packets with a DSCP of 12 for all traffic that has a source IP address of 199.93.238.83. • Mask the three least significant bits of the DSCP.	**access-list MyAccessList2 4 permit remark l2 mask ip host 199.93.238.83 any**
Replace the existing DSCP of the packet with a DSCP of 24 for all ICMP traffic that has a: • Source IP address of 2.2.2.2 • Destination IP address of 4.4.4.4	**access-list MyAccessList2 5 permit remark 24 icmp host 2.2.2.2 host 4.4.4.4**
Assign a priority of 6 to all TCP traffic that has a: • Source IP address in the 10.10.10 subnet • Destination IP address in the 11.11.11 subnet • Destination port of 1	**access-list MyAccessList2 6 permit use-priority 6 tcp 10.10.10.0 0.0.0.255 11.11.11.0 0.0.0.255 eq 1**
Use the layer 2 priority in the packet to classify all UDP traffic	**access-list MyAccessList2 7 permit use-l2 udp any any**
• Use the layer 2 priority in the packet to classify all TCP traffic that has a: — Source IP address in the 5.5.5 subnet — Destination IP address in the 6.6.6 subnet — Destination port that is less than 2 • Permit TCP connections that meet this criteria	**access-list MyAccessList2 8 permit use-l2 tcp 5.5.5.0 0.0.0.255 6.6.6.0 0.0.0.255 lt 2 established**
	2 of 3

Table 25-16. Examples: Extended ACL Rules

To...	Enter...
• Use the DSCP to classify all UDP traffic that has a: — Source IP address of 7.7.7.7 — Destination IP address of 8.8.8.8 — Destination port between 33 and 44 • Mask the three least significant bits of the DSCP	**access-list MyAccessList2 9 permit use-diffserv mask udp host 7.7.7.7 host 8.8.8.8 range 33 44**
• Assign a priority of 7 to all TCP traffic that has a: — Source IP address of 9.9.9.9 — Destination IP address of 3.3.3.3 — Destination port between 55 and 66 • Permit TCP connections that meet this criteria	**access-list MyAccessList2 10 permit use-priority 7 tcp host 9.9.9.9 host 3.3.3.3 range 55 66 established**
	3 of 3

Setting Up a Default ACL Rule

CLI Command

Use the **any** keyword in the **access list** command to set up a default ACL rule. The rule will be applied to all packets on the switch that do not match any other ACL rules. The command syntax is:

`(configure)#` **access-list <*access-list-name*> <*access-list-index*> {permit [{use-priority <*priority*> | use-diffserv [mask] | remark-diffserv <*dscp*> [mask] | use-l2}] | fwd1 | fwd2 | fwd3 | fwd4 | fwd5 | fwd6 | fwd7 | fwd8} any**

The default ACL rule must have the highest index in the ACL. To ensure that the switch never applies the default ACL rule to traffic that matches other ACL rules, Avaya recommends that you use an index of 512 for the default ACL rule.

For more information about how default ACL rules work, see “Classifying Traffic by Layer 3 or Layer 4 Characteristics” earlier in this chapter.

Examples

To...	Enter...
Use the DSCP in the packet to classify all traffic that does not match any other ACL rule.	**access-list MyAcessList1 512 permit use-diffserv any**
Replace the existing DSCP with a DSCP of 63 for all traffic that does not match any other ACL rule. The switch uses the new DSCP of 63 to classify the packets.	**access-list MyAcessList1 512 permit remark-diffserv 63 any**
Use the layer 2 priority in the packet to classify all traffic that does not match any other ACL rule.	**access-list MyAcessList1 512 permit use-l2 any**
Assign a priority of 4 to all traffic that does not match any other ACL rule.	**access-list MyAcessList1 512 permit use-priority 4 any**

Displaying ACL Rules

CLI Command

Use the **show access-lists** command to display the ACL rules in an ACL.

The syntax of this command is:

> show access-lists [<*access-list-name*>]

Ingress Policing

Policing makes it possible for you to limit the bandwidth for ingress queues. You limit the bandwidth by specifying the guaranteed bit rate for a port. If this bit rate is exceeded, the switch drops the excess packets.

For example, if you set policing on an ingress queue to be 5 Mbps, and traffic exceeds that 5 Mbps rate, all traffic that exceeds the 5Mbps is dropped.

Only 80-series modules that are licensed for routing support the policing feature.

The policing algorithm includes a normal burst threshold. This threshold sets the size of bursts that is guaranteed transfer.

The switch uses queue 0 to forward protocol packets (ARP, VRRP, OSPF, and so on) to the supervisor module. If you enable policing on queue 0, be sure to allocate the queue enough bandwidth for management packets and learned packets. Failure to allocate enough bandwidth to the queue may result in poor network performance. Do not disable this queue by allocating it 0 bits per second (bps). If you disable the queue, all protocol packets and learned packets are discarded before reaching the supervisor module.

For information about how to set up policing, see "Setting Up Policing" later in this chapter.

*** Note:** Avaya recommends that you do not set a port using policing as the source port or mirror port for a port mirror. When the switch limits the bandwidth of a port, packets are subject to random drop. If packets from a source port or mirror port are dropped, the mirror traffic may not match the source traffic.

This section contains the procedures:

- Setting Up Policing
- Displaying the Policing Settings

Setting Up Policing

Use the **set port police** command to enable or disable policing for ingress traffic on a port. For information on how policing works, see "Ingress Policing" earlier in this chapter.

The syntax for this command is:

`(configure)#` **set port police {{<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] | all-ports} queue <*queue*> {bit-rate <*rate*> normal-burst <*normal-burst*> | disable}**

Table 25-17. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, policing is enabled for all ports on the module in the slot that you specify.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py.* For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, policing is enabled for the port or range of ports on the module in the slot that you specify.
all-ports	All ports in the chassis. If you specify **all-ports**, policing is enabled on all modules in the chassis.
<*queue*>	The queue number, which can range from 0 to 7.
<*rate*>	The maximum bits per second that you want to assign to the queue. For Fabric mode 1, enter: • 0 to disable the queue **Or** • 220 Kbps to 1.5 Gbps For Fabric mode 2, enter: • 0 to disable the queue **Or** • 270 Kbps to 1.5 Gbps
<*normal-burst*>	This threshold sets the maximum size of burst that is guaranteed transfer. The normal burst can range from 0 to 15,000. Avaya recommends a setting of 4.
disable	Disables policing.

Displaying the Policing Settings

Use the **show port police** command to display the settings for policing. For information on how policing works, see "Ingress Policing" earlier in this chapter.

The syntax for this command is:

> show port police {<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}]

Table 25-18. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, the switch displays the policing settings for all ports on the module in the slot that you specify.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the switch displays the policing settings for the port or range of ports on the module in the slot that you specify.

Queue-Servicing Algorithms

The following queue-servicing algorithms are available for egress queues:

- Weighted fair queueing (WFQ)
- Strict Priority
- Class-based queueing (CBQ)
- Class-based weighted fair queueing (CBWFQ)

*** Note:** In earlier versions of the switch software, you could set ingress queues to use the weighted fair queuing (WFQ) and strict priority queue-servicing algorithms. In v5.3.1 and later, you can set only egress queues to use these queue-servicing algorithms. To service ingress queues, use the policing feature.

This section contains the following information and procedures:

- WFQ
- Strict Priority Queueing
- CBQ
- CBWFQ
- Management Traffic
- Setting Up WFQ
- Setting Up Strict Priority Queueing
- Setting Up CBQ
- Setting Up CBWFQ
- Displaying the Queue-Service Settings

WFQ

How WFQ Works

WFQ is the default queue-servicing algorithm. When a port is set to use the WFQ algorithm, each queue:

- Is assigned a weight increment. This value never changes.
- Maintains an accumulated weight. After the switch services a queue, its accumulated weight is reset to the value of its weight increment, and the accumulated weight for the other queues is increased by their respective weight increments.

The switch always services the queue that has the highest *accumulated weight*. If two queues have the same accumulated weight, the switch first services the queue that has the highest priority (0 – 7).

Table 25-19 lists the default weight increment for each queue.

Table 25-19. Default Weight Increments

Queue	Weight Increment
WFQ 0	1
WFQ 1	2
WFQ 2	4
WFQ 3	8
WFQ 4	16
WFQ 5	32
WFQ 6	64
WFQ 7	128

For information on how to set up WFQ, see "Setting Up WFQ" later in this chapter.

Example

Suppose that the accumulated weight of all queues is set at their weight increments.

1. The switch services queue 7 because its accumulated weight is 128.
2. The accumulated weight of queue 7 is reset to 128, and the accumulated weight of all other queues is increased by their weight increment.

   At this point, queue 7 and queue 6 both have an accumulated weight of 128.
3. The switch services queue 7 because it is the higher priority queue.
4. The accumulated weight of queue 7 is then reset to 128, and the accumulated weight of all other queues is increased by their weight increment.

   Now queue 6 has an accumulated weight of 192, and queue 7 has an accumulated weight of 128.
5. The switch services queue 6 because it has the higher accumulated weight.

Strict Priority Queueing

With strict priority queuing, the switch services the eight queues in order of their priority. The highest priority queue is serviced until it is empty, and then the lower priority queues are serviced sequentially until they are empty. For example, queue 7 must be empty before the switch services queue 6. Queue 6 must be empty before the switch services queue 5. Queue 5 must be empty before the switch services queue 4, and so on.

For information on how to set up strict priority queueing, see "Setting Up Strict Priority Queueing" later in this chapter.

CBQ

With the class-based queuing (CBQ) algorithm, you can specify a maximum bit rate, or bandwidth, for a queue. When that bit rate is exceeded, the switch drops all packets that exceed the bit rate.

The bit rate that you set should be the long-term average rate of transmission. Traffic that does not exceed this rate is guaranteed transfer.

For information on how to set up CBQ queueing, see "Setting Up CBQ" later in this chapter.

*** Note:** Avaya recommends that you do not set a port using CBQ as the source port or mirror port for a port mirror. When the switch limits the bandwidth of a port, packets are subject to random drop. If packets from a source port or mirror port are dropped, the mirror traffic may not match the source traffic.

CBWFQ

The class-based weighted fair queueing (CBWFQ) algorithm is a combination of the CBQ and WFQ algorithms. CBWFQ makes more complex management of traffic possible. Using CBWFQ, you can set:

- The maximum bit rate.

  The bit rate that you set should be the long-term average rate of transmission. Traffic that does not exceed this rate is guaranteed transfer.

- Burst thresholds

- The action that you want the switch to take when the bit rate exceeds the maximum bit rate.
  - Drop the packets

    Or

  - Forward the packets according to the weight of the queue

*** Note:** Avaya recommends that you do not set a port using CBWFQ as the source port or mirror port for a port mirror. When the switch limits the bandwidth of a port, packets are subject to random drop. If packets from a source port or mirror port are dropped, the mirror traffic may not match the source traffic.

Burst Thresholds

The CBWFQ algorithm has two burst thresholds that you can set. Both of thresholds are set in a number of bytes.

- Normal burst—This threshold sets the maximum size burst that is guaranteed transfer. Bursts that are smaller than this size are guaranteed transfer. Bursts that are larger than this size are either serviced by WFQ or dropped (whichever action that you specify). The default setting is servicing by WFQ and the default weight for the queues. Normal burst size can range from 0 to 15,000 bytes.
- Maximum burst—This threshold sets the maximum size burst that is serviced by WFQ once the normal burst has been exceeded. Bursts that are smaller than this size are serviced by WFQ. Bursts that are larger than this size are dropped. If you set this threshold to the same value as normal burst, the maximum burst capability is disabled. The maximum burst can range from the normal burst size to 15,000.

Management Traffic

The P580 and P882 use queue 7 to forward protocol packets (ARP, VRRP, OSPF, and so on) from the supervisor module.

If you enable queue servicing on this queue, be sure to allocate the queue enough bandwidth for management packets. Failure to allocate enough bandwidth to the queue may result in poor network performance. Do not disable this queue by allocating it 0 bps. If you disable this queue, all protocol packets are discarded.

Setting Up WFQ

Use the **set port queue service wfq** command to set a port, port range, or module to use weighted fair queueing (WFQ) queue servicing. WFQ is the default queue-servicing algorithm.

The syntax of the command is:

`(configure)#` **set port queue service {{<*mod-num*> | <*mod-swport-range*>} [...,{<*mod-num*> | <*mod-swport-range*>}] | all-ports} wfq {queue <*queue*> weight <*weight*> | default}**

Table 25-20. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, all ports on the module are set to use WFQ.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the port or range of ports that you specify is set to use WFQ.
all-ports	All ports in the chassis. If you specify **all-ports**, all ports on all modules in the chassis are set to use WQF.
<*queue*>	The queue number, which can range from 0 to 7.
<*weight*>	The weight that you want to assign to the queue. Weights can range from 1 to 254.
default	The default weights. For the specific default weight of each queue, see Table 25-19.

For information about how WFQ works, see "WFQ" earlier in this chapter.

Setting Up Strict Priority Queueing

CLI Command

Use the **set port queue service strict-priority** command to set a port, port range, or module to use strict priority queue servicing.

The syntax of the command is:

`(configure)#` **set port queue service {{<*mod-num*> | <*mod-swport-range*>} [...,{<*mod-num*> | <*mod-swport-range*>}] | all-ports} strict-priority**

Table 25-21. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, all ports on the module are set to use strict priority queueing.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the port or range of ports on the module that you specify is set to use strict priority queueing.
all-ports	All ports in the chassis. If you specify **all-ports**, all ports on all modules in the chassis are set to use strict priority queueing.

For information on how strict priority queueing works, see "Strict Priority Queueing" earlier in this chapter.

Example

To set ports 1 through 12 on module 5 to use strict priority queueing, enter:

set port queue service 5/1-12 strict-priority

Setting Up CBQ

Use the **set port queue service cbq** command to set a port, port range, or module to use class-based queuing (CBQ) queue servicing. The syntax of the command is:

`(configure)#` **set port queue service {{<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] | all-ports} cbq queue <*queue*> bit-rate <*rate*>**

Table 25-22. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, all ports on the module are set to use CBQ.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the port or range of ports that you specify is set use CBQ.
all-ports	All ports in the chassis. If you specify **all-ports**, all ports on all modules in the chassis are set to use CBQ.
<*queue*>	The queue number, which can range from 0 to 7.
<*rate*>	The maximum bits per second that you want to assign to the queue. If the switch is operating in *Fabric mode 1*, the rate can range from 220 Kbps to 1.5 Gbps. If the switch is operating in *Fabric mode 2*, the rate can range from 270 Kbps to 1.5 Gbps Entering a rate of 0 disables the queue.

For information about how the CBQ queue-servicing algorithm works, see "CBQ" earlier in this chapter.

Setting Up CBWFQ

Use the **set port queue service cbwfq** command to set a port, port range, or module to use class-based weighted fair queuing (CBWFQ) queue servicing. The syntax of the command is:

`(configure)#` **set port queue service {{<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] | all-ports} cbwfq queue <*queue*> bit-rate <*rate*> normal-burst <*normal-burst*> [exceed {drop | max-burst <*max-burst*> [weight <*weight*>]}]**

Table 25-23. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, all ports on the module are set to use CBQ.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the port or range of ports that you specify is set use CBQ.
all-ports	All ports in the chassis. If you specify **all-ports**, all ports on all modules in the chassis are set to use CBQ.
<*queue*>	The queue number, which can range from 0 to 7.
<*rate*>	The maximum bits per second that you want to assign to the queue. If the switch is operating in *Fabric mode 1*, the rate can range from 220 Kbps to 1.5 Gbps. If the switch is operating in *Fabric mode 2*, the rate can range from 270 Kbps to 1.5 Gbps Entering a rate of 0 disables the queue.
	1 of 3

Table 25-23. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*normal-burst*>	The maximum size of burst that is guaranteed transfer. Bursts that are smaller than this size are guaranteed transfer. Bursts that are larger than this size are either serviced by WFQ or dropped (whichever action that you specify). The default setting is servicing by WFQ and the default weight for the queues. The normal burst can range from 0 to 15,000 bytes. Avaya recommends a value of 6000. Enter this setting in a multiple of four. If you do not enter a multiple of four, the switch rounds down the number that you enter to a multiple of four. For example, if you enter a normal burst size of 43 byes, the switch converts the setting to 40 bytes. If you enter a normal burst size of 0,1,2 or 3, the switch stores a value of 0 and no data is forwarded from the queue.
exceed	The action that you want the switch to take if the bit rate exceeds the guaranteed bit rate that you specify. The switch can either drop packets or forward them based on the weight of the queue.
drop	Indicates that you want the switch to drop packets when the bit rate exceeds the guaranteed bit rate.
	2 of 3

Table 25-23. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*max-burst*>	The maximum size burst that is serviced by WFQ once the normal burst has been exceeded. Bursts that are smaller than this size are serviced by WFQ. Bursts that are larger than this size are dropped. If you set this threshold to the same value as normal burst, the maximum burst capability is disabled. The maximum burst can range from the normal burst size to 15,000.Avaya recommends a value of 6000. Increase the maximum burst setting as the burstiness of the traffic increases. **Note:** The maximum burst setting must be greater than or equal to the normal burst. Enter this setting in a multiple of four. If you do not enter a multiple of four, the switch rounds down the number that you enter to a multiple of four. For example, if you enter a maximum burst size of 43 byes, the switch converts the setting to 40 bytes.If you enter a maximum burst size of 0,1,2 or 3, the switch stores a value of 0 and no data is forwarded from the queue.
<*weight*>	The weight that you want to assign to the queue. Weights can range from 1to 254.
	3 of 3

For information about how CBWFQ works, see "CBWFQ" earlier in this chapter.

*** Note:** The switch does not drop packets on egress queue 7, even if you specify the **drop** keyword. This restriction ensures that traffic management packets are not discarded by the bandwidth-limiting functions. In spite of the ability of egress queue 7 to exceed its bandwidth limit, the queue is subject to a fairness mechanism that prevents other queues (6 through 0) from being starved. Note that, when you specify **drop**, the bandwidth of queue 7 is limited if lower priority queues require service and have not reached their bandwidth guarantee.

Displaying the Queue-Service Settings

CLI Command

Use the **show port queue service** command to display the settings for queue servicing. The syntax for this command is:

> show port queue service {<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}]

Table 25-24. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, the switch displays queue-service settings for all ports on the module in the slot that you specify.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the switch displays the queue-service settings for the port or range of ports that you specify.

Example

To display the queue service settings for port 1 on the module in slot 3, enter:

show port queue service 3/1

QoS Statistics

You can display QoS statistics for:

- All ports on a module
- A port or port range
- A specific queue on a port, port range, or module

Two sets of statistics are displayed for egress queues: port statistics and queue statistics. For ingress queues, only queue statistics are displayed.

Table 25-25 lists the statistics that are displayed for ingress queues. Table 25-27 lists the statistics that are displayed for egress ports. And Table 25-28 lists the statistics that are displayed for egress queues.

For information on how to display the QoS Statistics, see "Displaying QoS Statistics" later in this chapter.

Queue Statistics

Table 25-25 lists the statistics that the switch displays for ingress and egress queues.

Table 25-25. Ingress Queue Statistics for 10/100 Modules

Statistic	Explanation
Frames Dropped Queue Full	The total number of frames that were dropped because the queue was full. This number is cumulative since the last reset of the statistics.
Frames Enqueued	The total number of frames stored in the queue. This number is cumulative since the last reset of the statistics.
	1 of 3

Table 25-25. Ingress Queue Statistics for 10/100 Modules

Statistic	Explanation
Threshold (% Full)	The threshold for queue capacity that determines whether additional frames are counted toward **Frames Enqueued Above Threshold** or **Frames Enqueued Below Threshold**. If the number of frames in the queue exceeds this percentage of the queue capacity, the **Frames Enqueued Above Threshold** count is increased when another frame is added to the queue. If the number of frames in the queue does *not* exceed this percentage of the queue capacity, the **Frames Enqueued Below Threshold** count is increased when another frame is added to the queue.
Frames Enqueued Above Threshold	The total number of frames that were added to the queue when the number of frames in the queue exceeded the **Threshold (% Full)**. This number is cumulative since the last reset of the statistics.
Frames Enqueued Below Threshold	The total number of frames that were added to the queue when the number of frames in the queue did *not* exceed the **Threshold (% Full)**. This number is cumulative since the last reset of the statistics.
Peak Size (% Full)	The highest percentage of queue capacity that the queue has reached since the last reset of the statistics.
Frames Dropped Age Out	The total number of frames that have aged out of the queue (dropped from the queue). This number is cumulative since the last reset of the statistics.
	2 of 3

Table 25-25. Ingress Queue Statistics for 10/100 Modules

Statistic	Explanation
Frames Dequeued	The total number of frames that were removed from the queue and forwarded to the egress physical port for transmission. This number is cumulative since the last reset of the statistics. The total number of frames that were forwarded from the queue. The total number of frames that were removed from the queue and
Bytes Dequeued	The total number of bytes that were removed from the queue and forwarded to the egress physical port for transmission. This number is cumulative since the last reset of the statistics.
	3 of 3

Egress Port Statistics for 10/100 Modules

The egress port statistics for 10/100 modules are reported per 12 ports, not for individual physical ports. For example, if you display the QoS statistics for port 20 on a 10/100 module, the port statistics that are displayed reflect the cumulative statistics of ports 13 through 24, not for just port 20.

Table 25-26 lists the port ranges for which egress port statistics are reported on 10/100 modules.

Table 25-26. Port Ranges for 10/100 Egress Port Statistics

Module	Port Ranges for which Port Statistics Are Reported
24-port 10/100 modules	• 1 set of statistics for ports 1–12. • 1 set of statistics for ports 13–24.
48-port 10/100 modules	• 1 set of statistics for ports 1–12. • 1 set of statistics for ports 13–24. • 1 set of statistics for ports 25–36 • 1 set of statistics for ports 37–48

Table 25-27 lists the statistics that the switch displays for ports on 10/100 modules.

Table 25-27. Egress Port Statistics for 10/100 Modules

Statistic	Explanation
Frames Arriving Switch Fabric	The number of frames that the module received from the switch fabric. This number is cumulative since the last reset of the statistics.
Frames Dropped Runt	The number of frames that the port range dropped because they were smaller than the minimum frame size of 60 bytes. This number is cumulative since the last reset of the statistics.
Frames Dropped Parity Error	The number of frames that the port range dropped because they had a parity error. This number is cumulative since the last reset of the statistics.
Frames Dropped Flood Rate Limit	The number of frames that the port range dropped due to flood rate limiting. This number is cumulative since the last reset of the statistics.

Egress Port Statistics for Gigabit Modules

The egress port statistics for gigabit modules are reported per individual port. Table 25-28 lists the statistics that the switch displays for ports on gigabit modules.

Table 25-28. Egress Port Statistics for Gigabit Modules

Statistic	Explanation
Frames Arriving Switch Fabric	The number of frames that the module received from the switch fabric. This number is cumulative since the last reset of the statistics.
Unicast Frames Enqueued	The number of unicast frames that the port queued. This number is cumulative since the last reset of the statistics.
Unicast Bytes Enqueued	The number of unicast bytes that the port queued. This number is cumulative since the last reset of the statistics.
Multicast Frames Enqueued	The number of multicast frames that the port queued. This number is cumulative since the last reset of the statistics.

1 of 2

Table 25-28. Egress Port Statistics for Gigabit Modules

Statistic	Explanation
Multicast Bytes Enqueued	The number of multicast bytes that the port queued. This number is cumulative since the last reset of the statistics.
Broadcast Frames Enqueued	The number of broadcast frames that the port queued. This number is cumulative since the last reset of the statistics.
Broadcast Bytes Enqueued	The number of broadcast bytes that the port queued. This number is cumulative since the last reset of the statistics.
Frames Dropped Runt	The number of frames that the port dropped because they were smaller than the minimum frame size of 60 bytes. This number is cumulative since the last reset of the statistics.
Frames Dropped Parity Error	The number of frames that the port dropped because they had a parity error. This number is cumulative since the last reset of the statistics.
Frames Dropped Huge (> 24KBytes)	The number of frames that the port dropped because they were larger than 24,000 bytes. This number is cumulative since the last reset of the statistics.
Frames Dropped Keep Logic	The number of frames that the port dropped because of parity errors, frame sizes less than 60 bytes, flood rate limiting, a queue being full, and frame sizes larger than 24,000 bytes. This number is cumulative since the last reset of the statistics.
Frames Dropped Packet Memory Full	The number of frames that the port dropped because the queue memory was full. This number is cumulative since the last reset of the statistics.
Frames Forwarded to MAC	The number of frames that were removed from the queue and forwarded to the MAC address. This number is cumulative since the last reset of the statistics.
	2 of 2

Buffer Statistics for Egress Queues

You can also display the amount of memory that is allocated to each egress queue. The allocated amount of memory is displayed in both the number of buffers and the number of bytes. Buffers are 128-byte units of memory that store packets in the queues.

Table 25-29 lists the buffer statistics that you can display for egress queues. These settings are determined by the switch. You can display them but not change them.

Table 25-29. Buffer Statistics for Egress Queues

Statistic	Explanation
Total Queue Memory (buffers)	The total number of buffers that are allocated to all queues on the port.
Total Queue Memory (bytes)	The total number of bytes that are allocated to all queues on the port.
Age Out Interval	The maximum number of milliseconds that a queue can store a packet. If the switch does not forward the packet before this time period expires, the queue drops the packet.
Reserved Allocation (buffers)	The number of buffers that are allocated to a specific queue. These buffers are dedicated to the queue and cannot be used by another queue.
Reserved Allocation (bytes)	The number of bytes that are allocated to a specific queue. These bytes are dedicated to the queue and cannot be used by another queue.
Maximum Allocation (buffers)	The maximum number of buffers that can be allocated to a specific queue if extra buffers are available (not being used by another queue).
Maximum Allocation (bytes)	The maximum number of bytes that can be allocated to a specific queue if extra bytes are available (not being used by another queue).
Minimum Free Pool After Allocation (buffers)	The minimum number of empty buffers that the port must have available at all times. If queues exceed their **Reserved Allocation**, they can consume buffers in their **Maximum Allocation** only until the pool of empty buffers reaches this number.

This section contains procedures for the following tasks:

- Displaying QoS Statistics
- Resetting the QoS Statistics
- Displaying the Buffer Settings for Egress Queues

Displaying QoS Statistics

Use the **show port queue counters** to display QoS statistics. The syntax of this command is:

> show port queue counters {<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] {ingress | egress | all} [queue <*queue*>]

Table 25-30. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, the switch displays the QoS statistics for all ports on the module that you specify.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py*. For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the switch displays the QoS statistics for the port or range of ports that you specify.
{ingress | egress | all}	The direction of traffic that you want to view the Qos Statistics for. • Enter **ingress** to view the QoS statistics for ingress queues. • Enter **egress** to view the QoS statistics for egress queues. • Enter **all** to view the QoS statistics for both ingress and egress queues.
<*queue*>	The queue number, which can range from 0 to 7. If you do not specify a queue number, the switch displays all QoS statistics for the port or module.

For more information about the QoS statistics, see "Setting the Priority of a Physical Port" earlier in this chapter.

Resetting the QoS Statistics

Use the **reset port queue counters** to reset the queue statistics to 0. The syntax of the command is:

>reset port queue counters {<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] {ingress | egress | all} [queue <*queue*>]

Table 25-31. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, the switch resets the the QoS statistics for all ports on the module that you specify.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py.* For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the switch resets the QoS statistics for the port or range of ports that you specify.
{ingress | egress | all}	The direction of traffic that you want to reset the QoS statistics for. • Enter **ingress** to view the QoS statistics for ingress queues. • Enter **egress** to view the QoS statistics for egress queues. • Enter **all** to view the QoS statistics for both ingress and egress queues.
<*queue*>	The queue number, which can range from 0 to 7. If you do not specify a queue number, the switch resets the QoS statistics for all queues on the port.

Displaying the Buffer Settings for Egress Queues

Use the **show port queue buffer** command to display the amount of memory that is assigned to each queue.

The syntax of the command is:

`(configure)#` **show port queue buffer {{<*mod-num*> | <*mod-swport-range*>} [..., {<*mod-num*> | <*mod-swport-range*>}] | all-ports}**

Table 25-32. Keywords, Arguments, and Options

Keyword, Argument or Option	Definition
<*mod-num*>	The slot number of a module. If you specify <*mod-num*>, the switch displays the number of packet buffers that are allocated to the egress queues on all ports on the module that you specify.
<*mod-swport-range*>	The slot number of a module, and, either a port number, or a range of port numbers having the format *Px-Py.* For example: • To specify port 1 on the module in slot 3, enter **3/1**. • To specify ports 1 through 5 on the module in slot 3, enter **3/1-5**. If you specify <*mod-swport-range*>, the switch displays the number of packet buffers that are allocated to the egress queues on the port or range of ports that you specify.
{all-ports}	All ports in the chassis. If you specify **all-ports**, the switch displays the number of packet buffers that are allocated to the egress queues on all ports in the chassis.

A Upgrading the Application Software

Overview

Upgrading the switch software involves the following steps:

1. Backing Up the Current Software
2. Backing Up the Previous Configuration
3. Downloading Application Software
4. Setting the Startup Image
5. Synchronizing the Active and Standby Supervisor Modules
6. Resetting the Active Supervisor
7. Resetting the Standby Supervisor
8. Verifying the Upgrade

This appendix contains procedures for each of these steps.



CAUTION:

Before downgrading the switch from v6.*x* application software to v5.*x*, you must initialize NVRAM (**nvram initialize** command). If you do not initialize NVRAM before downgrading the application software, the switch may not respond after you reset it.

If the switch has multiple VLANs statically bound to hunt group ports, and you upgrade from application software earlier than v5.3, you must delete the hunt groups and set them up again. For information on how to perform this task, see "Upgrading a Switch with Multiple VLANs Statically Bound to Hunt Group Ports" later in this appendix.

For more information about the CLI commands that are mentioned in this appendix, see *Command Reference Guide for the Avaya P580 and P882 Multiservice Switches, Software Version 6.1*

Backing Up the Current Software

The Avaya Multiservice switches have two memory locations for storing the embedded switch software: APP1 and APP2. These two memory locations make it possible to store the current software in one APP location and download the new software to the other APP location.

Saving the current software in one APP location ensures that you could run this earlier version of software, if you encounter any problem with the new, downloaded software.

CAUTION:

When upgrading the switch from v5.*x* application software to v6.0 or later, you must download v6.*x* to the APP2 memory location. If you download v6.*x* to APP1, and store v5.*x* software in APP2, the 6.*x* software is corrupted. Once v6.*x* software is successfully stored in APP2, you can download v6.*x* to APP1.

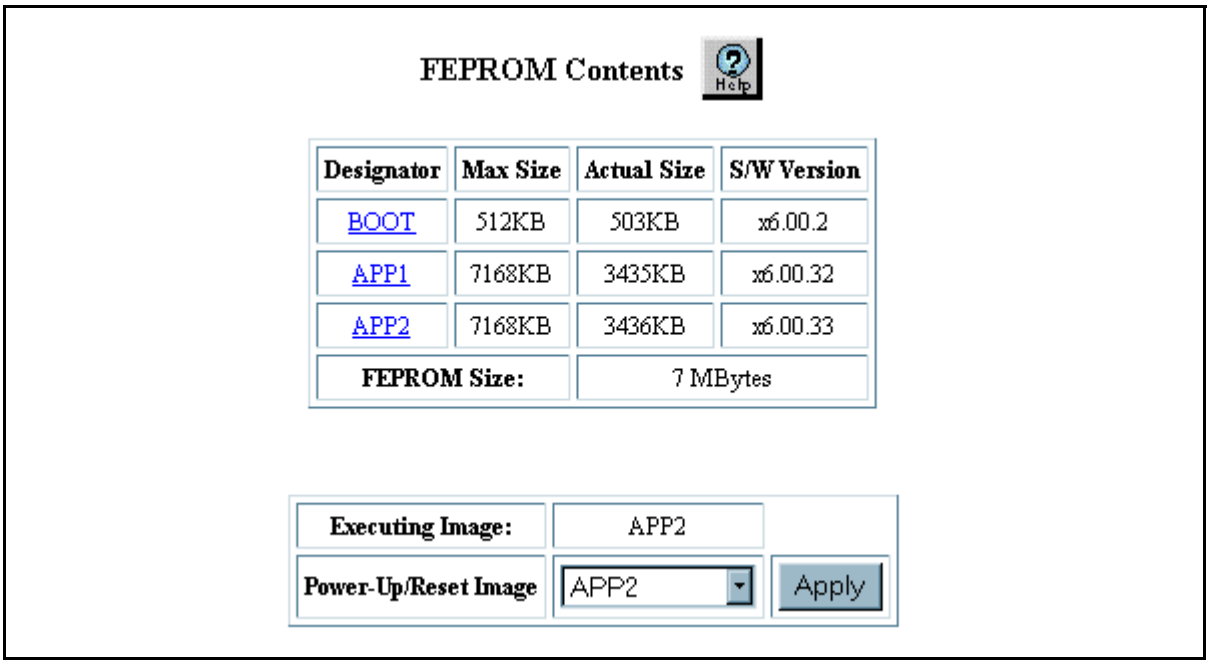
Web Agent Procedure

To determine the versions of application software that are stored in APP1 and APP2:

1. In the navigation pane, expand the **System** > **Memory** folders.
2. Select **FEPROM**.

   The FEPROM Contents Web page is displayed in the content pane (Figure A-1). The **S/W Version** field displays the version of application software that is stored in APP1 and APP2.

Figure A-1. FEPROM Contents Web page



3. Ensure that APP1 contains the software that you want to back up.

   When you download the new application software, you must download it to the APP2 memory location.

CLI Command

To determine the APP location from which the switch loads the previous version of software, use the following command:

> show flash

This command displays information about the application software that is currently stored in APP1 and APP2.

Backing Up the Previous Configuration

Before you upgrade the switch software, Avaya recommends that you back up the previous configuration to a file or directory on a TFTP server. You can use either the Web Agent or CLI to perform this task.

Web Agent Procedure

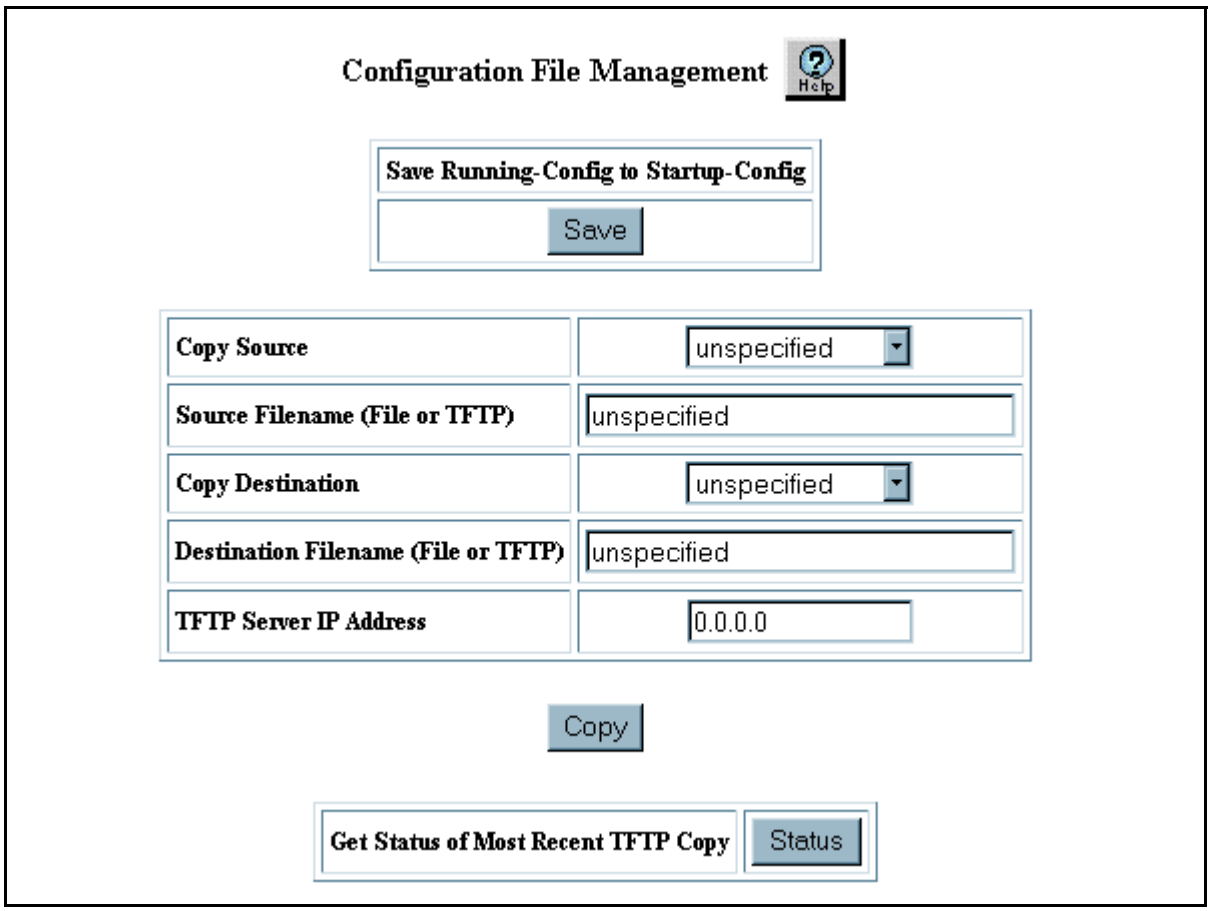
To back up the previous configuration by using the Web Agent:

1. In the navigation pane, expand the **System** > **Configuration** > **Configuration Files** folders.

2. Select **File Management**.

   The Configuration File Management Web page is displayed in the content pane (Figure A-2).

Figure A-2. Configuration File Management Web page



3. In the **Save Running-Config to Startup-Config** field, select **Save**.
4. In the **Copy Source** field, select **Startup-Config**.

*** Note:** The Source File name may be left unspecified

5. In the **Copy Destination** field, select **TFTP Server**.

*** Note:** You must specify a Destination file name.

6. In the **TFTP Server IP Address** field, enter the IP address of the TFTP server to which you want to copy the startup configuration file.
7. Click **Copy**.

CLI Command

To back up the current configuration to a TFTP server, use the following command:

copy startup-config tftp <*filename_opt_path*> <*tftp-server*>

Downloading Application Software

You can download new application software from a TFTP server to an APP memory location by using either the Web agent or the CLI.

CAUTION:

When upgrading the switch from v5.x application software to v6.0 or later, you must download v6.*x* to the APP2 memory location. If you download v6.*x* to APP1, and store v5.*x* software in APP2, the 6.*x* software is corrupted. Once v6.*x* software is successfully stored in APP2, you can download v6.*x* to APP1.

Web Agent Procedure

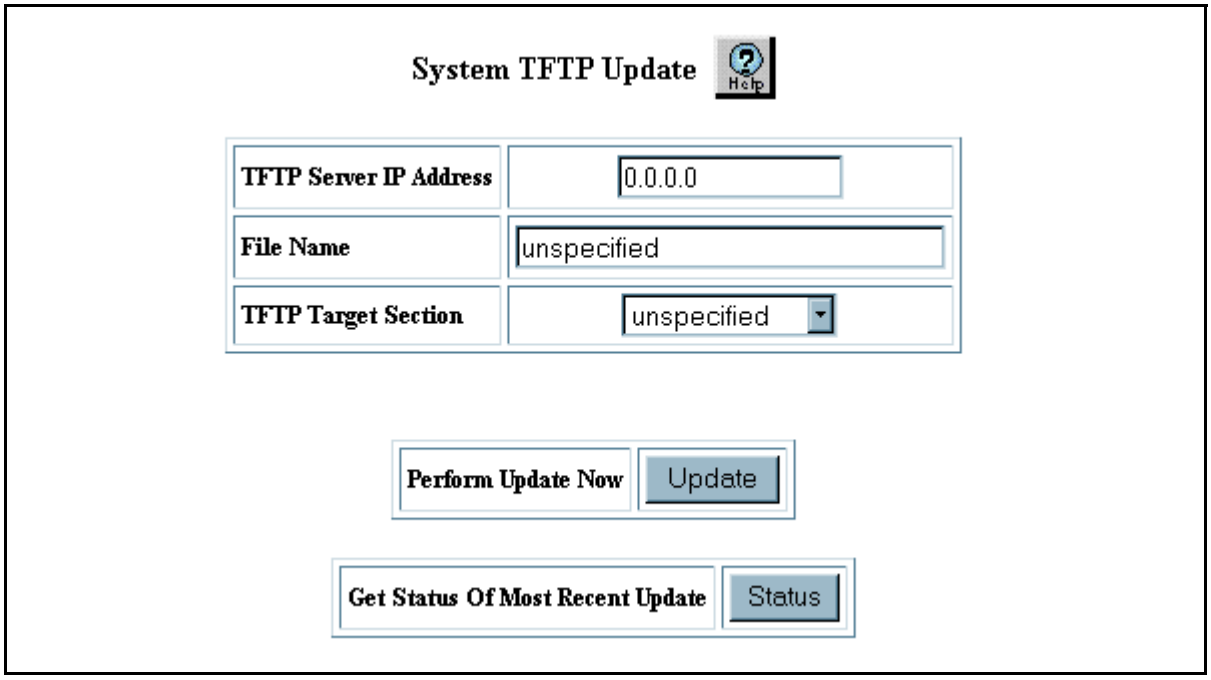
To download new application software from a TFTP server to an APP memory location by using the Web Agent:

*** Note:** The switch does not support Exceed TFTP server software.

1. In the navigation pane, expand the **System** > **Configuration** > **TFTP Update** folders.
2. Select **System**.

   The System TFTP Update Web page is displayed in the content pane (Figure A-3).

Figure A-3. System TFTP Update Web page



3. In the **TFTP Server IP Address** field, enter the IP address of the TFTP server on which the new application software is stored.
4. In the **File Name** field, enter the name of the binary file that you want to download.
5. In the **TFTP Target Section** field, select the memory location to which you want to download the new application software.
6. In the **Perform Update Now** field, select **Update**.
7. In the **Get Status of Most Recent Update** field, select **Status** to verify that the new image downloaded properly.

   The status of the update is displayed in the TFTP Update Status pane.

CLI Command

To download new application software from a TFTP server to an APP memory location by using the CLI, use the following command:

`(configure)#` **copy tftp flash {app1 | app2} {cardapp1 | cardapp2}** <***image_opt_path***> <***ip-addr***>

Setting the Startup Image

After you download the new software, you must set the switch to load the new software at startup. You can use either the Web Agent or CLI to perform this task.

Web Agent Procedure

To set the switch to load the new software at startup:

1. In the navigation pane, expand the **System** > **Memory** folders.
2. Select **FEPROM**.

   The FEPROM Contents Web page is displayed in the content pane (Figure A-1).
3. In the **Power Up/Reset Image** field, select the APP location to which you downloaded the new application software.
4. Select **APPLY**.

CLI Command

To set the switch to load the new software at startup, use the following CLI command:

`(configure)#` **boot system flash {app1 | app2 | cardapp1 | cardapp2}**

Synchronizing the Active and Standby Supervisor Modules

If the switch has a redundant supervisor module, you must synchronize the active and standby supervisor modules after you upgrade the application software on the active supervisor module. During this synchronization, the new boot code and application software are copied to the standby supervisor module.

For information on how to synchronize the active and standby supervisor modules, see "Synchronizing the Active and Standby CPUs" in Chapter 10, "Configuring Redundancy Options."

If both supervisors are running v6.0 or later application software, the switch performs the following checks after you synchronize the active and standby supervisors:

- Version of application software.
- Startup image (app 1, app 2, cardapp 1, or cardapp 2)

If any of these settings changed during the synchronization, the standby supervisor module automatically resets so that it is failover-ready. In earlier versions of software, you had to manually reset the standby supervisor, if these settings changed.

***Important:** If you are updating the boot code, you must manually reset the standby supervisor after synchronizing it with the active supervisor.

*** Note:** For this automatic reset feature to work, the active and standby supervisors must *both* be running v6.0 or later application software. The first time that you synchronize the supervisor modules after upgrading the active supervisor from a v5.*x* version of software to v6.*x*, you must manually reset the standby supervisor.

Resetting the Active Supervisor

After setting the switch to load the new software at startup, you must reset the switch. You can use either the Web Agent or CLI to perform this task.

If the switch has a redundant supervisor module, resetting the Active supervisor to load the new boot and application code forces the Active supervisor to Standby, and the Standby supervisor to the Active Supervisor mode. See "Resetting the Standby Supervisor" next to change the Supervisor modes back to their original state.

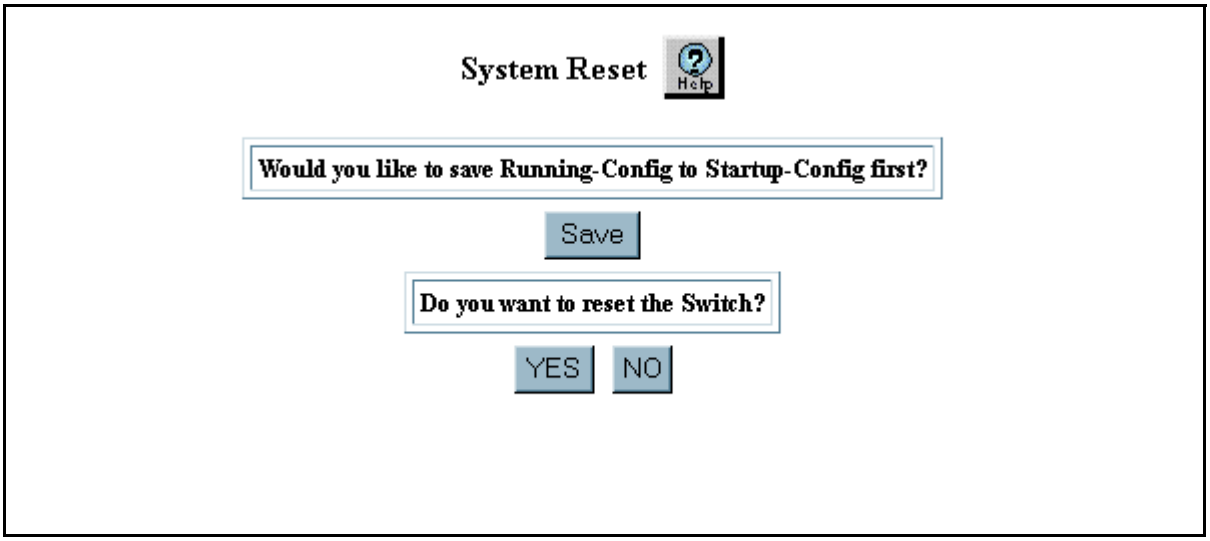
Web Agent Procedure

To reset the switch after downloading new application software:

1. In the navigation pane, expand the **System** folder.
2. Select **System Reset**.

   The System Reset Web page is displayed in the content pane (Figure A-4).

Figure A-4. System Reset Web Page



3. Select **Yes** in response to the question, **Do you want to reset the switch?**

   The switch resets and loads the new application software.

CLI Command

To reset the switch after downloading new application software, use the following CLI command:

```
(configure)# reset
```

Resetting the Standby Supervisor

The previous procedure, "Resetting the Active Supervisor," forces the Active Supervisor to Standby mode and the Standby Supervisor to the Active Supervisor mode. To recover back to the original state, with both supervisors running new code, reset the current Active supervisor module.

Web Agent Procedure

To reset the Active Supervisor to the Standby Supervisor:

1. In the navigation pane, expand the **System** folder.
2. Select **System Reset**.

   The System Reset Web page is displayed in the content pane (Figure A-4).

3. Select **Yes** in response to the question, **Do you want to reset the switch?**

CLI Command

To reset the switch after downloading new application software, use the following CLI command:

```
(configure)# reset
```

Verifying the Upgrade

To verify that the old Active supervisor is now the Standby supervisor, the LED display window (marquee) should scroll `Standby` and the software version.

Upgrading a Switch with Multiple VLANs Statically Bound to Hunt Group Ports

If the switch has multiple VLANs statically bound to hunt group ports and you want to upgrade from application software earlier than v5.3, you must:

1. Download the application software to which you want to upgrade.
2. Reset the switch to load this software.
3. Delete all hunt groups to which you want to statically bind multiple VLANs.
4. Use the **set vlan {<*vlan-id*> | name <*vlan-name*>} <*mod-swport-range*> [...,<*mod-swport-range*>]** CLI command to statically bind the VLANs to every port that you want to assign to a hunt group.
5. Assign the ports to hunt groups.
6. Save the running configuration to the startup configuration.

B Boot Mode

Boot Mode is a special mode for the Avaya P580 and P882 Multiservice switch. Boot mode is used to:

- Recover your password
- Download new operational code

This appendix contains the following sections:

- Accessing BOOT Mode
- Password Recovery
- Downloading New Operational Code

Accessing BOOT Mode

You can access BOOT mode during Power Up or when both operational images in APP1 and APP2 are corrupted.

Accessing BOOT Mode During Power Up

To access BOOT mode during power up:

1. Power-Off the Switch.
2. Connect a console to the supervisor Module (See Chapter 2, “Initial Configuration” of the *Installation Guide for the Avaya P550R, P580, P880, and P882 Multiservice Switches.*)
3. Press and Hold the DISPLAY button on the supervisor module as you Power-on the switch until the boot prompt displays on the console.

Accessing BOOT Mode with Corrupted Operational Images

The Avaya P580 and P882 Multiservice switch will automatically come up in BOOT mode if both operational images are corrupted.

You will see the following output on your console:

```
Starting the boot system, please wait.
Initializing the event subsystem ... done
Initializing the platform ... done

Starting up threads ...
    Periodic Task
    Network Interface
    Front Panel Display
    Download
    Command Line Parser

Boot process complete - system is now operational.

Boot Software Revision: V5.3

Press Ctrl-P for previous command, Ctrl-N for next command


download   Download to FEPROM
feprom     Flash EPROM configuration
initLogins Restore default passwords
ip         IP Configuration
manuf      Get/Set MAC address
nvram      Initializes NV RAM-diags run upon reboot
peek       Read from memory
ping       Ping host [tries] [delay]
poke       Write to memory
reset      Reset System

Boot>
```

Password Recovery

The following command is used to reset the switch to the factory default passwords.

initLogins

Syntax

```
Boot> initLogins set <0=no change, 1=factory default>
```

Description

Using a value of zero will retain the current passwords in NVRAM upon resetting the switch. Using a value of one will restore the factory default passwords in NVRAM upon resetting the switch.

Downloading New Operational Code

To download new operational code in BOOT mode into APP1 or APP2.

*** Note:** The Avaya P580 and P882 Multiservice switches do not have an IP address assigned to the console ethernet port when it comes up in BOOT mode. Therefore, a user must assign an IP address (and default gateway) to the ethernet console port prior to attempting the TFTP download of a new operational image.

1. Ensure the ethernet console port is connected to the network for a TFTP transfer.

2. Enter the following command to assign an IP address to the ethernet console port:

`Boot>` **ip address** <***ip address***> <***ip mask***>

3. Enter the following command to assign a default gateway to the ethernet console port (if necessary):

`Boot>` **ip default_gateway** <***ip address***>

4. Enter the following command to download new operational code to APP1 or APP2:

`Boot>` **download {app1 | app2}** <***serverIP***> <***filename***>

5. Enter the following command to set the switch to boot from the appropriate image stored in APP1 or APP2

`Boot>` **feprom set {app1 | app2}**

6. Enter the following command to reset the switch:

`Boot>` **reset**

C Supported MIB Groups

This appendix lists, by protocol, public and private MIBs that are supported by the Avaya P580 and P882 Multiservice switches. MIBs are categorized as follows:

- AppleTalk
- Application Software Management
- ATM Uplink
- Bridging
- DVMRP
- General Private MIBs
- IGMP
- IP
- IPX
- Monitoring
- SNMP

AppleTalk

Standard MIB rfc1243.mib, AppleTalk MIB.

Private MIB cjnatalk.mib, AppleTalk routing MIB.

Application Software Management

Private MIB Load.mib, MIB that defines upload, download, and copy of application software and configuration information.

ATM Uplink

Standard MIB lecmib.mib, MIB for LAN Emulation Client Management, as defined by the ATM Forum.

Private MIB avaya1483.mib, MIB for multi-protocol over ATM encapsulation, as defined in RFC1483.

Bridging

Standard MIBs

- rfc1493.mib, Bridge MIB.
- draft-ietf-bridge-rstp-mib-03.mib, MIB for IEEE 802.1w Rapid Spanning Tree Protocol.

DVMRP

Private MIB cjndvmrp.mib, DVMRP MIB.

General Private MIBs

- avayagen.mib, Avaya top-level MIB definitions.
- aveisprod.mib, Product OIDs for the Avaya EIS division.
- cjnroot.mib, Avaya P580 and P882 routing MIB top-level architecture.
- cjnSwitchRoot.mib, Avaya P580 and P882 routing MIB top-level architecture.
- genlic.mib, Avaya common license MIB
- ProminetMIB.txt, Main MIB for management of the P580 and P882. Provides chassis information, physical configuration, L2 configuration.

IGMP

Standard MIB draft-ietf-idmr-igmp-mib-07.mib, IGMP Management MIB.

Private MIB cjnipigmp.mib, IGMP MIB.

IP

Standard MIBs

- rfc2096.mib, IP Forwarding Table MIB.
- rfc2863.mib, Interface MIB.
- ianaiftype.mib, MIB for different interface types.
- rfc1213.mib, MIB II.
- rfc2011.mib, SNMPv2 MIB for IP.
- rfc2012.mib, SNMPv2 MIB for TCP.
- rfc2013.mib, SNMPv2 MIB for UDP.
- rfc1850.mib, OSPF MIB.
- rfc1724.mib, RIP-2 MIB extension.
- draft-ietf-vrrp-mib-04.mib, VRRP Version 2 Protocol.

Private MIBs

- l3fwd.mib, Layer 3 switching group added to Prominet.txt.
- cjnipifmgmt.mib, IP Interface Management MIB.
- cjnipv4.mib, IP version 4 mib extensions.
- cjnipv4serv.mib, IP version 4 services MIB.
- cjnospf.mib, OSPF mib.
- cjniprip.mib, RIP MIB.
- cjnpolicycap.mib, MIB that defines policy capabilities for the Lightweight Directory Access Protocol (LDAP).
- cjniplrrp.mib, Lucent Router Redundancy Protocol MIB.
- policy-mib.mib, MIB for access lists.

IPX

Private MIBs

- cjnipxifmgmt.mib, IPX interface management MIB.
- cjnipx.mib, IPX routing MIB.
- cjnipxrip.mib, IPX RIP MIB.
- cjnipxsap.mib, IPX SAP MIB.
- Ipx.mib, MIB for IPX.
- IpxRipSap.mib, MIB for RIP and SAP in IPX

Monitoring

Standard MIBs

- rfc1757.mib, Remote network monitoring MIB.
- rfc1513.mib, Token Ring RMON MIB.
- rfc2856.mib, MIB for HCRMON.
- draft-ietf-rmonmib-hcrmon-10.mib, Internet Draft RMON for High Capacity Networks.
- rfc2021.mib for the Trap Destination Table, SNMPv2 MIB for RMON2.
- rfc2613.mib, Module for managing remote monitoring device implementations for switched networks.
- rfc1398.mib, Ethernet-like MIB.

Private MIBs

- cjnportcopyext.mib, Port Copy Extensions MIB. Extension to RFC 2613 Port Copy Definitions.
- genutil.mib, Avaya generic utilization MIB, which includes CPU utilization and memory utilization.
- cjnutil.mib, Avaya specific utilization mib, which includes forwarding engine utilization.

SNMP

Standard MIBs

- snmpv2.mib, MIB for SNMPv2.
- snmp-framework.mib, SNMP management architecture MIB, RFC 3411.
- snmp-mpd.mib, MIB for message processing and dispatching, RFC 3412.
- snmp-usm.mib, Management information definitions for the SNMP user-based security model, RFC 3414.
- snmp-vacm.mib, Management information definitions for the view-based access control model for SNMP, RFC 3415.
- snmp-community.mib, Management information definitions for the SNMP Community Table, RFC 2576.
- rfc3413.mib, Management information definitions for Target, Notification, and Proxy MIBs.

D FCC Notice

FCC Notice — Class A Computing Device:

This equipment generates, uses, and may emit radio frequency energy. The equipment has been type tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules, which are designed to provide reasonable protection against such radio frequency interference. Operation of this equipment in a residential area may cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference. Any modifications to this device - unless expressly approved by the manufacturer - can void the user's authority to operate this equipment under part 15 of the FCC rules.

VCCI Notice — Class A Computing Device:

This equipment is in the Class A category (information equipment to be used in commercial and/or industrial areas) and conforms to the standards set by the Voluntary Control Council for Interference by Data Processing Equipment and Electronic Office Machines aimed at preventing radio interference in commercial and/or industrial areas. Consequently, when used in a residential area or in an adjacent area thereto, radio interference may be caused to radios and TV receivers. Read the instructions for correct handling.

CE Notice — Class A Computing Device:

Warning!

This is a Class A product. In a domestic environment, this product may cause radio interference, in which case the user may be required to take adequate measures.

Achtung!

Dieses ist ein Gerät der Funkstörgrenzwertklasse A. In Wohnbereichen können bei Betrieb dieses Gerätes Rundfunkstörungen auftreten, in welchen Fällen der Benutzer für entsprechende Gegenmaßnahmen verantwortlich ist.

Avertissement!

Cet appareil est un appareil de Classe A. Dans un environnement résidentiel cet appareil peut provoquer des brouillages radioélectriques. Dans ce cas, il peut être demandé à l'utilisateur de prendre les mesures appropriées.

Index

Numerics

10/100 port
  auto negotiation speed/duplex advertisement, 8-26, 8-27
10/100 port parameter
  auto-negotiation mode, 8-26
  category, 8-26
  duplex mode, 8-26
  flow control mode, 8-26
  name, 8-26
  port PACE priority, 8-27
  rate limit burst size, 8-27
  rate limit mode, 8-27
  rate limit rate, 8-27
  speed mode, 8-26
10Base-T crossover patch cables, 2-5
10-port 100BASE-FX, 1-11
12-port 10/100BASE-TX, 1-11
3Com mapping table
  switch port configuration parameters, 8-33

A

AARP invalid PDU
  AppleTalk global statistics, 19-24
AARP reply rx
  AppleTalk global statistics, 19-24
AARP reply tx
  AppleTalk global statistics, 19-24
AARP req rx
  AppleTalk global statistics, 19-24
AARP req tx
  AppleTalk global statistics, 19-24
access list
  AppleTalk NBP filter parameters, 19-15
  AppleTalk zone filter parameters, 19-20
access list name
  IP access list parameters, 13-8, 13-10
access list rules
  IP interface, 13-6
access lists
  IP interface, 13-1
access rule index
  IP access list parameters, 13-8, 13-10
access type
  IP access list parameters, 13-8, 13-10
Accessing BOOT Mode, B-1
active alarm table
  viewing, 21-2, 21-3
active alarms
  checking, 21-2
active backpressure, 8-26
add IP interface parameters
  admin. state, 12-8
  ICMP redirect, 12-10
  interface, 12-8
  IP routing, 12-9
  MAC format, 12-9
  mask, 12-8
  multicast protocol, 12-10
  NetBIOS UDP rebroadcast, 12-10
  network address, 12-8
  OSPF, 12-9
  proxy ARP, 12-10
  RIP, 12-9
  VLAN, 12-8
  VRRP, 12-10
adding a user, 2-14, 2-17, 2-19, 2-21
adding address forward table entries manually, 9-19
Adding Entries to the AFT Manually Using the Web Agent, 9-19
addr owner override
  VRRP configuration parameters, 12-78
address age time
  definition, 9-9
address forwarding table, 9-1
  adding entries manually, 9-19
  capacity, 9-13
address forwarding table parameters
  group, 9-17
  index, 9-17
  MAC address, 9-17
  persistence, 9-18
  port, 9-17
  priority, 9-17
  status, 9-18
  TblInst, 9-17
  valid, 9-17
address mask
  DVMRP interface statistical parameters, 12-60
address table instance

delete all learned entries, 9-13
delete invalid learned entries, 9-13
address table instance entry type
learned, 9-12
management, 9-12
multicast, 9-12
self, 9-12
address table instance hash table
auto increment, 9-12
size, 9-12
address table instance parameters
bucket info, 9-12
entry type, 9-12
entry validity, 9-12
hash table, 9-12
total number of entries, 9-11
VLAN association, 9-11
address table sizes
controlling reconfiguration, 9-13
address table, age timer
configuring, 9-9
address table, super age timer
configuring, 9-9
admin state
AppleTalk interface parameters, 19-6
IPX interface, 16-8
admin. state
add IP interface parameters, 12-8
adv. address
IRDP parameters, 12-84
advertisement
OSPF summaries parameters, 15-17
advertisement life time
IRDP parameters, 12-84
advertisement timer
VRRP configuration parameters, 12-77
advertisements received
VRRP statistical parameters, 12-80
advertisements sent
VRRP statistical parameters, 12-80
age
LSA detail, 15-27
OSPF link state database parameter, 15-26
age interval
Layer 3 cache configuration, 22-3
age timer
buffer management table parameters, 24-5
aged entries
forwarding cache (FE) parameters, 22-11
ageout
persistence entry, 9-18
aggregating bandwidth, 6-13
aging
Layer 3 cache configuration, 22-3
aging interval multiplier
IPX SAP interface parameters, 18-3
alarms
log size, setting, 21-10
system events, activating, 21-10
all module ports configuration, 8-28
allow learning
switch port configuration parameters, 8-32
always be group membership querier
IGMP interface parameters, 12-34
APP1
CPU redundancy configuration, 10-17
APP2
CPU redundancy configuration, 10-17
AppleTalk
creating an NBP filter, 19-14
ethernet versions supported, 19-1
event class parameters, 21-6
AppleTalk ARP cache table
viewing, 19-29, 19-30
AppleTalk ARP cache table statistics
interface, 19-30
MAC address, 19-30
network range, 19-30
node, 19-30
TTL, 19-30
type, 19-30
AppleTalk global routing
enabling, 19-4
AppleTalk global statistics
AARP invalid PDU, 19-24
AARP reply rx, 19-24
AARP reply tx, 19-24
AARP req rx, 19-24
AARP req tx, 19-24
Bcast error, 19-24
checksum error, 19-24
config addr error, 19-25
config zone error, 19-25
DDP forward counter, 19-24
DDP input counter, 19-24
DDP local counter, 19-24
DDP output counter, 19-24
DDP output long, 19-24
DDP output short, 19-24
echo reply rx, 19-24
echo req rx, 19-24
echo req tx, 19-24
no client, 19-24
no route, 19-24
RTMP RDR Rx, 19-25
RTMP RDR Tx, 19-25

RTMP Rq Rx, 19-25
RTMP Rq Tx, 19-25
RTMP Rsp Rx, 19-25
RTMP Rsp Tx, 19-25
Short PDU in error, 19-24
too long, 19-24
too short, 19-24
TTL expired, 19-24
viewing, 19-23, 19-25
Zip Ext Reply Rx, 19-25
Zip Ext Reply Tx, 19-25
Zip GNI Rq Rx, 19-25
Zip GNI Rq Tx, 19-25
Zip GNI Rsp Rx, 19-25
Zip GNI Rsp Tx, 19-25
Zip Query Rx, 19-25
Zip Query Tx, 19-25
Zip Reply Rx, 19-25
Zip Reply Tx, 19-25
AppleTalk interface
deleting, 19-9
AppleTalk interface parameters
admin state, 19-6
encapsulation type, 19-5
interface, 19-5
metric, 19-5
network number, 19-6
network range end, 19-6
network range start, 19-6
node, 19-6
VLAN, 19-5
zone, 19-6
AppleTalk interface statistics
interface, 19-26
network number, 19-26
network range, 19-26
node, 19-26
seed, 19-26
state, 19-26
AppleTalk interface, creating, 19-4, 19-7
AppleTalk interface, editing, 19-7, 19-9
AppleTalk NBP filter
editing, 19-16, 19-17
AppleTalk NBP filter parameters
access list, 19-15
interface, 19-16
name, 19-15
type, 19-16
AppleTalk NBP table
viewing, 19-33
AppleTalk NBP table parameters
index, 19-33
interface, 19-33
name, 19-33
type, 19-33
zone, 19-33
AppleTalk overview, 19-1
AppleTalk protocols
supported, 19-2
AppleTalk route table
viewing, 19-27
AppleTalk route table statistics
current number of routes, 19-29
interface, 19-28
metric, 19-27
network range, 19-27
next hop, 19-28
owner, 19-28
peak number of routes, 19-29
state, 19-27
zones, 19-28
AppleTalk route table statistics, viewing, 19-28
AppleTalk routing overview, 19-1
AppleTalk routing, configuring, 19-3
AppleTalk static route
creating, 19-10
deleting, 19-13, 19-14
editing, 19-12, 19-13
network number, 19-11
network range end, 19-11
network range start, 19-11
node, 19-11
type, 19-11
zone, 19-11
AppleTalk statistics
viewing, 19-22
AppleTalk zone filter
creating, 19-19
editing, 19-20
AppleTalk zone filter parameters
access list, 19-20
interface, 19-20
name, 19-20
type, 19-20
AppleTalk zone table
viewing, 19-31
AppleTalk zone table statistics
current number of zones, 19-32
index, 19-31
name, 19-31
network range, 19-31
peak number of zones, 19-32
AppleTalk zone table statistics, viewing, 19-32
application
multicast session client port parameters, 20-15
multicast session parameters, 20-13

router port display parameters, 20-9
area
LSA detail, 15-27
OSPF interface parameters, 15-10
OSPF summaries parameters, 15-17
OSPF virtual link parameters, 15-14
area ID
OSPF area parameters, 15-6
OSPF link parameters, 15-21
OSPF link state database parameter, 15-26
OSPF link state database search parameters, 15-25
OSPF statistical parameters, 15-20
area type
OSPF area parameters, 15-6
ARP cache
searching, 12-49
ARP cache entries, 12-13
ARP cache search parameters
IP address, 12-50
VLAN, 12-50
AS border route
OSPF global configuration parameters, 15-4
AS border Rtrs
OSPF statistical parameters, 15-20
ASBR, 15-4
ATM
overview, 1-10
attached router ID 1 & 2
LSA detail, 15-28
auth key
RIP interface parameters, 14-5
VRRP configuration parameters, 12-77
auth type
RIP interface parameters, 14-5
authentication
OSPF interface parameters, 15-11
OSPF virtual link parameters, 15-14
authentication failure
event class parameters, 21-6
authentication key
OSPF virtual link parameters, 15-14
Authorization Type, 12-77
auto increment
address table instance hash table parameter, 9-12
auto increment HT size
VLAN configuration parameters, 6-10
auto negotiation speed/duplex advertisement
10/100 port, 8-26, 8-27
auto-creation of virtual links
OSPF global configuration parameters, 15-4
automatic client pruning
automatic client pruning, 20-7
automatic router pruning
automatic router pruning, 20-7
automatic session pruning
automatic session pruning, 20-7
automatic VLAN creation, 6-4, 6-7
Cisco Catalyst 5000 parameters, 8-35
switch port configuration parameters, 8-31
auto-negotiation
Fast Ethernet module, 8-24
auto-negotiation mode
10/100 port parameter, 8-26

B

bad advertisements received
VRRP statistical parameters, 12-80
bad packets received
IPX SAP interface statistical parameters, 17-9, 18-13
RIP statistical parameters, 14-10, 14-12
bad routes received
RIP statistical parameters, 14-10, 14-12
baud rate
console serial port settings, 2-28
Bcast error
AppleTalk global statistics, 19-24
BDR address
OSPF link parameters, 15-21
BDR choice
OSPF neighbors parameters, 15-23
bind to all
Virtual LAN binding options, 8-36
bind to received
Virtual LAN binding options, 8-36
Binding a Port to more than one VLAN with the CLI, 6-6
binding type
VLAN switch port table parameters, 6-12
blocking state (spanning tree bridge port), 7-16
BOOTP relay agent, 12-24, 12-25
BOOTP/DHCP in discards
IP routing global statistics, 12-46
BOOTP/DHCP in hops exceeded
IP routing global statistics, 12-46
BOOTP/DHCP in requests
IP routing global statistics, 12-45
BOOTP/DHCP in responses
IP routing global statistics, 12-46
BOOTP/DHCP out requests
IP routing global statistics, 12-46
BOOTP/DHCP out responses
IP routing global statistics, 12-46
BOOTP/DHCP relay agent

IP global configuration parameters, 12-12
border Rtrs
OSPF statistical parameters, 15-20
bridge forward delay
spanning tree bridge level parameters, 7-12
bridge hello time
spanning tree bridge level parameters, 7-12
bridge max age
spanning tree bridge level parameters, 7-12
bridge port
spanning tree bridge port parameters, 7-15
bridge port parameters
bridge port, 7-15
designated bridge, 7-16
designated cost, 7-16
designated port, 7-16
designated root, 7-16
forward transitions, 7-16
name, 7-15
port, 7-15
state, 7-16
bridge status
event class parameters, 21-6
broadcast packets
Ethernet interface statistical parameters, 23-4
bucket info
address table instance parameters, 9-12
buffer and queue management, overview, 1-20
buffer detailed configuration parameters
memory, 24-5
buffer management parameters
fabric port buffers, 24-4
physical port buffers, 24-4
buffer management table parameters
age timer, 24-5
congestion drops, 24-6
high and normal overflow drops, 24-6
high and normal stale drops, 24-6
high priority allocation, 24-6
high priority service ratio, 24-6
priority threshold, 24-6
buffers
managing, 24-2, 24-7
bytes
Ethernet interface statistical parameters, 23-4

C

Cajun-Service-Type required
RADIUS configuration parameters, 4-28
category
10/100 port parameter, 8-26
gigabit port parameters, 8-15, 8-19
CGMP packet reception statistics
CGMP snooping parameters, 20-34
CGMP snooper, configuring/viewing
configuring CGMP snooper, 20-33, 20-34, 20-35
CGMP snooping parameters
CGMP packet reception statistics, 20-34
enable state, 20-34
intelligent multicast session statistics, 20-34
checking active alarms, 21-2
checksum
IPX datagram fields, 16-2
LSA detail, 15-27
OSPF link state database parameter, 15-26
checksum error
AppleTalk global statistics, 19-24
Cisco Catalyst 5000 parameters
port VLAN, 8-34
Cisco Catalyst 5000, parameters
automatic VLAN creation, 8-35
trunk mode, 8-34
VLAN binding, 8-34
VTP snooping, 8-35
CLI, 9-19
regaining access, 2-32
CLI configuration
supervisor module, 2-2
client port
multicast session parameters, 20-13
client pruning timer
client pruning timer, 20-7
clients per VLAN, displaying, 20-31, 20-32
collisions, 23-7
Command, 19-25
comparison value
Layer 3 route cache search parameters, 22-8
config addr error
AppleTalk global statistics, 19-25
config zone error
AppleTalk global statistics, 19-25
configuration
event class parameters, 21-6
invalid copies, 2-44
saving, A-3
configuration file management
copy destination, 2-43
copy source, 2-42
destination filename, 2-43
save running configuration to startup configuration, 2-42
source filename, 2-42
TFTP server IP address, 2-43
configuration files, 2-38
configuration revision number

VTP snooping parameters, 6-31
configure, 20-9, 20-10
configure ports
gigabit module, 8-12, 8-16, 8-21, 8-27
configure source
port mirroring information parameters, 23-9
configure static router ports
static router ports, 20-9, 20-10
configuring
switch port, 8-29
configuring a DNS client, 11-1
configuring a Fast Ethernet module, 8-21
configuring a redundant CPU, 10-10
configuring a redundant element, 10-9
configuring all VLAN ports, 6-10
configuring an IPX router, 16-4, 16-5
configuring AppleTalk routing, 19-3
configuring event logs, protocol, 21-8, 21-10
Configuring Event Notification, 21-4
configuring event notification, 21-4
configuring flow control mode
gigabit module, 8-14, 8-18, 8-24
configuring IGMP, 12-31
configuring IGMP interfaces, 12-32
Configuring IGMP Snooping, 20-20, 20-21, 20-26
configuring IGMP snooping, 20-21, 20-22
configuring IP global routing, 12-11
configuring IP interface access lists, 13-1
configuring IP interfaces, 12-6
Configuring IPX Routing Globally, 16-4
configuring IPX SAP interfaces, 18-1
Configuring LDAP Setting from the Web Agent, 12-86, 12-88
configuring multiple interfaces on a VLAN, 12-14
configuring physical port parameters
gigabit ports, 8-15, 8-19
configuring RADIUS client, 4-26
configuring RIP global configuration, 14-1
configuring RMON mirror port, 23-7
configuring service port
gigabit module, 8-14, 8-18, 8-23
configuring the address table age timer, 9-9
configuring the address table super age timer, 9-9
Configuring the Age Timer and Super Age Timer Using the Web Agent, 9-10
Configuring the IP Interface for the PPP Console, 12-91
configuring the LGMP server, 20-22, 20-25, 20-26, 20-27
configuring the PPP console static route, 12-90
configuring the serial port console
PPP console, 2-30
configuring the supervisor module
using the CLI, 2-2
configuring the temperature system, 3-13
configuring user port
gigabit module, 8-14, 8-18, 8-23
Configuring VLAN Parameters, 6-10
configuring VRRP, 12-72
congestion drops
buffer management table parameters, 24-6
connecting a modem, 2-29
overview, 2-29
console serial port
changing, 2-26
console serial port settings, changing
baud rate, 2-28
data bits, 2-28
flow control, 2-28
parity, 2-28
stop bits, 2-28
consumer signal
LDAP statistics, 12-90
controlling address table size reconfiguration, 9-13
conventions, 0-xxii
cooling system status
displaying, 3-16
copy destination
configuration file management, 2-43
copy source
configuration file management, 2-42
copying
configuration
invalid combinations, 2-44
from the switch or server to a configuration file, 2-42
running configuration to startup configuration, 2-40, 2-41, 2-44
startup configuration to running configuration, 2-42
to a TFTP server, 2-42
to the switch, 2-42
copying files, 2-43
copying from/to a TFTP server, 2-43
copying source files, 2-42
cost
IP static route parameters, 12-20
OSPF interface parameters, 15-11
CPU redundancy configuration
APP1, 10-17
APP2, 10-17
power-up/reset image, 10-16
startup config, 10-17
status, 10-16
CRC alignment errors
Ethernet interface statistical parameters, 23-5

create a Custom Access Type, 2-13
creating, 12-19
Creating 3Com Mapping Tables Using the Web Agent, 6-13
creating a BOOTP/DHCP server entry, 12-24, 12-25
creating a multinet interface, 12-15
creating a static client port, 20-18, 20-19
creating a VRRP router, 12-74
creating access rules for filtering traffic between subnets, 13-14
creating an AppleTalk interface, 19-4, 19-7
creating an AppleTalk static route, 19-10
creating an AppleTalk zone filter, 19-19
creating an IP interface, 12-4, 12-10
creating an IPX interface, 16-5, 16-9
creating an NBP filter
AppleTalk, 19-14
creating an VRRP virtual router, 12-74, 12-78
Creating Hunt Groups, 6-21
Creating Hunt Groups Using the CLI, 6-23
creating IP interface access list rules, 13-6
creating IPX RIP filters, 17-3, 17-8
creating IPX SAP filters, 18-3, 18-4
Creating IPX SAP Name Filters Using the Web Agent, 18-4
creating IPX SAP network filters, 18-7, 18-11
creating IPX static routes, 16-9, 16-11
Creating IPX Static Routes Using the CLI, 16-11
creating IPX static services, 16-12, 16-15
Creating IPX Static Services Using the CLI, 16-15
creating OSPF summaries, 15-16, 15-17
Creating OSPF Summaries Using the Web Agent, 15-16
creating OSPF virtual links, 15-12, 15-14
Creating Standard Access Rules, 13-6
creating static ARP entries
IP static route, 12-22
Creating Static Client Ports, 20-18
Creating Trusted RIP Neighbors Using the Web Agent, 14-6
Crossbar, 1-4, 1-6
crossbar features supported, 1-4, 1-6
current entries
forwarding cache (FE) parameters, 22-11
current number of routes
AppleTalk route table statistics, 19-29
DVMRP routing statistical parameters, 12-63
IP routing table statistical parameters, 12-49
current number of zones
AppleTalk zone table statistics, 19-32
current power available
power system statistical parameters, 3-15

D

data bits
console serial port settings, 2-28
day
summer time hours configuration, 3-9
DD number
OSPF neighbors parameters, 15-23
DDP forward counter
AppleTalk global statistics, 19-24
DDP input counter
AppleTalk global statistics, 19-24
DDP local counter
AppleTalk global statistics, 19-24
DDP output counter
AppleTalk global statistics, 19-24
DDP output long
AppleTalk global statistics, 19-24
DDP output short
AppleTalk global statistics, 19-24
dead interval
OSPF interface parameters, 15-11
OSPF virtual link parameters, 15-14
debug
protocol event log settings, 21-9
default metric
RIP interface parameters, 14-5
default password, 2-3
default route
RIP interface parameters, 14-5
delete a custom access type, 2-16
Delete Custom Access Type, 2-16
deleting an AppleTalk interface, 19-9
Deleting an AppleTalk Interface Using the CLI, 19-9
Deleting an AppleTalk Static Route Using the Web Agent, 19-13
deleting an OSPF area, 15-7, 15-8
deleting AppleTalk static routes, 19-13, 19-14
deleting OSPF virtual links, 15-15
Deleting Static Client Ports, 20-19
deleting static client ports, 20-19
Deleting Static Client Ports Using the CLI, 20-19
description
event and shutdown log entries, 21-12
designated bridge
spanning tree bridge port parameters, 7-16
designated cost
spanning tree bridge port parameters, 7-16
designated forwarder
DVMRP route table parameters, 12-65
designated forwarder(s) table parameters
forwarder cost to source network, 12-66
forwarder network address, 12-66

forwarding interface, 12-66
designated port
spanning tree bridge port parameters, 7-16
designated root
spanning tree bridge port parameters, 7-16
destination address
Layer 3 route cache search parameters, 22-8
destination filename
configuration file management, 2-43
destination group address
DVMRP downstream link parameters, 12-70
DVMRP multicast forwarding cache parameters, 12-68
DVMRP upstream source parameters, 12-71
IGMP local multicast forwarding cache parameters, 12-57
upstream prune information, 12-69
destination network
IPX datagram fields, 16-3
destination node
IPX datagram fields, 16-3
destination port
Layer 3 route cache search parameters, 22-8
destination port number
DVMRP upstream source parameters, 12-71
destination socket
IPX datagram fields, 16-3
destination subnet
IP access list parameters, 13-10
detail link
OSPF link state database parameter, 15-26
diagnostics
for switch hardware, 10-4
direction
IPX RIP filter parameters, 17-5
IPX SAP filter parameters, 18-5
IPX SAP network filter parameters, 18-9
displaying
intelligent multicasting, 20-8
IPX routing global statistics, 16-16, 16-18
displaying DVMRP global statistics, 12-57
displaying DVMRP multicast forwarding cache, 12-67
displaying DVMRP routing table statistics, 12-62
displaying frame forwarding cache statistics, 22-4
Displaying Global IP Routing Statistics, 12-40
displaying global IP routing statistics, 12-40, 12-49
displaying IGMP global statistics, 12-51
displaying IGMP local multicast forwarding cache, 12-56
displaying IPX RIP interface statistics, 10-1
displaying Layer 3 forwarding cache, 22-9
displaying multicast forwarding cache, 12-68
Displaying OSPF Links Using the Web Agent, 15-21
displaying OSPF neighbors, 15-22
Displaying OSPF Neighbors Using the Web Agent, 15-22
displaying OSPF statistics, 15-19
displaying router ports, 20-8
displaying the cooling system status, 3-16
displaying the DVMRP designated forwarder table, 12-66
displaying the DVMRP routing table, 12-63
Displaying the Forwarding Cache, 22-9
displaying the IGMP multicast group table, 12-55
displaying the power system statistics, 3-15
DNS client
configuring, 11-1
document conventions, 0-xxii
documentation
feedback, 0-xxiv
documentation feedback, 0-xxiv
domain name
VTP snooping parameters, 6-31
downloaded image
viewing, A-7
downloading image from TFTP server to an APP location, A-5
downstream dependent router
DVMRP route table parameters, 12-65
downstream interface
DVMRP multicast forwarding cache parameters, 12-69
downstream interfaces
IGMP local multicast forwarding cache parameters, 12-57
DR address
OSPF link parameters, 15-21
DR choice
OSPF neighbors parameters, 15-23
DR priority
OSPF interface parameters, 15-10
OSPF neighbors parameters, 15-23
DR router ID
OSPF link parameters, 15-21
DUAL-LAYER SPANNING TREE, 7-5
duplex mode
10/100 port parameter, 8-26
duplicate add attempts
forwarding cache (FE) parameters, 22-11
DVMRP
configuring, 12-35
displaying global statistics, 12-57
modifying information, 12-38
DVMRP designated forwarder table
displaying, 12-66

DVMRP downstream dependent router
DVMRP supported major/minor version, 12-67
found on interface, 12-67
router is SNMP manageable, 12-67
router network address, 12-67
router received probe from this router, 12-67
router supports generation ID function, 12-67
router supports prune function, 12-67
DVMRP downstream interface
DVMRP downstream link parameters, 12-70
DVMRP downstream link parameters
destination group address, 12-70
DVMRP downstream interface, 12-70
interface is pruned, 12-71
interface type, 12-71
prune expiration, 12-71
source subnetwork, 12-70
DVMRP downstream routers
viewing, 12-67
DVMRP global statistical parameters
graft acknowledgments received, 12-58
graft acknowledgments transmitted, 12-59
graft messages received, 12-58
graft messages transmitted, 12-59
probe message received, 12-58
probe messages transmitted, 12-59
prune messages received, 12-58
prune messages transmitted, 12-59
report messages received, 12-58
report messages transmitted, 12-59
unknown messages codes encountered, 12-58
DVMRP interface
modifying, 12-38
DVMRP interface parameters
interface, 12-38
interface metric, 12-39
interface scope, 12-39
interface type, 12-39
IP address, 12-38
IP address mask, 12-38
tunnel endpoint address, 12-39
DVMRP interface statistical parameters
address mask, 12-60
IGMP querier interface, 12-60
interface, 12-60
invalid routes received, 12-61
metric, 12-60
neighbor DVMRP routers, 12-61
network address, 12-60
next probe message in sec, 12-60
state, 12-60
type, 12-60
unrecognized packets received, 12-60
DVMRP interface statistics
examining, 12-59
DVMRP multicast forwarding cache
displaying, 12-67
DVMRP multicast forwarding cache parameters
destination group address, 12-68
downstream interface, 12-69
invalid flows from upstream, 12-69
next pruned downstream interface to timeout, 12-69
packets forwarded through cache entry, 12-69
source address mask, 12-69
source subnetwork, 12-68
upstream interface, 12-69
upstream interface is pruned, 12-69
upstream neighbor router address, 12-69
upstream sources, 12-69
DVMRP neighbor routers
DVMRP supported major/minor version, 12-62
expiration period, 12-62
found on interface, 12-62
neighbor is SNMP manageable, 12-62
neighbor network address, 12-62
neighbor received probe from this router, 12-62
neighbor supports generation ID function, 12-62
neighbor supports MTRACE requests, 12-62
neighbor supports prune function, 12-62
viewing, 12-61
DVMRP parameters
maximum number of routes, 12-37
DVMRP route table parameters
designated forwarder, 12-65
downstream dependent router, 12-65
expiration period in sec, 12-64
reporting router, 12-64
reporting router interface, 12-64
route metric, 12-64
source network, 12-64
source network mask, 12-64
upstream router, 12-64
DVMRP routing statistical parameters
current number of routes, 12-63
number of trigger routes, 12-63
number of valid routes, 12-63
DVMRP routing table
displaying, 12-63
DVMRP routing table statistics
displaying, 12-62
DVMRP supported major/minor version
DVMRP downstream dependent router, 12-67
DVMRP neighbor routers, 12-62
DVMRP upstream interface
upstream prune information, 12-69
DVMRP upstream routers
found on interface, 12-65

router cost to source network, 12-65
router network address, 12-65
viewing, 12-65
DVMRP upstream source parameters
destination group address, 12-71
destination port number, 12-71
flow source address, 12-71
flow upstream interface, 12-71
payload protocol type, 12-71
source port number, 12-71
DVMRP upstream sources, 12-71

E

echo reply rx
AppleTalk global statistics, 19-24
echo req rx
AppleTalk global statistics, 19-24
echo req tx
AppleTalk global statistics, 19-24
editing an AppleTalk static route, 19-12, 19-13
editing an AppleTalk zone filter, 19-20
editing an NBP filter, 19-16, 19-17
editing AppleTalk interfaces, 19-7, 19-9
editing general system information, 3-2
egress rules
VLAN, 6-6
enable state
CGMP snooping parameters, 20-34
enable state, 20-7
IGMP snooping, 20-21
LGMP client configuration parameters, 20-29
LGMP server configuration parameters, 20-24
RADIUS configuration parameters, 4-27
enabling
enabling IRDP, 12-83
enabling AppleTalk global routing, 19-4
enabling gigabit ports, 8-13, 8-17, 8-22
enabling SNTP, 3-3, 3-4
Enabling the LGMP Server, 20-22
enabling the rate limit mode
Fast Ethernet module, 8-25
enabling VRRP, 12-72
encapsulation type
AppleTalk interface parameters, 19-5
end
one-time summer time hours configuration, 3-11
summer time hours configuration, 3-8
end network
IPX RIP filter parameters, 17-5
entries removed due to route changes
forwarding cache (FE) parameters, 22-11
entries removed due to route deletes
forwarding cache (FE) parameters, 22-11
entry expiration period in sec
IGMP group membership table parameters, 12-56
entry type
address table instance parameters, 9-12
entry validity
address table instance parameters, 9-12
E-option
OSPF neighbors parameters, 15-23
error
protocol event log settings, 21-9
Ethernet interface statistical parameters
broadcast packets, 23-4
bytes, 23-4
CRC alignment errors, 23-5
fragments, 23-6
interval start, 23-4
jabbers, 23-6
multicast packets, 23-5
oversize packets, 23-6
packets, 23-4
sample, 23-4
undersize packets, 23-5
utilization, 23-4
Ethernet interface statistics
configuring, 23-3
Ethernet statistics, RMON
viewing, 23-1, 23-7
event and shutdown log entries
description, 21-12
event ID, 21-12
log ID, 21-12
severity, 21-12
time stamp, 21-12
type, 21-12
event and shutdown logs
viewing, 21-11, 21-12
event class parameters
AppleTalk, 21-6
authentication failure, 21-6
bridge status, 21-6
configuration, 21-6
fan status, 21-6
LDAP, 21-6
OSPF, 21-6
resource, 21-6
RIP, 21-6
service port status, 21-6
start, 21-6
switch fabric, 21-6
system, 21-6
temperature status, 21-6
user port status, 21-6

event configuration
accessing, 21-4, 21-10
event ID
event and shutdown log entries, 21-12
event log (switch event)
definition, 21-3
event log wraps
event statistics, 21-14
event notification
configuring, 21-4
event statistics
event log wraps, 21-14
events dropped due to event system queue full, 21-14
events dropped due to overload of event system, 21-14
viewing, 21-7, 21-10, 21-13
events dropped due to event system queue full
event statistics, 21-14
events dropped due to overload of event system
event statistics, 21-14
examining DVMRP interface statistics, 12-59
examining IP routing table statistics, 12-48
example of routing with Layer 2/3, 1-23
expiration period
DVMRP neighbor routers, 12-62
expiration period in sec
DVMRP route table parameters, 12-64
Extensive Fault Tolerance, 1-18
external LSA count
OSPF statistical parameters, 15-20

F

fabric port buffers
buffer management parameters, 24-4
failed add attempts
forwarding cache (FE) parameters, 22-11
failed element, 10-4
fan status
event class parameters, 21-6
Fast Ethernet module
auto-negotiation mode, enabling, 8-24
configuring, 8-21
rate limit mode, enabling, 8-25
Fast Ethernet module, configuring
pace priority mode, 8-25
rate limit rate, 8-25
Fast Ethernet Modules, 1-11
fault
protocol event log settings, 21-9
files
copying, 2-43
startup.txt, 2-38
filter/suppress
IPX RIP filter parameters, 17-5
IPX SAP filter parameters, 18-6
IPX SAP network filter parameters, 18-9
filtering
index, 9-17
MAC address, 9-17
persistence, 9-18
port, 9-17
priority, 9-17
filtering on the MAC address, 9-14, 9-15
filtering on the VLAN, 9-15
filtering traffic between subnets, 13-14
Filters
short-lived IP protocol, 12-15
FIRE fabric port
frame forwarding statistical parameters, 22-5
flooding., 6-1
flow control
console serial port settings, 2-28
flow control mode
10/100 port parameter, 8-26
gigabit port parameters, 8-15, 8-20
flow source address
DVMRP upstream source parameters, 12-71
flow upstream interface
DVMRP upstream source parameters, 12-71
forward delay
spanning tree bridge level parameters, 7-12
forward transitions
spanning tree bridge port parameters, 7-16
forwarder cost to source network
designated forwarder(s) table parameters, 12-66
forwarder network address
designated forwarder(s) table parameters, 12-66
forwarding cache
configuring, 22-2, 22-4
displaying, 22-9
forwarding cache (FE) parameters
aged entries, 22-11
current entries, 22-11
duplicate add attempts, 22-11
entries removed due to route changes, 22-11
entries removed due to route deletes, 22-11
failed add attempts, 22-11
lookup hits, 22-11
lookup levels, 22-11
lookup misses, 22-11
memory use in bytes, 22-11
mode, 22-11
PRE, 22-11
total entries, 22-11

type, 22-11
forwarding cache statistics
monitoring, 22-4
forwarding interface
designated forwarder(s) table parameters, 12-66
forwarding rules
VLAN, 6-5
forwarding state (spanning tree bridge port), 7-16
found on interface
DVMRP downstream dependent router, 12-67
DVMRP neighbor routers, 12-62
DVMRP upstream routers, 12-65
fragments
Ethernet interface statistical parameters, 23-6
Frame Classification, 1-17
frame format
VLAN switch port table parameters, 6-12
frame forwarding cache
displaying statistics, 22-4
frame forwarding statistical parameters
FIRE fabric port, 22-5
L2 frame count, 22-6
L3 drop frames, 22-6
L3 frame cache hits, 22-5
L3 slow path frames, 22-6
L3 total (T2), 22-5
percent cache hits, 22-5
percent drops, 22-6
percent slow path, 22-6
RX frame count (T2), 22-6
frame tags
switch port configuration parameters, 8-31
frame type
IPX interface, 16-8
front panel LEDs
interpreting, 21-1
front panel LEDs (10/100 module)
port, 21-2

G

general system information, setting
device contact, 3-2
device location, 3-2
switch name, 3-2
get nearest server reply
IPX SAP interface parameters, 18-3
get nearest server reply delay
IPX SAP interface parameters, 18-3
Gigabit Ethernet modules (L2/L3), 1-11
gigabit module
configure ports, 8-12, 8-16, 8-21, 8-27
gigabit module, configuring
flow control mode, 8-14, 8-18, 8-24
port name, 8-13, 8-17, 8-22
ports, 8-13, 8-17, 8-22
service port, 8-14, 8-18, 8-23
user port, 8-14, 8-18, 8-23
gigabit port name
configuring, 8-13, 8-17, 8-22
gigabit port parameters
category, 8-15, 8-19
flow control mode, 8-15, 8-20
port name, 8-15, 8-19
port PACE priority, 8-20
gigabit ports
configuring physical port parameters, 8-15, 8-19
enabling, 8-13, 8-17, 8-22
Gigabit-Speed Modules, 1-11
global configuration, 20-5, 20-9
intelligent multicast, 20-5, 20-9
GLOBAL DISABLE, 7-6
global RIP
RIP global configuration parameters, 14-2
Globally Enabling VRRP, 12-72
GNS requests received
IPX SAP interface statistical parameters, 18-13
GNS responses sent
IPX SAP interface statistical parameters, 18-13
graft acknowledgments received
DVMRP global statistical parameters, 12-58
graft acknowledgments transmitted
DVMRP global statistical parameters, 12-59
graft messages received
DVMRP global statistical parameters, 12-58
graft messages transmitted
DVMRP global statistical parameters, 12-59
group
address forwarding table parameters, 9-17
RADIUS configuration parameters, 4-27
group address
IGMP group membership table parameters, 12-56
group created on
IGMP group membership table parameters, 12-56
group join requests received
IGMP interface statistical parameters, 12-54
group leave requests received
IGMP interface statistical parameters, 12-54
group member interface
IGMP group membership table parameters, 12-56
group membership queries received
IGMP global statistic parameters, 12-52
group membership queries transmitted, IGMP global statistic parameters, 12-52
group membership reports received
IGMP global statistic parameters, 12-52

group multicast protocol
IGMP group membership table parameters, 12-56
group reporter address
IGMP group membership table parameters, 12-56
group reports received
IGMP interface statistical parameters, 12-54

H

hardware requirements for routing, 12-3, 16-4
hash mode
Layer 3 cache configuration, 22-3
hash table
address table instance parameters, 9-12
hello interval
OSPF interface parameters, 15-10
OSPF virtual link parameters, 15-14
hello time
spanning tree bridge level parameters, 7-12
high and normal overflow drops
buffer management table parameters, 24-6
high and normal stale drops
buffer management table parameters, 24-6
high priority allocation
buffer management table parameters, 24-6
high priority service ratio
buffer management table parameters, 24-6
high-preference static routes
IP global configuration parameters, 12-13
high-priority traffic, 24-1
hops
IPX RIP filter parameters, 17-6
IPX route table parameter, 16-20
IPX SAP filter parameters, 18-6
IPX SAP network filter parameters, 18-9
IPX service table parameter, 16-24
IPX static route, 16-10
IPX static service parameter, 16-14
hour
summer time hours configuration, 3-9
hunt group
considerations, 6-20
switch port configuration parameters, 8-32
switch port parameter, 8-41
hunt group members
viewing details, 6-27
hunt groups, 1-17
aggregate bandwidth, 6-13
before configuring, 6-13

I

ICMP in address mask reply
IP routing global statistics, 12-44
ICMP in address mask requests
IP routing global statistics, 12-44
ICMP in destination unreachable
IP routing global statistics, 12-44
ICMP in echo replies
IP routing global statistics, 12-44
ICMP in echo requests
IP routing global statistics, 12-44
ICMP in errors
IP routing global statistics, 12-44
ICMP in messages
IP routing global statistics, 12-43
ICMP in parameter problems
IP routing global statistics, 12-44
ICMP in redirects
IP routing global statistics, 12-44
ICMP in source quenches
IP routing global statistics, 12-44
ICMP in time exceeds
IP routing global statistics, 12-44
ICMP in timestamp replies
IP routing global statistics, 12-44
ICMP in timestamp requests
IP routing global statistics, 12-44
ICMP out address mask reply
IP routing global statistics, 12-45
ICMP out address mask requests
IP routing global statistics, 12-45
ICMP out destination unreachable
IP routing global statistics, 12-44
ICMP out echo reply
IP routing global statistics, 12-45
ICMP out echo requests
IP routing global statistics, 12-45
ICMP out errors
IP routing global statistics, 12-44
ICMP out messages
IP routing global statistics, 12-44
ICMP out parameter problems
IP routing global statistics, 12-44
ICMP out redirects
IP routing global statistics, 12-45
ICMP out source quenches
IP routing global statistics, 12-45
ICMP out time exceeded
IP routing global statistics, 12-44
ICMP out timestamp reply
IP routing global statistics, 12-45
ICMP out timestamp request
IP routing global statistics, 12-45
ICMP redirect
add IP interface parameters, 12-10
IEEE 802.3x full duplex flow control, 1-10

IEEE 802.3X PAUSE, 24-2
IEEE802.1D SPANNING TREE, 7-4
IGMP
configuring, 12-31
Global statistics, 12-51, 12-53, 12-55
overview, 12-31
IGMP global statistic parameters
group membership queries received, 12-52
group membership queries transmitted, 12-52
group membership reports received, 12-52
unknown messages received, 12-52
IGMP global statistics
displaying, 12-51
IGMP group membership table parameters
entry expiration period in sec, 12-56
group address, 12-56
group created on, 12-56
group member interface, 12-56
group multicast protocol, 12-56
group reporter address, 12-56
IGMP interface parameters
always be group membership querier, 12-34
IGMP version, 12-33
interface, 12-33
IP address, 12-33
IP address mask, 12-33
maximum groups, 12-33
neighbor querier timeout interval, 12-34
process leave packets, 12-34
query request interval, 12-34
query response interval, 12-34
robustness variable, 12-34
IGMP interface statistical parameters
group join requests received, 12-54
group leave requests received, 12-54
group reports received, 12-54
IGMP version, 12-53
interface, 12-53
IP address, 12-53
IP address mask, 12-53
neighbor querier timeout in sec, 12-54
next query request in sec, 12-54
number of current groups, 12-54
query messages received, 12-54
query messages transmitted, 12-54
robustness variable, 12-54
state, 12-53
this router is group membership querier, 12-54
unknown messages received, 12-54
IGMP interface statistics
viewing, 12-52, 12-53, 12-55
IGMP interfaces
configuring, 12-32
IGMP local multicast forwarding cache
displaying, 12-56
IGMP local multicast forwarding cache parameters
destination group address, 12-57
downstream interfaces, 12-57
invalid flows from upstream, 12-57
packets forwarded through cache entry, 12-57
source address mask, 12-57
source subnetwork, 12-57
upstream interface, 12-57
upstream source, 12-57
IGMP multicast group table
displaying, 12-55
IGMP Only Multicast Forwarding, 12-31
IGMP querier interface
DVMRP interface statistical parameters, 12-60
IGMP snooping
configuring, 20-21, 20-22
IGMP snooping parameters
enable state, 20-21
intelligent multicast session statistics, 20-21
new client ports added, 20-21
new router ports added, 20-21
new sessions created, 20-21
router ports removed, 20-21
sessions destroyed, 20-21
IGMP version
IGMP interface parameters, 12-33
IGMP interface statistical parameters, 12-53
in max hops exceeded
IPX global parameters, 16-18
index
address forwarding table parameters, 9-17
AppleTalk NBP table parameters, 19-33
AppleTalk zone table statistics, 19-31
info
protocol event log settings, 21-9
ingress rules
VLAN operation, 6-4
initial hash table size
VLAN configuration parameters, 6-10
intelligent multicast, 20-17, 20-18, 20-22, 20-25, 20-26, 20-27, 20-30, 20-31, 20-32, 20-33, 20-34, 20-35
session search, 20-10, 20-13
intelligent multicast global configuration parameters, 20-7
intelligent multicast session statistics
CGMP snooping parameters, 20-34
IGMP snooping, 20-21
LGMP client configuration parameters, 20-29
LGMP client display per VLAN parameters, 20-32
LGMP server configuration parameters, 20-25

LGMP server display per VLAN parameters, 20-27
intelligent multicasting, 20-5, 20-9, 20-10
interface
add IP interface parameters, 12-8
AppleTalk ARP cache table statistics, 19-30
AppleTalk interface parameters, 19-5
AppleTalk interface statistics, 19-26
AppleTalk NBP filter parameters, 19-16
AppleTalk NBP table parameters, 19-33
AppleTalk route table statistics, 19-28
AppleTalk zone filter parameters, 19-20
creating a VRRP virtual router, 12-74, 12-78
DVMRP interface parameters, 12-38
DVMRP interface statistical parameters, 12-60
IGMP interface parameters, 12-33
IGMP interface statistical parameters, 12-53
IP interface parameters, 12-9
IP routing table search parameters, 12-47
IPX interface, 16-8
IPX RIP filter parameters, 17-5
IPX route table parameter, 16-20
IPX routing table search parameters, 16-19
IPX SAP filter parameters, 18-5
IPX SAP interface parameters, 18-2
IPX SAP interface statistical parameters, 17-9, 18-12
IPX SAP network filter parameters, 18-8
IPX service table parameter, 16-24
IPX service table search parameter, 16-22
IPX static route, 16-10
IPX static service parameter, 16-14
IRDP parameters, 12-84
OSPF interface parameters, 15-10
RIP interface parameters, 14-4
RIP statistical parameters, 14-11
VRRP configuration parameters, 12-76
VRRP statistical parameters, 12-79
interface is pruned
DVMRP downstream link parameters, 12-71
upstream prune information, 12-70
interface metric
DVMRP interface parameters, 12-39
interface scope
DVMRP interface parameters, 12-39
interface type
DVMRP downstream link parameters, 12-71
DVMRP interface parameters, 12-39
upstream prune information, 12-70
interfaces
adding to NBP filter, 19-17, 19-18, 19-22
adding to zone filter, 19-21, 19-22
deleting from a zone filter, 19-21, 19-22
deleting from NBP filter, 19-17, 19-18, 19-22
interpreting front panel LEDs, 21-1
interpreting IPX SAP interface statistics, 17-8, 17-9, 18-12, 18-13
interval start
Ethernet interface statistical parameters, 23-4
intra-area routes
IP global configuration parameters, 12-13
invalid entry
permanent, definition, 9-18
invalid flows from upstream
DVMRP multicast forwarding cache parameters, 12-69
IGMP local multicast forwarding cache parameters, 12-57
invalid routes received
DVMRP interface statistical parameters, 12-61
IP
creating OSPF summaries, 15-16, 15-17
creating OSPF virtual links, 15-12, 15-14
creating RIP trusted neighbors, 14-6, 14-7
displaying OSPF links, 15-21
displaying OSPF neighbors, 15-22
displaying OSPF statistics, 15-19
DVMRP configuration, 12-35
interface configuration, 12-6
modifying OSPF areas, 15-9
multicast forwarding, 12-30
multicast routing, 12-30
multicast statistics, 12-51
multicast, overview, 12-30
OSPF area creation, 15-5
OSPF global configuration, 15-3
RIP configuring interfaces, 14-3
RIP global configuration, 14-1
RIP statistics, 14-11, 14-12
searching the OSPF link state database, 15-24, 15-28
Searching the routing table, 12-46
short-lived protocol filters, 12-15
IP access list parameters
access list name, 13-8, 13-10
access rule index, 13-8, 13-10
access type, 13-8, 13-10
destination subnet, 13-10
protocol ID, 13-10
TCP established, 13-11
TCP/UDP destination port, 13-11
TCP/UDP source port, 13-11
IP access rules
creating rules for filtering traffic between subnets, 13-14
filtering traffic between subnets example, 13-14

filtering Web traffic example, 13-13
IP address
ARP cache search parameters, 12-50
DVMRP interface parameters, 12-38
IGMP interface parameters, 12-33
IGMP interface statistical parameters, 12-53
IP routing table search parameters, 12-47
IP static ARP parameters, 12-23
OSPF interface parameters, 15-10
OSPF link parameters, 15-21
OSPF neighbors parameters, 15-23
RIP statistical parameters, 14-9, 14-11
static multicast session configuration parameters, 20-16
static multicast session parameters, 20-18
VRRP configuration parameters, 12-76
VRRP statistical parameters, 12-79
IP address mask
DVMRP interface parameters, 12-38
IGMP interface parameters, 12-33
IGMP interface statistical parameters, 12-53
IP and IPX router, 1-21
IP ARP
searching ARP cache, 12-49
IP forward datagrams
IP routing global statistics, 12-42
IP fragmentation creates
IP routing global statistics, 12-43
IP fragmentation failures, 12-43
IP fragmentation OKs
IP routing global statistics, 12-43
IP global configuration parameters
BOOTP/DHCP relay agent, 12-12
high-preference static routes, 12-13
intra-area routes, 12-13
IP multicast forwarding, 12-12
IP routing, 12-12
IP source routing, 12-12
limit proxy ARP to same network, 12-12
local routes, 12-13
low-preference static routes, 12-14
maximum number of ARP cache entries, 12-13
maximum number of routes, 12-13
OSPF external routes, 12-13
OSPF inter-area routes, 12-13
RIP routes, 12-14
route preference by protocol, 12-13
use default route for proxy ARP, 12-12
VRRP, 12-12
IP global routing
configuring, 12-11
IP global routing statistics
displaying, 12-40, 12-49
IP in address errors
IP routing global statistics, 12-42
IP in delivers
IP routing global statistics, 12-42
IP in discards
IP routing global statistics, 12-42
IP in header errors
IP routing global statistics, 12-41
IP in receives
IP routing global statistics, 12-41
IP in unknown protocols
IP routing global statistics, 12-42
IP interface
creating, 12-4, 12-10
IP interface access list rules
creating, 13-6
IP interface access lists
configuring, 13-1
IP interface parameters
interface, 12-9
IP interfaces
configuring, 12-6
IP multicast
overview, 12-30
IP multicast forward datagrams
IP routing global statistics, 12-45
IP multicast forwarding
IP global configuration parameters, 12-12
IP multicast in discard
IP routing global statistics, 12-45
IP multicast in receives
IP routing global statistics, 12-45
IP out discards
, 12-43
IP out no routes
IP routing global statistics, 12-43
IP out requests
IP routing global statistics, 12-42
IP reassembly failures
IP routing global statistics, 12-43
IP reassembly OKs
IP routing global statistics, 12-43
IP reassembly required
IP routing global statistics, 12-43
IP reassembly timeout
IP routing global statistics, 12-43
IP routing
add IP interface parameters, 12-9
creating a BOOTP/DHCP server entry, 12-24, 12-25
IP global configuration parameters, 12-12
statistics, 12-40
IP routing discards

IP routing global statistics, 12-43
IP routing global statistics
BOOTP/DHCP in discards, 12-46
BOOTP/DHCP in hops exceeded, 12-46
BOOTP/DHCP in requests, 12-45
BOOTP/DHCP in responses, 12-46
BOOTP/DHCP out requests, 12-46
BOOTP/DHCP out responses, 12-46
ICMP in address mask reply, 12-44
ICMP in address mask requests, 12-44
ICMP in destination unreachable, 12-44
ICMP in echo replies, 12-44
ICMP in echo requests, 12-44
ICMP in errors, 12-44
ICMP in messages, 12-43
ICMP in parameter problems, 12-44
ICMP in redirects, 12-44
ICMP in source quenches, 12-44
ICMP in time exceeds, 12-44
ICMP in timestamp replies, 12-44
ICMP in timestamp requests, 12-44
ICMP out address mask reply, 12-45
ICMP out address mask requests, 12-45
ICMP out destination unreachable, 12-44
ICMP out echo reply, 12-45
ICMP out echo requests, 12-45
ICMP out errors, 12-44
ICMP out messages, 12-44
ICMP out parameter problems, 12-44
ICMP out redirects, 12-45
ICMP out source quenches, 12-45
ICMP out time exceeded, 12-44
ICMP out timestamp reply, 12-45
ICMP out timestamp request, 12-45
IP forward datagrams, 12-42
IP fragmentation creates, 12-43
IP fragmentation failures, 12-43
IP fragmentation OKs, 12-43
IP in address errors, 12-42
IP in delivers, 12-42
IP in discards, 12-42
IP in header errors, 12-41
IP in receives, 12-41
IP in unknown protocols, 12-42
IP multicast forward datagrams, 12-45
IP multicast in discard, 12-45
IP multicast in receives, 12-45
IP out discards, 12-43
IP out no routes, 12-43
IP out requests, 12-42
IP reassembly failures, 12-43
IP reassembly OKs, 12-43
IP reassembly required, 12-43
IP reassembly timeout, 12-43
IP routing discards, 12-43
UDP in datagrams, 12-45
UDP in errors, 12-45
UDP in no ports, 12-45
UDP out datagrams, 12-45
IP routing search table
deleting local entries from routing table, 12-48
IP routing statistics, 12-40
IP routing table
searching, 12-46
IP routing table search parameters
interface, 12-47
IP address, 12-47
source, 12-47
IP routing table statistical parameters
current number of routes, 12-49
peak number of routes, 12-49
RIP queries, 12-49
RIP route changes, 12-49
total routes added, 12-49
total routes deleted, 12-49
IP routing table statistics
examining, 12-48
IP source routing
IP global configuration parameters, 12-12
IP static ARP parameters
IP address, 12-23
MAC address, 12-23
IP static route
creating static ARP entries, 12-22
IP static route parameters
cost, 12-20
mask, 12-20
network address, 12-20
next-hop address, 12-20
preference, 12-20
IP static routes
creating, 12-19
IPX, C-3
creating static routes, 16-9, 16-11
creating static services, 16-12, 16-15
datagram structure, 16-2
deleting static routes, 16-11
deleting static service, 16-15, 16-16
examining the route table, 16-20
examining the service table, 16-23
global statistics, 16-16, 16-18
interface configuration, 16-5, 16-9
modifying IPX RIP filters, 17-3
modifying static routes, 16-12
modifying static services, 16-15
network number, 16-1

node number, 16-1
RIP interfaces, 17-1
routing table statistics, 16-20
searching the route table, 16-18
searching the service table, 16-21
socket number, 16-1
IPX datagram fields
checksum, 16-2
destination network, 16-3
destination node, 16-3
destination socket, 16-3
packet length, 16-2
packet type, 16-3
source network, 16-3
source node, 16-3
source socket, 16-3
transport control, 16-2
IPX datagram structure, 16-2
IPX forward datagrams
IPX global parameters, 16-17
IPX global parameters, 16-5, 16-9
in max hops exceeded, 16-18
IPX forward datagrams, 16-17
IPX in checksum errors, 16-18
IPX in delivery, 16-17
IPX in discards, 16-17
IPX in header errors, 16-17
IPX in ping replies, 16-18
IPX in ping request, 16-18
IPX in receives, 16-17
IPX in unknown sockets, 16-18
IPX Netbios receives, 16-17
IPX out discards, 16-18
IPX out no routes, 16-18
IPX out packets, 16-18
IPX out ping replies, 16-18
IPX out ping requests, 16-18
IPX out requests, 16-18
maximum number of services, 16-5
IPX in checksum errors
IPX global parameters, 16-18
IPX in delivery
IPX global parameters, 16-17
IPX in discards
IPX global parameters, 16-17
IPX in header errors
IPX global parameters, 16-17
IPX in ping replies
IPX global parameters, 16-18
IPX in ping request
IPX global parameters, 16-18
IPX in receives
IPX global parameters, 16-17
IPX in unknown sockets
IPX global parameters, 16-18
IPX interface
admin state, 16-8
frame type, 16-8
interface, 16-8
network address, 16-8
node address, 16-8
RIP, 16-8
ticks, 16-8
type 20 packet propagation, 16-9
VLAN, 16-8
IPX interface MIBs, C-3
IPX Netbios receives
IPX global parameters, 16-17
IPX out discards
IPX global parameters, 16-18
IPX out no routes
IPX global parameters, 16-18
IPX out packets
IPX global parameters, 16-18
IPX out ping replies
IPX global parameters, 16-18
IPX out pint requests
IPX global parameters, 16-18
IPX out requests
IPX global parameters, 16-18
IPX overview, 16-1
IPX RIP filter
suppress advertising of all networks example, 17-6, 17-7
IPX RIP filter parameters
direction, 17-5
end network, 17-5
filter/suppress, 17-5
hops, 17-6
interface, 17-5
precedence, 17-5
start network, 17-5
ticks, 17-6
IPX RIP filters
creating, 17-3, 17-8
IPX RIP filters, modifying, 17-3
IPX RIP interface statistics
displaying, 10-1
IPX route table
examining, 16-20
IPX route table parameters
hops, 16-20
interface, 16-20
network, 16-20
next-hop MAC address, 16-20
source, 16-20

ticks, 16-20
TTL, 16-20
IPX route table search, 16-18
IPX route table statistics
examining, 16-20
IPX router
configuring, 16-4, 16-5
IPX routing
IPX routing, 16-5
IPX routing global statistics, displaying, 16-16, 16-18
IPX routing table search parameters
interface, 16-19
network number, 16-19
source, 16-19
IPX routing table statistics
current number of routes, 16-21
peak number of routes, 16-21
route add failures, 16-21
IPX SAP filter parameters
direction, 18-5
filter/suppress, 18-6
hops, 18-6
interface, 18-5
name, 18-5
precedence, 18-5
type, 18-5
IPX SAP filters
creating, 18-3, 18-4
filtering all services except NDS example, 18-10
suppress advertising of print servers example, 18-6
IPX SAP interface parameters
aging interval multiplier, 18-3
get nearest server reply, 18-3
get nearest server reply delay, 18-3
interface, 18-2
mode, 18-3
network number, 18-2
periodic update interval, 18-2
triggered updates, 18-3
use interpacket gap, 18-2
use max packet size, 18-2
IPX SAP interface statistical parameters
bad packets received, 17-9, 18-13
GNS requests received, 18-13
GNS responses sent, 18-13
interface, 17-9, 18-12
network number, 17-9, 18-12
non-triggered updates sent, 17-9, 18-13
requests received, 17-9, 18-13
state, 17-9, 18-12
updates received, 17-9, 18-13
IPX SAp interface statistical parameters
triggered updates sent, 17-9, 18-12
IPX SAP interface statistics
interpreting, 17-8, 17-9, 18-12, 18-13
IPX SAP interfaces
configuring, 18-1
IPX SAP network filter parameters
direction, 18-9
filter/suppress, 18-9
hops, 18-9
interface, 18-8
net, 18-8
precedence, 18-8
type, 18-9
IPX SAP network filters
creating, 18-7, 18-11
IPX service table
examining, 16-23
searching, 16-21
IPX service table parameter
hops, 16-24
interface, 16-24
network, 16-24
next-hop MAC address, 16-24
node, 16-24
socket, 16-24
source, 16-24
TTL, 16-24
type, 16-24
IPX service table parameters
name, 16-23
IPX service table search parameter
service name, 16-23
service type, 16-23
IPX service table search parameters
interface, 16-22
source, 16-22
IPX static route
hops, 16-10
interface, 16-10
network, 16-10
next-hop node, 16-10
ticks, 16-10
IPX static route parameters, 16-10
IPX static routes
creating, 16-9, 16-11
deleting, 16-11
modifying, 16-12
IPX static service parameter
hops, 16-14
interface, 16-14
network, 16-14
next-hop node, 16-14
node, 16-14

service name, 16-14
socket, 16-14
type, 16-14
IPX static service parameters, 16-13
IPX static services
creating, 16-12, 16-15
IPX syntax example, 16-1
IRDP, 12-83
enabling on an interface, 12-83
IRDP overview, 12-83
IRDP parameters
adv. address, 12-84
advertisement life time, 12-84
interface, 12-84
IRDP state, 12-84
max. advertisement interval, 12-84
minimum advertisement interval, 12-84
network address, 12-84
preferences, 12-84
IRDP state
IRDP parameters, 12-84

J

jabbers
Ethernet interface statistical parameters, 23-6

K

key
OSPF interface parameters, 15-11
known mode
switch port configuration parameters, 8-32, 8-33

L

L2 frame count
frame forwarding statistical parameters, 22-6
L3 drop frames
frame forwarding statistical parameters, 22-6
L3 frame cache hits
frame forwarding statistical parameters, 22-5
L3 slow path frames
frame forwarding statistical parameters, 22-6
L3 total (T2)
frame forwarding statistical parameters, 22-5
LAN router
configuration, 12-4
last change
LDAP statistics, 12-89
Layer 2
event log, 21-3
redundant switch matrix and controller module, 1-18
shutdown log, 21-3
Layer 2 Fast Ethernet modules
20-port 10/100BASE-TX, 1-11
Layer 2 Gigabit Ethernet modules, 1-11
Layer 2 routing
overview, 1-23
Layer 2 switching, 1-21
Layer 2/3 routing example, 1-23
Layer 3 (learned) traffic routing
overview, 1-22
Layer 3 (not learned) traffic routing
overview, 1-22
Layer 3 cache configuration
age interval, 22-3
aging, 22-3
hash mode, 22-3
maximum entries, 22-3
Layer 3 dynamic intelligent multicasting, 20-4
Layer 3 forwarding cache
displaying, 22-9
Layer 3 forwarding cache (FE) parameters
aged entries, 22-11
current entries, 22-11
duplicate add attempts, 22-11
entries removed due to route changes, 22-11
entries removed due to route deletes, 22-11
failed add attempts, 22-11
lookup hits, 22-11
lookup levels, 22-11
lookup misses, 22-11
memory use in bytes, 22-11
mode, 22-11
PRE, 22-11
total entries, 22-11
type, 22-11
Layer 3 overview, 20-4
Layer 3 route cache search parameters
comparison value, 22-8
destination address, 22-8
destination port, 22-8
PRE, 22-8
protocol, 22-8
rule number, 22-8
source address, 22-8
source port, 22-8
VLAN, 22-8
Layer 3 routing
overview, 1-22
LDAP
configuring, 12-86
configuring settings, 12-86, 12-88
configuring statistics, 12-89, 12-90
event class parameters, 21-6

settings, 12-86
LDAP configuration parameters
primary server IP address, 12-87
primary server port, 12-87
search base, 12-88
secondary server IP address, 12-87
secondary server port, 12-88
LDAP statistics
consumer signal, 12-90
last change, 12-89
producer signal, 12-89
learned entries
delete invalid learned entries, 9-13
learning state (spanning tree bridge port), 7-16
LEDs
port (10/100 module), 21-2
LGMP client configuration parameters
enable state, 20-29
intelligent multicast session statistics, 20-29
LGMP clients, 20-29
LGMP message reception statistics, 20-29
LGMP client display per VLAN parameters
intelligent multicast session statistics, 20-32
LGMP message reception statistics, 20-32
VLAN, 20-32
LGMP client, configuring
configuring/viewing a LGMP client, 20-30
LGMP clients
LGMP client configuration parameters, 20-29
LGMP message reception statistics
LGMP client configuration parameters, 20-29
LGMP client display per VLAN parameters, 20-32
LGMP server display per VLAN parameters, 20-27
LGMP message statistics
LGMP server configuration parameters, 20-24
LGMP message transmission statistics
LGMP server configuration parameters, 20-25
LGMP messages transmission statistics
LGMP server display per VLAN parameters, 20-27
LGMP server, 20-25
LGMP server configuration, 20-22, 20-25, 20-26, 20-27
LGMP server configuration parameters
enable state, 20-24
intelligent multicast session statistics, 20-25
LGMP message reception statistics, 20-24
LGMP message transmission statistics, 20-25
LGMP servers, 20-24
proxy mode, 20-24
robustness variable, 20-24
router report time, 20-24
server ID priority, 20-24
LGMP server display per VLAN parameters
intelligent multicast session statistics, 20-27
LGMP message reception statistics, 20-27
LGMP messages transmission statistics, 20-27
state, 20-27
VLAN, 20-27
LGMP servers
LGMP server configuration parameters, 20-24
limit proxy ARP to same network
IP global configuration parameters, 12-12
link state database, 15-24, 15-28
link state database search parameters
area ID, 15-25
router ID, 15-25
type, 15-25
links
switch port parameter, 8-41
load MIB
overview, C-1
Load Share Function, 6-17
local ext type
OSPF global configuration parameters, 15-4
local routes
IP global configuration parameters, 12-13
log ID
event and shutdown log entries, 21-12
logging into
Web Agent, 2-9
lookup hits
forwarding cache (FE) parameters, 22-11
lookup levels
forwarding cache (FE) parameters, 22-11
lookup misses
forwarding cache (FE) parameters, 22-11
low-preference static routes
IP global configuration parameters, 12-14
LS ID
LSA detail, 15-27
OSPF link state database parameter, 15-26
LSA checksum sum
OSPF statistical parameters, 15-20
LSA chksum sum
OSPF statistical parameters, 15-20
LSA detail
age, 15-27
area, 15-27
attached router ID 1 & 2, 15-28
checksum, 15-27
LS ID, 15-27
network mask, 15-28
router ID, 15-27
sequence, 15-27

type, 15-27
LSAs
OSPF statistical parameters, 15-20

M

MAC Address
search, 9-14, 9-15
MAC address
address forwarding table parameters, 9-17
AppleTalk ARP cache table statistics, 19-30
filtering, 9-14, 9-15
IP static ARP parameters, 12-23
multicast session parameters, 20-12
static multicast session configuration parameters, 20-17
static multicast session parameters, 20-18
switch port parameter, 8-41
MAC address value
adding, 9-19
MAC format
add IP interface parameters, 12-9
managing buffers and queues, 24-2, 24-7
mask
add IP interface parameters, 12-8
IP static route parameters, 12-20
OSPF summaries parameters, 15-17
master
OSPF neighbors parameters, 15-23
max age
spanning tree bridge level parameters, 7-12
max packets per second
port mirroring information parameters, 23-9
max. packets per second
mirror port configuration parameters, 23-11
max.advertisement interval
IRDP parameters, 12-84
maximum entries
Layer 3 cache configuration, 22-3
maximum groups
IGMP interface parameters, 12-33
maximum number of ARP cache entries
IP global configuration parameters, 12-13
maximum number of paths
OSPF global configuration parameters, 15-4
maximum number of routes
DVMRP parameters, 12-37
IP global configuration parameters, 12-13
maximum number of routes, 16-5
maximum number of services
IPX global parameters, 16-5
MC-option
OSPF neighbors parameters, 15-23
MD5 key ID
OSPF interface parameters, 15-11
OSPF virtual link parameters, 15-14
memory
buffer detailed configuration parameters, 24-5
memory use in bytes
forwarding cache (FE) parameters, 22-11
metric
AppleTalk interface parameters, 19-5
AppleTalk route table statistics, 19-27
DVMRP interface statistical parameters, 12-60
MIB groups supported, C-3
MIBs
AppleTalk, C-1
bridge MIBs, C-2
DVMRP, C-2
IGMP, C-2
IP forwarding/routing table, C-3
private IPX interface MIBs, C-3
Prominent, C-2
SNMPv2, C-4
minimum adv. interval
IRDP parameters, 12-84
minutes
summer time hours configuration, 3-9
mirror port
configuring
configuring a mirror port, 23-7, 23-12, 23-17
port mirroring information parameters, 23-9
RMON mirror port configuration parameters, 23-10
switch port configuration parameters, 8-33
mirror port configuration parameters
max. packets per second, 23-11
mode
forwarding cache (FE) parameters, 22-11
IPX SAP interface parameters, 18-3
RIP interface parameters, 14-4
spanning tree bridge level parameters, 7-11
modem
connecting, 2-29
modem configuration string parameters
&D0, 2-31
CD follows carrier, 2-31
E0, 2-31
S0=1, 2-31
software flow control, 2-31
modify IPX static routes, 16-12
modifying DVMRP interface information, 12-38
modifying IGMP interfaces, 12-32
Modifying OSPF Summaries, 15-18
modifying OSPF virtual links, 15-15

Modifying the DVMRP Global Configuration Using the CLI, 12-39
module features, 1-10
modules & ports, 6-30, 8-12, 8-16, 8-21, 8-28, 8-29, 8-39
monitoring switch performance
IP routing statistics, 12-40
OSPF, 15-18
Monitoring the Forwarding Cache Statistics, 22-4
month
summer time hours configuration, 3-9
multicast forwarding
description, 12-30
multicast forwarding cache
displaying, 12-68
multicast packets
Ethernet interface statistical parameters, 23-5
multicast protocol
add IP interface parameters, 12-10
multicast routing
description, 12-30
multicast session client port parameters
application, 20-15
port, 20-15
multicast session parameters
application, 20-13
client port, 20-13
MAC address, 20-12
session ID, 20-12
type, 20-12
VLAN, 20-12
multicast statistics
IP, 12-51
multi-gigabit connections example, 6-13
multilayer Fast Ethernet modules
20-port 10/100BASE-TX, 1-11
multi-link trunk groups
features, 6-14
multinetting, 12-14, 16-5

N

name
AppleTalk NBP filter parameters, 19-15
AppleTalk NBP table parameters, 19-33
AppleTalk zone filter parameters, 19-20
AppleTalk zone table statistics, 19-31
IPX SAP filter parameters, 18-5
IPX service table parameter, 16-23
spanning tree bridge port parameters, 7-15
switch port parameter, 8-41
NBP filter
adding interfaces, 19-17, 19-18, 19-22
deleting interfaces, 19-17, 19-18, 19-22
editing, 19-16, 19-17
neighbor DVMRP routers
DVMRP interface statistical parameters, 12-61
neighbor is SNMP manageable
DVMRP neighbor routers, 12-62
neighbor network address
DVMRP neighbor routers, 12-62
neighbor querier timeout in sec
IGMP interface statistical parameters, 12-54
neighbor querier timeout interval
IGMP interface parameters, 12-34
neighbor received probe from this router
DVMRP neighbor routers, 12-62
neighbor supports generation ID function
DVMRP neighbor routers, 12-62
neighbor supports MTRACE requests
DVMRP neighbor routers, 12-62
neighbor supports prune function
DVMRP neighbor routers, 12-62
net
IPX SAP network filter parameters, 18-8
NetBIOS UDP rebroadcast
add IP interface parameters, 12-10
network
IPX route table parameter, 16-20
IPX service table parameter, 16-24
IPX static route, 16-10
IPX static service parameter, 16-14
network address
add IP interface parameters, 12-8
DVMRP interface statistical parameters, 12-60
IP static route parameters, 12-20
IPX interface, 16-8
IRDP parameters, 12-84
OSPF summaries parameters, 15-17
RIP interface parameters, 14-4
network mask
LSA detail, 15-28
network number
AppleTalk interface parameters, 19-6
AppleTalk interface statistics, 19-26
AppleTalk static route, 19-11
IPX routing table search parameters, 16-19
IPX SAP interface parameters, 18-2
IPX SAP interface statistical parameters, 17-9, 18-12
network range
AppleTalk ARP cache table statistics, 19-30
AppleTalk interface statistics, 19-26
AppleTalk route table statistics, 19-27
AppleTalk zone table statistics, 19-31
network range end
AppleTalk interface parameters, 19-6

AppleTalk static route, 19-11
network range start
AppleTalk interface parameters, 19-6
AppleTalk static route, 19-11
new client ports added
IGMP snooping, 20-21
new router ports added
IGMP snooping, 20-21
new sessions created
IGMP snooping, 20-21
next hop
AppleTalk route table statistics, 19-28
next probe message in sec
DVMRP interface statistical parameters, 12-60
next pruned downstream interface to timeout
DVMRP multicast forwarding cache parameters, 12-69
next query request in sec
IGMP interface statistical parameters, 12-54
next-hop address
IP static route parameters, 12-20
next-hop MAC address
IPX route table parameter, 16-20
IPX service table parameter, 16-24
next-hop node
IPX static route, 16-10
IPX static service parameter, 16-14
no client
AppleTalk global statistics, 19-24
no route
AppleTalk global statistics, 19-24
node
AppleTalk ARP cache table statistics, 19-30
AppleTalk interface parameters, 19-6
AppleTalk interface statistics, 19-26
AppleTalk static route, 19-11
IPX service table parameter, 16-24
IPX static service parameter, 16-14
node address
IPX interface, 16-8
non-tagged access mode, definition, 1-10
non-triggered updates sent
IPX SAP interface statistical parameters, 17-9, 18-13
RIP statistical parameters, 14-9, 14-11
N-option
OSPF neighbors parameters, 15-23
normal-priority traffic, 24-1
Notes, Cautions, and Warnings, 0-xxiii
number of current groups
IGMP interface statistical parameters, 12-54
number of routes
IPX routing table statistic, 16-21
number of triggered routes
DVMRP routing statistical parameters, 12-63
number of valid routes
DVMRP routing statistical parameters, 12-63

O

one-time summer time hours configuration, 3-10
end, 3-11
start, 3-11
Online documentation
adding files to a Web server, 0-xxi
Online Help, 0-xxi
Online help
accessing help files from the server, 0-xxi
adding files to a Web server, 0-xxi
OPQ option
OSPF neighbors parameters, 15-23
originate LSA count
OSPF statistical parameters, 15-20
OSPF
add IP interface parameters, 12-9
creating areas, 15-5
creating summaries, 15-16, 15-17
creating virtual links, 15-12, 15-14
displaying links, 15-21
displaying neighbors, 15-22
event class parameters, 21-6
monitoring switch performance, 15-18
OSPF global configuration parameters, 15-4
searching the link state database, 15-24, 15-28
OSPF area parameters
area ID, 15-6
area type, 15-6
stub metric, 15-6
translate 7 into 5, 15-6
type 3 ASE filter, 15-6
OSPF areas
deleting, 15-7, 15-8
OSPF external routes
IP global configuration parameters, 12-13
OSPF global configuration
configuring, 15-3
OSPF global configuration parameters
AS border route, 15-4
auto-creation of virtual links, 15-4
local ext type, 15-4
maximum number of paths, 15-4
OSPF, 15-4
RIP ext type, 15-4
router ID, 15-4
SPF hold time, 15-4
SPF suspend, 15-4

static ext type, 15-5
static low ext type, 15-5
OSPF inter-area routes
IP global configuration parameters, 12-13
OSPF interface parameters
area, 15-10
authentication, 15-11
cost, 15-11
dead interval, 15-11
DR priority, 15-10
hello interval, 15-10
interface, 15-10
IP address, 15-10
key, 15-11
MD5 key ID, 15-11
poll interval, 15-11
retransmit interval, 15-10
transit delay, 15-10
OSPF interfaces
modifying, 15-9
OSPF link parameters
area ID, 15-21
BDR address, 15-21
DR address, 15-21
DR router ID, 15-21
IP address, 15-21
state, 15-21
type, 15-21
OSPF link state database parameter
age, 15-26
area ID, 15-26
checksum, 15-26
detail link, 15-26
LS ID, 15-26
router ID, 15-26
sequence, 15-26
type, 15-26
OSPF link state database search parameters
area ID, 15-25
router ID, 15-25
type, 15-25
OSPF neighbors parameters
BDR choice, 15-23
DD number, 15-23
DR choice, 15-23
DR priority, 15-23
E-option, 15-23
IP address, 15-23
master, 15-23
MC-option, 15-23
N-option, 15-23
OPQ-option, 15-23
router ID, 15-23
state, 15-23
T-option, 15-23
OSPF packet tracing
enabling, 21-8, 21-10
OSPF state
OSPF statistical parameters, 15-19
OSPF statistical parameters
area ID, 15-20
AS border Rtrs, 15-20
border Rtrs, 15-20
external LSA count, 15-20
LSA checksum sum, 15-20
LSA chksum sum, 15-20
LSAs, 15-20
originate LSA count, 15-20
OSPF state, 15-19
OSPF version, 15-19
receive new LSA count, 15-20
router ID, 15-20
SPF runs, 15-20
OSPF statistics
displaying, 15-19
OSPF summaries parameters
advertisement, 15-17
area, 15-17
mask, 15-17
network address, 15-17
OSPF version
OSPF statistical parameters, 15-19
OSPF virtual link parameters
area, 15-14
authentication, 15-14
authentication key, 15-14
dead interval, 15-14
hello interval, 15-14
MD5 key ID, 15-14
retransmit interval, 15-14
router ID, 15-14
transit delay, 15-14
OSPF virtual links
deleting, 15-15
modifying, 15-15
overflow drops
buffer management table parameters, 24-6
oversize packets
Ethernet interface statistical parameters, 23-6
overview, 1-17
ATM, 1-10
intelligent multicasting, 20-4
overview of Layer 2 traffic routing, 1-23
overview of Layer 3 (learned) traffic routing, 1-22
overview of Layer 3 (not learned) traffic routing, 1-22
overview of Layer 3 routing, 1-22

owner
AppleTalk route table statistics, 19-28

P

packet length
IPX datagram fields, 16-2
packet tracing
enabling, 21-8, 21-10
packet type
IPX datagram fields, 16-3
packets
Ethernet interface statistical parameters, 23-4
packets forwarded through cache entry
DVMRP multicast forwarding cache parameters, 12-69
IGMP local multicast forwarding cache parameters, 12-57
parity
console serial port settings, 2-28
path cost
spanning tree bridge port information parameters, 7-18
payload protocol type
DVMRP upstream source parameters, 12-71
peak number of routes
AppleTalk route table statistics, 19-29
IP routing table statistical parameters, 12-49
IPX routing table statistic, 16-21
peak number of zones
AppleTalk zone table statistics, 19-32
PER VLAN SPANNING TREE, 7-5
percent cache hits
frame forwarding statistical parameters, 22-5
percent drops
frame forwarding statistical parameters, 22-6
percent slow path
frame forwarding statistical parameters, 22-6
performing a system reset, 3-17
periodic update interval
IPX SAP interface parameters, 18-2
permanent
invalid entry, 9-18
persistence entry, 9-18
persistence
address forwarding table parameters, 9-18
persistence entry
ageout, definition, 9-18
permanent, definition, 9-18
physical port buffers
buffer management parameters, 24-4
piggyback port
port mirroring information parameters, 23-9
port mirroring parameter, 23-11
poll interval
OSPF interface parameters, 15-11
port
address forwarding table parameters, 9-17
multicast session client port parameters, 20-15
RMON mirror port configuration parameters, 23-10
router port display parameters, 20-9
spanning tree bridge port parameters, 7-15
switch port parameter, 8-41
VLAN switch port table parameters, 6-11
port density, 1-3, 1-5
port LED (10/100 module)
behavior, 21-2
port mirror, 23-8
port mirroring information parameters, 23-9
configure source, 23-9
max packets per second, 23-9
mirror port, 23-9
piggyback port, 23-9
sampler type, 23-9
source port, 23-9
port mirroring parameters, 23-11
port name
gigabit port parameters, 8-15, 8-19
VLAN switch port table parameters, 6-11
port numbers
pointer to a well-known list, 13-11
port PACE priority
10/100 port parameter, 8-27
gigabit port parameters, 8-20
port VLAN
Cisco Catalyst 5000 parameters, 8-34
switch port configuration parameters, 8-31
switch port parameter, 8-41
port-based VLANs, 6-2
power supply
power system statistical parameters, 3-15
power system statistical parameters
current power available, 3-15
power supply, 3-15
status, 3-15
total system power, 3-15
type, 3-15
power system statistics
displaying, 3-15
power-up/reset image
CPU redundancy configuration, 10-16
PPP console
configuring, 2-30
configuring a network mask, 12-91
configuring an IP address, 12-91

PPP console settings
baud rate, 2-28
flow control, 2-28
PPP console static route
configuring, 12-90
PPP serial port console, 2-34
regaining access to the CLI, 2-32
PRE
forwarding cache (FE) parameters, 22-11
Layer 3 route cache search parameters, 22-8
precedence
IPX RIP filter parameters, 17-5
IPX SAP filter parameters, 18-5
IPX SAP network filter parameters, 18-8
preference
IP static route parameters, 12-20
preferences
IRDP parameters, 12-84
primary controller
replacing, 10-8
primary server
RADIUS configuration parameters, 4-27
primary server IP address
LDAP configuration parameters, 12-87
primary server port
LDAP configuration parameters, 12-87
priority
address forwarding table parameters, 9-17
spanning tree bridge level parameters, 7-12
spanning tree bridge port information parameters, 7-18
VRRP configuration parameters, 12-76
priority queuing, definition, 1-10
priority threshold
buffer management table parameters, 24-6
probe message received
DVMRP global statistical parameters, 12-58
probe messages transmitted
DVMRP global statistical parameters, 12-59
process leave packets
IGMP interface parameters, 12-34
producer signal
LDAP statistics, 12-89
Prominent MIBs, C-2
protocol
Layer 3 route cache search parameters, 22-8
protocol event log
configuring, 21-8, 21-10
protocol event log settings
debug, 21-9
error, 21-9
fault, 21-9
info, 21-9
trace, 21-9
warning, 21-9
protocol ID
IP access list parameters, 13-10
protocol numbers
pointer to a complete list, 13-10
proxy ARP
add IP interface parameters, 12-10
IP global configuration parameters, 12-12
proxy mode
LGMP server configuration parameters, 20-24
prune expiration
DVMRP downstream link parameters, 12-71
prune expiration time
upstream prune information, 12-70
prune messages received
DVMRP global statistical parameters, 12-58
prune messages transmitted
DVMRP global statistical parameters, 12-59
purge TTL
RIP global configuration parameters, 14-2

Q

query messages received
IGMP interface statistical parameters, 12-54
query messages transmitted
IGMP interface statistical parameters, 12-54
query request interval
IGMP interface parameters, 12-34
query response interval
IGMP interface parameters, 12-34
queues
managing, 24-2, 24-7
service ratio, 24-2

R

RADIUS client
configuring, 4-26
RADIUS configuration parameters
Cajun-Service-Type required, 4-28
enable state, 4-27
group, 4-27
primary server, 4-27
realm, 4-27
retry number, 4-28
retry time, 4-28
secondary server, 4-27
shared secret, 4-27
source IP address, 4-27
UDP port, 4-28
RADIUS secondary server
shared secret, 4-27

rate limit burst size
10/100 port parameter, 8-27
rate limit mode
10/100 port parameter, 8-27
rate limit rate
10/100 port parameter, 8-27
realm
RADIUS configuration parameters, 4-27
receive new LSA count
OSPF statistical parameters, 15-20
receive version
RIP interface parameters, 14-4
recurring
summer time hours configuration, 3-8
redundant CPU
configuring, 10-10
redundant element, 10-4
configuring, 10-9
Related Documentation, 0-xx
Related documents, 0-xx
Remote Authentication Dial-In User Service (RADIUS) Client Support, 11-1
replacing the primary controller, 10-8
report messages received
DVMRP global statistical parameters, 12-58
report messages transmitted
DVMRP global statistical parameters, 12-59
reporting router
DVMRP route table parameters, 12-64
reporting router interface
DVMRP route table parameters, 12-64
requests received
IPX SAP interface statistical parameters, 17-9, 18-13
resetting the switch, A-8
Resetting the Switch Using the Web Agent, A-8
resource
event class parameters, 21-6
retransmit interval
OSPF interface parameters, 15-10
OSPF virtual link parameters, 15-14
retry number
RADIUS configuration parameters, 4-28
retry time
RADIUS configuration parameters, 4-28
RIP
add IP interface parameters, 12-9
creating trusted RIP neighbors, 14-6, 14-7
event class parameters, 21-6
IPX interface, 16-8
RIP, 16-5
trusted neighbors, 14-6, 14-7
RIP ext type
OSPF global configuration parameters, 15-4
RIP global configuration, 14-1
RIP global configuration parameters
global RIP, 14-2
purge TTL, 14-2
triggered updates, 14-3
update pkt delay, 14-3
update timer, 14-2
RIP interface parameters
auth key, 14-5
auth type, 14-5
default metric, 14-5
default route, 14-5
interface, 14-4
mode, 14-4
network address, 14-4
receive version, 14-4
split horizon, 14-4
version, 14-4
RIP interfaces
configuring, 14-3
RIP packet tracing
enabling, 21-8, 21-10
RIP queries
IP routing table statistical parameters, 12-49
RIP route changes
IP routing table statistical parameters, 12-49
RIP routes
IP global configuration parameters, 12-14
RIP statistical parameters
bad packets received, 14-10, 14-12
bad routes received, 14-10, 14-12
interface, 14-11
IP address, 14-9, 14-11
non-trigger updates sent, 14-9, 14-11
state, 14-11
triggered updates sent, 14-11
updates received, 14-10, 14-12
RIP statistics
interpreting, 14-11, 14-12
RMON
configuring a mirror port, 23-7, 23-12, 23-17
RMON mirror port configuration parameters
mirror port, 23-10
piggyback port, 23-11
port, 23-10
sampler type, 23-11
RMON statistics
viewing, 23-1, 23-7
robustness variable
IGMP interface parameters, 12-34
IGMP interface statistical parameters, 12-54
LGMP server configuration parameters, 20-24

route add failures
IPX routing table statistic, 16-21
route metric
DVMRP route table parameters, 12-64
route preference by protocol
IP global configuration parameters, 12-13
router cost to source network
DVMRP upstream routers, 12-65
router ID
LSA detail, 15-27
OSPF global configuration parameters, 15-4
OSPF link state database parameter, 15-26
OSPF link state database search parameters, 15-25
OSPF neighbors parameters, 15-23
OSPF statistical parameters, 15-20
OSPF virtual link parameters, 15-14
router is SNMP manageable
DVMRP downstream dependent router, 12-67
router network address
DVMRP downstream dependent router, 12-67
DVMRP upstream routers, 12-65
router port display parameters
application, 20-9
port, 20-9
VLAN, 20-9
router ports, 20-8
router ports removed
IGMP snooping, 20-21
router pruning timer
router pruning timer, 20-7
router received probe from this router
DVMRP downstream dependent router, 12-67
router report time
LGMP server configuration parameters, 20-24
router supports generation ID function
DVMRP downstream dependent router, 12-67
router supports prune function
DVMRP downstream dependent router, 12-67
routing
hardware requirements, 12-3, 16-4
Layer 2 overview, 1-23
Layer 3 (learned) overview, 1-22
Layer 3 (not learned) overview, 1-22
Layer 3 overview, 1-22
routing configuration
minimum requirements, 12-3
Routing with Layer 2 and Layer 3 Modules, 1-21
RTMP RDR Rx
AppleTalk global statistics, 19-25
RTMP RDR Tx
AppleTalk global statistics, 19-25
RTMP Rq Rx
AppleTalk global statistics, 19-25
RTMP Rq Tx
AppleTalk global statistics, 19-25
RTMP Rsp Rx
AppleTalk global statistics, 19-25
RTMP Rsp Tx
AppleTalk global statistics, 19-25
rule number
Layer 3 route cache search parameters, 22-8
running configuration
viewing, 2-39
running.txt, 2-38
RX frame count (T2)
frame forwarding statistical parameters, 22-6

S

sample
Ethernet interface statistical parameters, 23-4
sampler type
port mirroring information parameters, 23-9
RMON mirror port configuration parameters, 23-11
SAP
SAP, 16-5, 16-9
save running configuration to startup configuration
configuration file management, 2-42
saving your current configuration, A-3
search base
LDAP configuration parameters, 12-88
Searching the OSPF Link State Database Using the CLI, 15-28
Searching the OSPF Link State Database Using the Web Agent, 15-24
searching the routing cache, 22-6
Searching the Routing Cache for an Entry, 22-7, 22-9, 22-10
Searching the Switch AFT Using the CLI, 9-19
secondary server
RADIUS configuration parameters, 4-27
secondary server IP address
LDAP configuration parameters, 12-87
secondary server port
LDAP configuration parameters, 12-88
seed
AppleTalk interface statistics, 19-26
Select, 12-84
Selecting the Image for Reboot Using the Web Agent, A-7
sequence
LSA detail, 15-27
OSPF link state database parameter, 15-26
server ID priority

LGMP server configuration parameters, 20-24
service name
IPX service table search parameter, 16-23
IPX static service parameter, 16-14
service port status
event class parameters, 21-6
service ratio (queues), 1-20, 24-2
service type
IPX service table search parameter, 16-23
session ID
multicast session parameters, 20-12
session pruning timer
session pruning timer, 20-7
session search
intelligent multicast, 20-10, 20-13
sessions destroyed
IGMP snooping, 20-21
set recurring
summer time hours, 3-5, 3-9
setting all ports in a module, 8-28
setting summer time hours, 3-5
one time, 3-10
setting switch ports on all modules, 8-39
setting the system clock, 3-3
setting up user accounts, 2-19
severity
event and shutdown log entries, 21-12
shared secret
RADIUS configuration parameters, 4-27
RADIUS secondary server, 4-27
Short PDU in error
AppleTalk global statistics, 19-24
Short-lived IP protocol filters, 12-15
shutdown log (switch event)
definition, 21-3
size
address table instance hash table parameter, 9-12
SNMP Administration, 2-27, 2-30
SNTP
enabling, 3-3, 3-4
socket
IPX service table parameter, 16-24
IPX static service parameter, 16-14
source
IP routing table search parameters, 12-47
IPX route table parameter, 16-20
IPX routing table search parameters, 16-19
IPX service table parameter, 16-24
IPX service table search parameter, 16-22
source address
Layer 3 route cache search parameters, 22-8
source address mask
DVMRP multicast forwarding cache parameters, 12-69
IGMP local multicast forwarding cache parameters, 12-57
source filename
configuration file management, 2-42
source IP address
RADIUS configuration parameters, 4-27
source network
DVMRP route table parameters, 12-64
IPX datagram fields, 16-3
source network mask
DVMRP route table parameters, 12-64
source node
IPX datagram fields, 16-3
source port
Layer 3 route cache search parameters, 22-8
port mirroring information parameters, 23-9
source port number
DVMRP upstream source parameters, 12-71
source socket
IPX datagram fields, 16-3
source subnetwork
DVMRP downstream link parameters, 12-70
DVMRP multicast forwarding cache parameters, 12-68
IGMP local multicast forwarding cache parameters, 12-57
upstream prune information, 12-69
spanning tree
path cost, 7-18
spanning tree bridge level parameters
bridge forward delay, 7-12
bridge hello time, 7-12
bridge max age, 7-12
forward delay, 7-12
hello time, 7-12
max age, 7-12
mode, 7-11
priority, 7-12
spanning tree bridge port
blocking state, 7-16
forwarding state, 7-16
learning state, 7-16
spanning tree bridge port information parameters
path cost, 7-18
priority, 7-18
spanning tree bridge port parameters
bridge port, 7-15
designated bridge, 7-16
designated cost, 7-16
designated port, 7-16
designated root, 7-16
forward transitions, 7-16

name, 7-15
port, 7-15
state, 7-16
spanning tree mode
switch port configuration parameters, 8-32
speed mode
10/100 port parameter, 8-26
SPF hold time
OSPF global configuration parameters, 15-4
SPF runs
OSPF statistical parameters, 15-20
SPF suspend
OSPF global configuration parameters, 15-4
split horizon
RIP interface parameters, 14-4
stale drops
buffer management table parameters, 24-6
STAP mode
switch port parameter, 8-41
start
event class parameters, 21-6
one-time summer time hours configuration, 3-11
summer time hours configuration, 3-8
start network
IPX RIP filter parameters, 17-5
startup config
CPU redundancy configuration, 10-17
startup configuration
viewing, 2-39
startup.txt, 2-38
state
AppleTalk interface statistics, 19-26
AppleTalk route table statistics, 19-27
DVMRP interface statistical parameters, 12-60
IGMP interface statistical parameters, 12-53
IPX SAP interface statistical parameters, 17-9, 18-12
LGMP server display per VLAN parameters, 20-27
OSPF link parameters, 15-21
OSPF neighbors parameters, 15-23
RIP statistical parameters, 14-11
spanning tree bridge port parameters, 7-16
VRRP statistical parameters, 12-80
static
Virtual LAN binding options, 8-36
static address configuration
adding entries to the address forwarding table, 9-19
static client port
creating, 20-18, 20-19
static client ports
deleting, 20-19
static multicast session parameters, 20-18
static ext type
OSPF global configuration parameters, 15-5
static low ext type
OSPF global configuration parameters, 15-5
static multicast session configuration parameters
IP address, 20-16
MAC address, 20-17
VLAN, 20-17
static multicast session parameters
IP address, 20-18
MAC address, 20-18
static client ports, 20-18
VLAN, 20-18
static multicast session, creating
creating a static multicast session, 20-15, 20-17
static multicast sessions, configuring
configuring static multicast sessions, 20-17, 20-18
static route
PPP console, 12-90
static route, AppleTalk, 19-10
static routes
IPX
deleting, 16-11
status
address forwarding table parameters, 9-18
CPU redundancy configuration, 10-16
power system statistical parameters, 3-15
stop bits
console serial port settings, 2-28
stub metric
OSPF area parameters, 15-6
summer time hours
one-time setting, 3-10
set recurring, 3-5, 3-9
setting, 3-5
summer time hours configuration
day, 3-9
end, 3-8
hour, 3-9
minutes, 3-9
month, 3-9
recurring, 3-8
start, 3-8
week, 3-8
super age time
definition, 9-9
supervisor module configuration
CLI, 2-2
supported browsers, 2-9
switch
resetting, A-8
switch configuration

managing, 2-40
switch description, 1-1
switch fabric
event class parameters, 21-6
switch features
crossbar switch fabric, overview, 1-4, 1-6
switch port
configuring, 8-29
switch port configuration parameters
3Com mapping table, 8-33
allow learning, 8-32
automatic VLAN creation, 8-31
frame tags, 8-31
hunt group, 8-32
known mode, 8-32, 8-33
mirror port, 8-33
port VLAN, 8-31
spanning tree mode, 8-32
trunk mode, 8-31
VLAN binding, 8-31
VTP snooping, 8-31
switch port parameters
hunt group, 8-41
links, 8-41
MAC address, 8-41
name, 8-41
port, 8-41
port VLAN, 8-41
STAP mode, 8-41
trunk mode, 8-41
viewing, 8-41
VLAN classification, 8-41
switch ports, single command, 8-39
Synchronizing the Primary and Redundant CPUs, 10-12
system
event class parameters, 21-6
system information
editing, 3-2
system information, changing
contact, 3-3
location, 3-2
system reset
performing, 3-17

T

TblInst
address forwarding table parameters, 9-17
TCP established
IP access list parameters, 13-11
TCP/UDP destination port
IP access list parameters, 13-11
TCP/UDP source port
IP access list parameters, 13-11
technical support, 0-xx
temperature status
event class parameters, 21-6
temperature system
configuring, 3-13
terminal setup, 2-2
TFTP server
copying, 2-43
TFTP server IP address
configuration file management, 2-43
TFTP transfer status
viewing, 2-45
this router is group membership querier
IGMP interface statistical parameters, 12-54
Throughput, 1-4, 1-6
ticks
IPX RIP filter parameters, 17-6
IPX route table parameter, 16-20
IPX static route, 16-10
time of state change
VRRP statistical parameters, 12-80
time stamp
event and shutdown log entries, 21-12
times this VR became master
VRRP statistical parameters, 12-80
too long
AppleTalk global statistics, 19-24
too short
AppleTalk global statistics, 19-24
T-option
OSPF neighbors parameters, 15-23
total entries
forwarding cache (FE) parameters, 22-11
total number of entries
address table instance parameters, 9-11
total routes added
IP routing table statistical parameters, 12-49
total routes deleted
IP routing table statistical parameters, 12-49
total system power
power system statistical parameters, 3-15
trace
protocol event log settings, 21-9
traffic flooding, 23-8
traffic mirror port, 23-10
transit delay
OSPF interface parameters, 15-10
OSPF virtual link parameters, 15-14
translate 7 into 5
OSPF area parameters, 15-6
transport control

IPX datagram fields, 16-2
triggered updates
IPX SAP interface parameters, 18-3
RIP global configuration parameters, 14-3
triggered updates sent
IPX SAP interface statistical parameters, 17-9, 18-12
RIP statistical parameters, 14-11
trunk mode
Cisco Catalyst 5000, parameters, 8-34
switch port configuration parameters, 8-31
switch port parameter, 8-41
TTL
AppleTalk ARP cache table statistics, 19-30
IPX route table parameter, 16-20
IPX service table parameter, 16-24
purge, 14-2
TTL expired
AppleTalk global statistics, 19-24
TTY console settings
baud rate, 2-28
data bits, 2-28
flow control, 2-28
parity, 2-28
stop bits, 2-28
tunnel endpoint address
DVMRP interface parameters, 12-39
type
AppleTalk ARP cache table statistics, 19-30
AppleTalk NBP filter parameters, 19-16
AppleTalk NBP table parameters, 19-33
AppleTalk static route, 19-11
AppleTalk zone filter parameters, 19-20
DVMRP interface statistical parameters, 12-60
event and shutdown log entries, 21-12
forwarding cache (FE) parameters, 22-11
IPX SAP filter parameters, 18-5
IPX SAP network filter parameters, 18-9
IPX service table parameter, 16-24
IPX static service parameter, 16-14
LSA detail, 15-27
multicast session parameters, 20-12
OSPF link parameters, 15-21
OSPF link state database parameter, 15-26
OSPF link state database search parameters, 15-25
power system statistical parameters, 3-15
type 20 packet propagation
IPX interface, 16-9
type 3 ASE filter
OSPF area parameters, 15-6

U

UDP in datagrams
IP routing global statistics, 12-45
UDP in errors
IP routing global statistics, 12-45
UDP out datagrams
IP routing global statistics, 12-45
UDP port
RADIUS configuration parameters, 4-28
undersize packets
Ethernet interface statistical parameters, 23-5
unknown messages codes encountered
DVMRP global statistical parameters, 12-58
unknown messages received
IGMP global statistic parameters, 12-52
IGMP interface statistical parameters, 12-54
unrecognized packets received
DVMRP interface statistical parameters, 12-60
UPD in no ports
IP routing global statistics, 12-45
update pkt delay
RIP global configuration parameters, 14-3
update timer
RIP global configuration parameters, 14-2
update timestamp
VTP snooping parameters, 6-31
updater identity
VTP snooping parameters, 6-31
updater time stamp
VLAN exchange parameters, 6-31
updates received
IPX SAP interface statistical parameters, 17-9, 18-13
RIP statistical parameters, 14-10, 14-12
upstream interface
DVMRP multicast forwarding cache parameters, 12-69
IGMP local multicast forwarding cache parameters, 12-57
upstream interface is pruned
DVMRP multicast forwarding cache parameters, 12-69
upstream neighbor router address
DVMRP multicast forwarding cache parameters, 12-69
upstream prune information
destination group address, 12-69
DVMRP upstream interface, 12-69
interface is pruned, 12-70
interface type, 12-70
prune expiration time, 12-70
source subnetwork, 12-69

upstream router
DVMRP route table parameters, 12-64
upstream source
IGMP local multicast forwarding cache parameters, 12-57
upstream sources, 12-71
DVMRP multicast forwarding cache parameters, 12-69
use default route
use default route, 16-5
use default route for proxy ARP
IP global configuration parameters, 12-12
use interpacket gap
IPX SAP interface parameters, 18-2
use max packet size
IPX SAP interface parameters, 18-2
user accounts
setting up, 2-19
user port status
event class parameters, 21-6
using dial-up networking
using dial-up networking, 2-34
Using the Web Agent, 20-8
utilization
Ethernet interface statistical parameters, 23-4

V

valid
address forwarding table parameters, 9-17
version
RIP interface parameters, 14-4
view the active alarm table, 21-2, 21-3
Viewing, 12-89
viewing AppleTalk global statistics, 19-23, 19-25
Viewing AppleTalk Global Statistics Using the CLI, 19-25
viewing AppleTalk route table statistics, 19-28
viewing AppleTalk statistics, 19-22
viewing AppleTalk zone table statistics, 19-32
viewing DVMRP downstream routers, 12-67
viewing DVMRP neighbor routers, 12-61
viewing DVMRP upstream routers, 12-65
viewing Ethernet statistics, 23-1, 23-7
viewing event and shutdown logs, 21-11, 21-12
viewing event statistics, 21-7, 21-10, 21-13
viewing IGMP interface statistics, 12-52, 12-53, 12-55
Viewing LDAP Statistics, 12-89
viewing RMON statistics, 23-1, 23-7
viewing switch port parameters, 8-41
Viewing Switch Port Parameters Using the CLI, 8-42
viewing the AppleTalk ARP cache table, 19-29, 19-30
viewing the AppleTalk NBP table, 19-33
Viewing the AppleTalk Route Table, 19-26
viewing the AppleTalk route table, 19-27
viewing the AppleTalk zone table, 19-31
viewing the downloaded image, A-7
viewing the running configuration, 2-39
viewing the script execution log file
script execution log file
viewing, 2-39
viewing the startup configuration, 2-39
Viewing the Status of a TFTP Transfer, 2-45
viewing the TFTP transfer status, 2-45
viewing VRRP statistics, 12-78, 12-80
Virtual LAN binding options
bind to all, 8-36
bind to received, 8-36
static, 8-36
Virtual LANs
flood pruning overview, 1-16
virtual LANs
creating a VLAN, 6-7, 6-10
parameters, 6-9
virtual router
VRRP, 12-74, 12-78
VLAN
add IP interface parameters, 12-8
AppleTalk interface parameters, 19-5
ARP cache search parameters, 12-50
configuring all ports, 6-10
creating, 6-7, 6-10
egress rules, 6-6
filtering, 9-15
forwarding rules, 6-5
IPX interface, 16-8
Layer 3 route cache search parameters, 22-8
LGMP client display per VLAN parameters, 20-32
LGMP server display per VLAN parameters, 20-27
multicast session parameters, 20-12
operation, 6-4
router port display parameters, 20-9
search, 9-15
static multicast session configuration parameters, 20-17
static multicast session parameters, 20-18
VLAN association
address table instance parameters, 9-11
VLAN binding
Cisco Catalyst 5000 parameters, 8-34
switch port configuration parameters, 8-31
VLAN classification
switch port parameter, 8-41
VLAN configuration parameters

auto increment HT size, 6-10
initial hash table size, 6-10
VLAN considerations, 6-7
VLAN exchange parameters
updater timestamp, 6-31
VLAN operation
ingress rules, 6-4
VLAN Operational Rules, 8-7
VLAN spanning tree bridge level parameters
bridge forward delay, 7-12
bridge hello time, 7-12
bridge max age, 7-12
forward delay, 7-12
hello time, 7-12
max age, 7-12
mode, 7-11
priority, 7-12
VLAN switch port table parameters
binding type, 6-12
frame format, 6-12
port, 6-11
port name, 6-11
VLAN trunking, definition, 1-10
VR ID
VRRP configuration parameters, 12-76
VRRP statistical parameters, 12-79
VRRP
add IP interface parameters, 12-10
enabling, 12-72
IP global configuration parameters, 12-12
VRRP configuration overview, 12-72
VRRP configuration parameters
addr owner override, 12-78
advertisement timer, 12-77
auth key, 12-77
interface, 12-76
IP address, 12-76
priority, 12-76
VR ID, 12-76
VRRP overview
configuring, 12-72
VRRP router
creating, 12-74
VRRP statistical parameters
advertisements received, 12-80
advertisements sent, 12-80
bad advertisements received, 12-80
interface, 12-79
IP address, 12-79
state, 12-80
time of state change, 12-80
times this VR became master, 12-80
VR ID, 12-79
VRRP statistics
viewing, 12-78, 12-80
VRRP virtual router
creating, 12-74, 12-78
VTP snooping
Cisco Catalyst 5000 parameters, 8-35
switch port configuration parameters, 8-31
VTP snooping parameters
configuration revision number, 6-31
domain name, 6-31
update timestamp, 6-31
updater identity, 6-31
VTP snooping state, 6-31
VTP snooping state
VTP snooping parameters, 6-31

W

warning
protocol event log settings, 21-9
Web Agent
logging into, 2-9
setting up user accounts, 2-19
week
summer time hours configuration, 3-8

X

Xon/Xoff
flow control, 2-28

Z

Zip Ext Reply Rx
AppleTalk global statistics, 19-25
Zip Ext Reply Tx
AppleTalk global statistics, 19-25
Zip GNI Rq Rx
AppleTalk global statistics, 19-25
Zip GNI Rq Tx
AppleTalk global statistics, 19-25
Zip GNI Rsp Rx
AppleTalk global statistics, 19-25
Zip GNI Rsp Tx
AppleTalk global statistics, 19-25
Zip Query Rx
AppleTalk global statistics, 19-25
Zip Query Tx
AppleTalk global statistics, 19-25
Zip Reply Rx
AppleTalk global statistics, 19-25
Zip Reply Tx
AppleTalk global statistics, 19-25
zone

AppleTalk interface parameters, 19-6
AppleTalk NBP table parameters, 19-33
AppleTalk static route, 19-11
zone filter
adding interfaces, 19-21, 19-22
creating, 19-19
deleting interfaces, 19-21, 19-22
editing, 19-20
zones
AppleTalk route table statistics, 19-28